Appendix A

Conducting Literature Reviews and Finding Information

Social Research Foundations—How to Conduct a Literature Review

Case Study: Arrest and Domestic Violence

How do we find prior research on questions of interest? You may already know some of the relevant material from prior coursework or your independent reading, but that would not be enough. When you are about to launch an investigation of a new research question, you must apply a very different standard than when you are studying for a test or just seeking to learn about domestic violence. You need to find reports of previous investigations that sought to answer the same research question that you wish to answer, not just those that were about a similar topic. If there are no prior studies of exactly the same research question on which you wish to focus, you should seek to find reports from investigations of very similar research questions. Once you have located reports from prior research similar to the research you wish to conduct, you may expand your search to include investigations about related topics or studies that used similar methods.

Sometimes, you will find that someone else has already reviewed the literature on your research question in a special review article or book chapter. For example, Chalk and Garner (2001) published an excellent review of the research on arrest and domestic violence in the journal *New Directions for Evaluation*. Most of the research articles you find will include a literature review. These reviews can help a lot, but they are no substitute for reviewing the literature yourself. No one but you can decide what is relevant for your research question and the research circumstances you will be

facing: the setting you will study, the timing of your study, the new issues that you want to include in your study, and your specific methods. And you cannot depend on any published research review for information on the most recent work. New research results about many questions appear continually in scholarly journals and books, in research reports from government agencies and other organizations, and on website all over the world; you will need to check for new research such as this yourself.

Finding Information

Conducting a thorough search of the research literature and then reviewing critically what you have found is an essential foundation for any research project. Fortunately, much of this information can be identified online, without leaving your desktop, and an increasing number of published journal articles can be downloaded directly to your own computer (depending on your particular access privileges). But just because there is a lot available online does not mean that you need to find it all. Keep in mind that your goal is to find reports of prior research investigations, which means that you should focus on scholarly journals that choose articles for publication after they have been reviewed by other social scientists: "refereed" journals. Newspaper and magazine articles just will not do, although you may find some that raise important issues or even that summarize social science research investigations.

Every year, the Web offers more and more useful material, including indexes of the published research literature. You may find copies of particular rating scales, reports from research in progress, papers that have been presented at professional conferences, and online discussions of related topics. This section reviews the basic procedures for finding relevant research information in both the published literature and on the Web, but keep in mind that the primary goal is to identify research articles published in refereed journals.

Searching the Literature

The social science literature should be consulted at the beginning and end of an investigation. Even while an investigation is in progress, consultations with the literature may help to resolve methodological problems or facilitate supplementary explorations. As with any part of the research process, the method you use will affect the quality of your results. You should try to ensure that your search method includes each of the steps below.

Specify your research question. Your research question should be neither so broad that hundreds of articles are judged relevant nor so narrow that you miss important literature. "Is informal social control effective?" is probably too broad. "Does informal social control reduce rates of burglary in large cities?" is probably too narrow. "Is informal social control more effective in reducing crime rates than policing?" provides about the right level of specificity.

Identify appropriate bibliographic databases to search. *Criminal Justice Abstracts* and *Sociological Abstracts* may meet many of your needs, but if you are studying a question about medical consequences or other health issues, you should also search in *Medline*, the database for searching the medical literature. If your focus is on mental health, you will also want to include a search in the online *Psychological Abstracts* database, *PsycINFO*, or the version that also contains the full text of articles since 1985, *PsycARTICLES*. To find articles that refer to a previous publication, such as Sherman and Berk's (1984) study of the police response to domestic violence, the *Social Science Citation Index* would be helpful. In addition, the search engine Google now offers anyone with Web access Google Scholar (which indexes and searches the full text of selected journals) and Google Print (which digitizes and searches the full text of the books that are owned by selected research libraries). (At the time this book went to press, the Google Print project was on hold due to copyright concerns raised by some publishers, while the search engine and directory Yahoo! was starting a similar venture that focused only on older books that are no longer covered by copyright law; Hafner, 2005: C1.)

Choose a search technology. For most purposes, an online bibliographic database that references the published journal literature will be all you need. However, searches for unpopular topics or very recent literature may require that you also search Websites or bibliographies of relevant books.

Create a tentative list of search terms. List the parts and subparts of your research question and any related issues that you think are important: "informal social control," "policing," "influences on crime rates," and perhaps "community cohesion and crime." List the authors of relevant studies. Specify the most important journals that deal with your topic.

Narrow your search. The sheer number of references you find can be a problem. For example, searching for "social capital" resulted in 2,293 citations in *Sociological Abstracts*. Depending on the database you are working with and the purposes of your search, you may want to limit your search to English-language publications, to journal articles rather than conference papers or dissertations (both of which are more difficult to acquire), and to materials published in recent years.

Refine your search. Learn as you go. If your search yields too many citations, try specifying the search terms more precisely. If you have not found much literature, try using more general terms. Whatever terms you search first, do not consider your search complete until you have tried several different approaches and have seen how many articles you find. A search for "domestic violence" in *Sociological Abstracts* on September 11, 2005, yielded 1,569 hits; adding "effects" or "influences" as required search terms dropped the number of hits to 370.

Use Boolean search logic. It is often a good idea to narrow down your search by requiring that abstracts contain combinations of words or phrases that include more of the specific details of your research question. Using the Boolean connector "and" allows you to do this, whereas using the connector "or" allows you to find abstracts containing different words that mean the same thing. Exhibit A.1 provides an example.

Use appropriate subject descriptors. Once you have found an article that you consider to be appropriate, take a look at the "descriptors" field in the citation (see Exhibit A.2). You can then redo your search after requiring that the articles be classified with some or all of these descriptor terms.

Check the results. Read the titles and abstracts you have found, and identify the articles that appear to be most relevant. If possible, click on these article titles and generate a list of their references. See if you find more articles that are relevant to your research question but that you have missed so far. You will be surprised (I always am) at how many important articles your initial online search missed.

Read the articles. Now, it is time to find the full text of the articles of interest. If you are lucky, some of the journals you need will be available to patrons of your library in online versions, and you will be able to link to the full text just by clicking on a "full text" link. But many journals, specific issues of some journals, or both will be available only in print; in this case, you will have to find them in your library or order a copy through interlibrary loan.

Refer to a good book for even more specific guidance. Fink's (2004) *Conducting Research Literature Reviews: From the Internet to Paper* is an excellent guide.

You may be tempted to write up a "review" of the literature based on reading the abstracts or using only those articles available online, but you will be selling yourself short. Many crucial details about methods, findings, and theoretical implications will be found only in the body of the article, and many important articles will not be available online. To understand, critique, and really learn from previous research studies, you must read the important articles, no matter how you have to retrieve them.

Exhibit A.1 Use of Boolean Connectors in a Literature Search

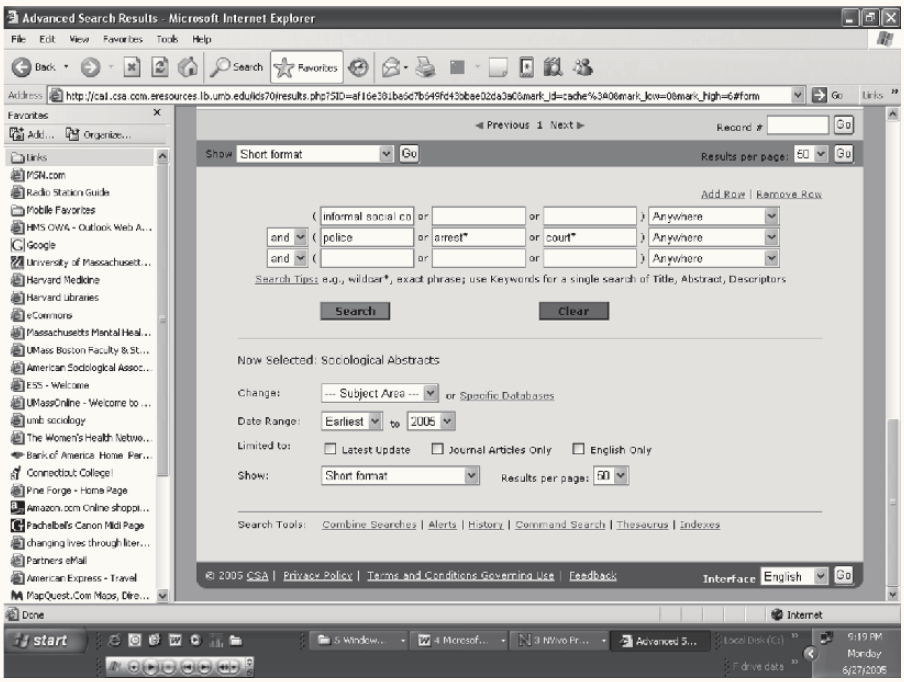

Exhibit A.2 Checking Standard Subject Matter Descriptors

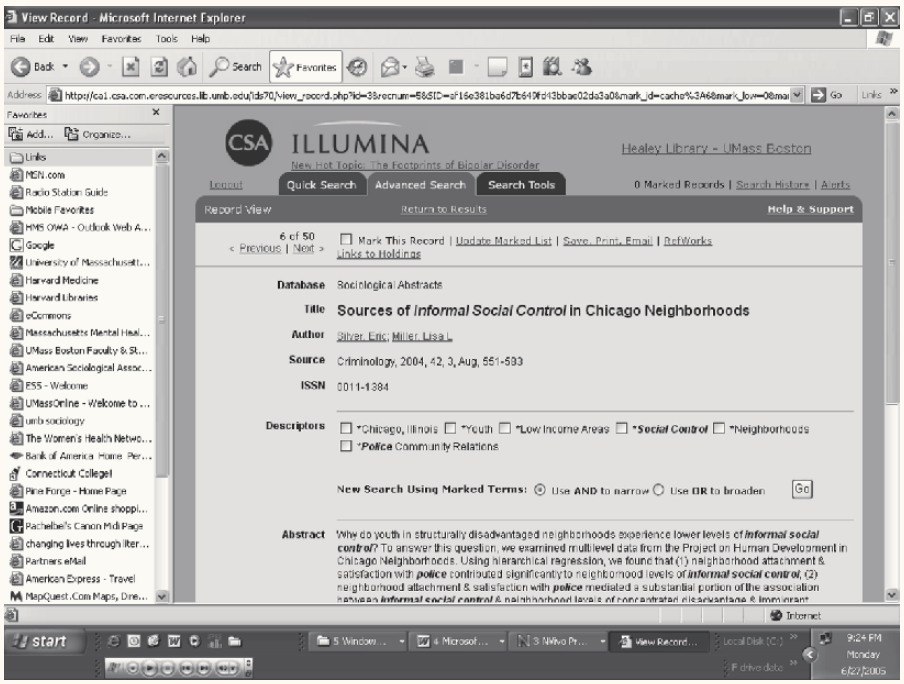

If you have done your job well, you will now have more than enough literature as background for your own research, unless it is on a very obscure topic (see Exhibit A.3). (Of course, ultimately your search will be limited by the library holdings you have access to and by the time you have to order or find copies of journal articles, conference papers, and perhaps dissertations that you cannot obtain online.) At this point, your main concern is to construct a coherent framework in which to develop your research question, drawing as many lessons as you can from previous research. You may use the literature to identify a useful theory and hypotheses to be reexamined, to find inadequately studied specific research questions, to explicate the disputes about your research question, to summarize the major findings of prior research, and to suggest appropriate methods of investigation.

Be sure to take notes on each article you read, organizing your notes into standard sections: Theory, Methods, Findings, Conclusions. In any case, write the literature review so that it contributes to your study in some concrete way; do not feel compelled to discuss an article just because you have read it. Be judicious. You are conducting only one study of one issue; it will only obscure the value of your study if you try to relate it to every tangential point in related research.

Don't think of searching the literature as a one-time-only venture, something that you leave behind as you move on to your "real" research. You may encounter new questions or unanticipated problems as you conduct your research or as you burrow deeper into the literature. Searching the literature again to determine what others have found in response to these questions or what steps they have taken to resolve these problems can yield substantial improvements in your own research. There is so much literature on so many topics that often it is not possible to figure out in advance every subject you should search the literature for or what type of search would be most beneficial.

Another reason to make searching the literature an ongoing project is that the literature is always growing. During the course of one research study, whether it takes only one semester or several years, new findings will be published and relevant questions will be debated. Staying attuned to the literature and checking it at least when you are writing up your findings may save your study from being outdated.

Exhibit A.3 **A Search in *Sociological Abstracts* on Informal Control**

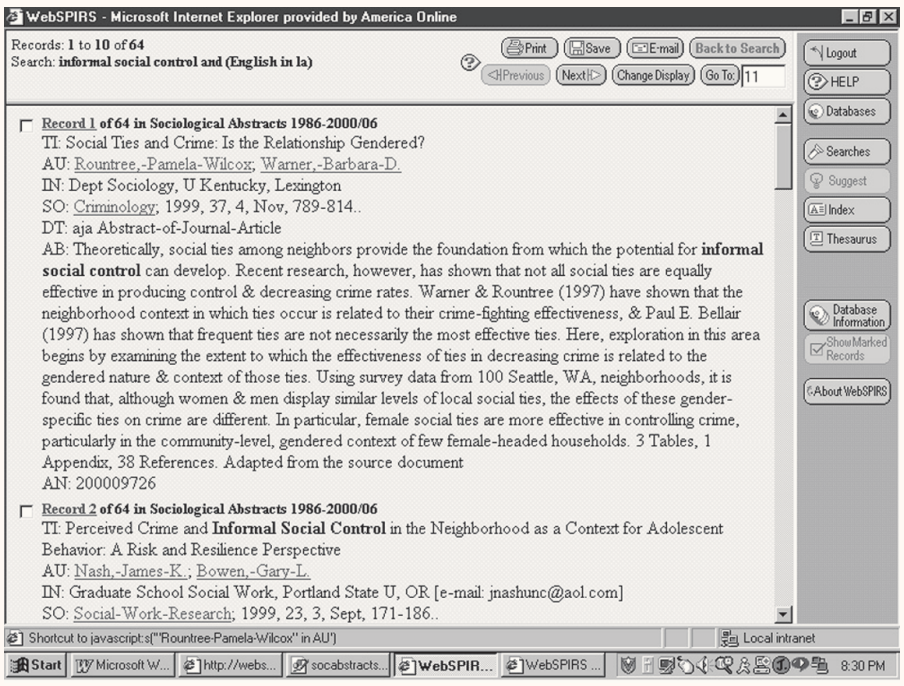

Searching the Web

The World Wide Web provides access to vast amounts of information of many different sorts (Ó Dochartaigh, 2002). You can search the holdings of other libraries and download the complete text of government reports, some conference papers, and newspaper articles. You can find policies of local governments, descriptions of individual social scientists and particular research projects, and postings of advocacy groups. It is also hard to avoid finding a lot of information in which you have no interest, such as commercial advertisements, third-grade homework assignments, or college course syllabi. In 1999, there were already about 800 million publicly available pages of information on the Web (Davis 1999). Today, there may be as many as 15 billion pages on the Web (Novak, 2003).

After you are connected to the Web with a browser such as Microsoft Internet Explorer or Mozilla Firefox, you can use three basic strategies for finding information: direct addressing (i.e., typing in the address, or uniform resource locator [URL], of a specific site); browsing (i.e., reviewing online lists of Websites); and searching (i.e., Google is currently the most popular search engine for searching the Web). For some purposes, you will need to use only one strategy; for other purposes, you will want to use all three.

Exhibit A.4 illustrates the first problem that you may encounter when searching the Web: the sheer quantity of resources that are available. It is a much bigger problem than when searching bibliographic databases. On the Web, less is usually more. Limit your inspection of Websites to the first few pages that turn up in your list (they are ranked by relevance). See what those first pages contain, and then try to narrow your search by including some additional terms. Putting quotation marks around a phrase that you want to search will also help to limit your search; for example, searching for "informal social control" on Google (on September 11, 2005) produced 31,100 sites, compared to the roughly 15,500,000 sites retrieved when we omitted the quotes wherein Google searched "informal" and "social" and "control."

Exhibit A.4 **A Search in *Sociological Abstracts* on Informal Control**

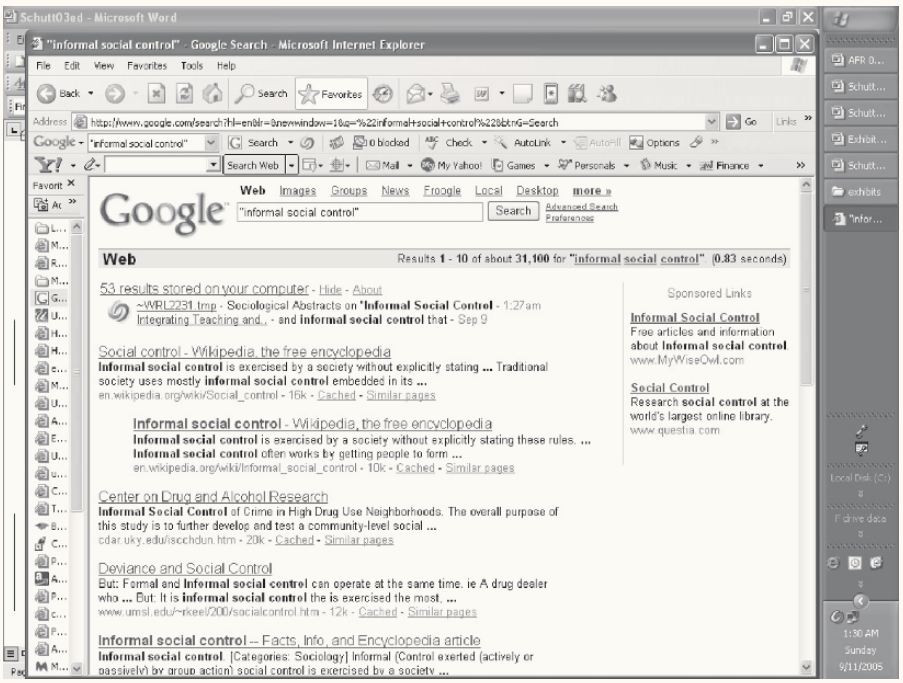

Remember the following warnings when you conduct searches on the Web:

- *Clarify your goals.* Before you begin the search, jot down the terms that you think you need to search for as well as a statement of what you want to accomplish with your search. This will help to ensure that you have a sense of what to look for and what to ignore.
- *Quality is not guaranteed.* Anyone can post almost anything, so the accuracy and adequacy of the information you find are always suspect. There is no journal editor or librarian to evaluate quality and relevance.
- *Anticipate change.* Websites that are not maintained by stable organizations can come and go very quickly. Any search will result in attempts to link to some URLs that no longer exist.
- *One size does not fit all.* Different search engines use different procedures for indexing Websites. Some attempt to be all-inclusive, whereas others aim to be selective. As a result, you can get different results from different search engines (e.g., Google or Yahoo!) even though you are searching for the same terms.
- *Be concerned about generalizability.* You might be tempted to characterize police department policies by summarizing the documents you find at police department Websites. But how many police departments are there? How many have posted their policies on the Web? Are these policies representative of all police departments? To answer all these questions, you would have to conduct a research project just on the Websites themselves.
- *Evaluate the sites.* There is a lot of stuff out there, so how do you know what is good? Some Websites contain excellent advice and pointers on how to differentiate the good from the bad.
- *Avoid Web addiction.* Another danger of the enormous amount of information available on the Web is that one search will lead to another and to another and so on. There are always more possibilities to explore and one more interesting source to check. Establish boundaries of time and effort to avoid the risk of losing all sense of proportion.
- *Cite your sources.* Using text or images from Web sources without attribution is plagiarism. It is the same as copying someone else's work from a book or article and pretending that it is your own. Record the Web address (URL), the name of the information provider, and the date on which you obtain material from the site. Include this information in a footnote to the material that you use in a paper.

Reviewing Research

Effective review of the prior research you find is an essential step in building the foundation for new research. You must assess carefully the quality of each research study, consider the implications of each article for your own plans, and expand your thinking about your research question to take account of new perspectives and alternative arguments. It is through reviewing the literature and using it to extend and sharpen your own ideas and methods that you become a part of the social science community. Instead of being just one individual studying an issue that interests you, you are building on an ever-growing body of knowledge that is being constructed by the entire community of scholars.

The research information you find on various Web sites comes in a wide range of formats and represents a variety of sources. *Caveat emptor* (buyer beware) is the watchword when you search the Web; following review guidelines such as those we have listed will minimize, but not eliminate, the risk of being led astray. By contrast, the published scholarly journal literature that you find in databases such as *Sociological Abstracts* and *Psychological Abstracts* follows a much more standard format and has been subject to a careful review process. There is some variability in the contents of these databases: some journals publish book reviews, comments on prior articles, dissertation abstracts, and conference papers. However, most literature you will find on a research topic in these databases represents peer-reviewed articles reporting analyses of data collected in a research project. These are the sources on which you should focus. This section

concentrates on the procedures you should use for reviewing these articles. These procedures also can be applied to reviews of research monographs: books that provide much more information from a research project than that contained in a journal article.

Reviewing the literature is really a two-stage process. In the first stage, you must assess each article separately. This assessment should follow a standard format such as that represented by the "Questions to Ask About a Research Article" in Appendix B. However, you should keep in mind that you cannot adequately understand a research study if you just treat it as a series of discrete steps, involving a marriage of convenience among separate techniques. Any research project is an integrated whole, so you must be concerned with how each component of the research design influenced the others, for example, how the measurement approach might have affected the causal validity of the researcher's conclusions and how the sampling strategy might have altered the quality of measures.

The second stage of the review process is to assess the implications of the entire set of articles (and other materials) for the relevant aspects of your research question and procedures and then to write an integrated review that highlights these implications. Although you can find literature reviews that consist simply of assessments of one published article after another—that never get beyond stage one in the review process—your understanding of the literature and the quality of your own work will be much improved if you make the effort to write an integrated review.

In the next two sections, we will show how you might answer many of the questions in Appendix B as we review a research article about domestic violence. We will then show how the review of a single article can be used within an integrated review of the body of prior research on this research question. Because at this early point in the text you will not be familiar with all the terminology used in the article review, you might want to read through the more elaborate article review in Appendix C later in the course.

A Single-Article Review: Formal and Informal Deterrents to Domestic Violence

Anthony Pate and Edwin Hamilton at the National Police Foundation designed one of the studies funded by the U.S. Department of Justice to replicate the Minneapolis Domestic Violence Experiment. This section examines the article that resulted from that replication, which was published in the *American Sociological Review* (Pate & Hamilton, 1992). The numbers in brackets refer to the article review questions in Appendix B.

The Research Question. Like Sherman and Berk's (1984) original Minneapolis study, Pate and Hamilton's (1992) Metro-Dade spouse assault experiment sought to test the deterrent effect of arrest in domestic violence cases, but with an additional focus on the role of informal social control [1]. The purpose of the study was explanatory because the goal was to explain variation in the propensity to commit spouse abuse [2]. Deterrence theory provided the theoretical framework for the study, but this framework was broadened to include the proposition by Williams and Hawkins (1986) that informal sanctions such as stigma and the loss of valued relationships augment the effect of formal sanctions such as arrest [3]. Pate and Hamilton's literature review referred, appropriately, to the original Sherman and Berk research, to the other studies that attempted to replicate the original findings, and to research on informal social control [4].

There is no explicit discussion of ethical guidelines in the article, although reference is made to a more complete unpublished report [6]. Clearly, important ethical issues had to be considered, given the experimental intervention in the police response to serious assaults, but the adherence to standard criminal justice procedures suggests attention to the welfare of victims as well as the rights of suspects. We will consider these issues in more detail later in this chapter.

The Research Design. Developed as a follow-up to the original Minneapolis experiment, the Metro-Dade experiment exemplified the guidelines for scientific research that were presented in Chapter 2 [5]. It was designed systematically, with careful attention to specification of terms and clarification of assumptions, and focused on the possibility of different outcomes rather than certainty about one preferred outcome. The major concepts in the study, formal and informal deterrence, were defined clearly [7] and then measured with straightforward indicators: arrest or nonarrest for formal deterrence and marital status and employment status for informal deterrence. However, the

specific measurement procedures for marital and employment status were not discussed, and no attempt was made to determine whether they captured adequately the concept of informal social control.

Three hypotheses were stated and also related to the larger theoretical framework and prior research [8]. The study design focused on the behavior of individuals [11] and collected data over time, including records indicating subsequent assault up to 6 months after the initial arrest [12]. The project's experimental design was used appropriately to test for the causal effect of arrest on recidivism [13]. The research project involved all eligible cases, rather than a sample of cases, but there were a number of eligibility criteria that narrowed the ability to generalize these results to the entire population of domestic assault cases in the Metro-Dade area or elsewhere [14]. There is a brief discussion ofthe 92 eligible cases that were not given the treatment to which they were assigned, but it does not clarify the reasons for the misassignment [15].

The Research Findings and Conclusion. Pate and Hamilton's (1992) analysis of the Metro-Dade experiment was motivated by concern with effect of social context because the replications in other cities of the original Minneapolis domestic violence experiment had not had consistent results [19]. Their analysis gave strong support to the expectation that informal social control processes are important: As they had hypothesized, arrest had a deterrent effect on suspects who were employed but not on those who were unemployed (see Exhibit A.5). However, marital status had no such effect [20]. The subsequent discussion of these findings gives no attention to the implications of the lack of support for the effect of marital status [21], but the study represents an important improvement over earlier research that had not examined informal sanctions [22]. The need for additional research is highlighted, and the importance of the findings for social policy are discussed: Pate and Hamilton suggest that their finding that arrest deters only those who have something to lose (e.g., a job) must be taken into account when policing policies are established [23].

Overall, the Pate and Hamilton (1992) study represents an important contribution to understanding how informal social control processes influence the effectiveness of formal sanctions such as arrest. Although the use of a population of actual spouse assault cases precluded the use of very sophisticated measures of informal social control, the experimental design of the study and the researchers' ability to interpret the results in the context of several comparable experiments distinguishes this research as exceptionally worthwhile. It is not hard to understand why these studies continue to stimulate further research and ongoing policy discussions.

Exhibit A.5 **Percentage of Suspects With a Subsequent Assault by Employment Status and Arrest Status**

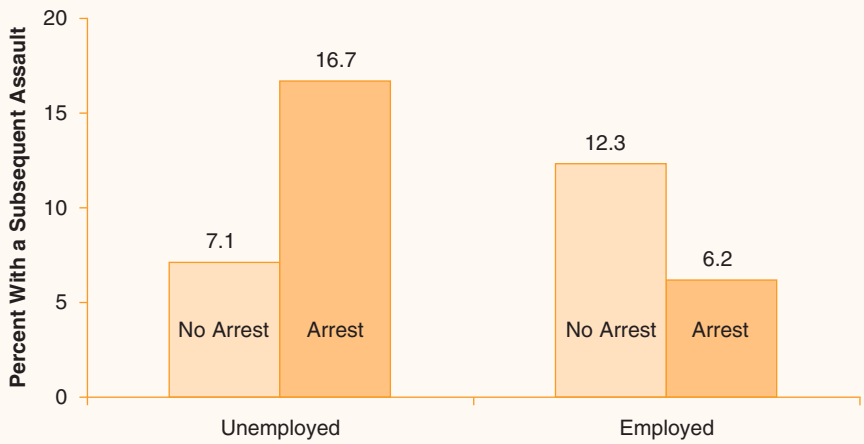

An Integrated Literature Review: When Does Arrest Matter?

The goal of the second stage of the literature review process is to integrate the results of your separate article reviews and develop an overall assessment of the implications of prior research. The integrated literature review should accomplish three goals: (1) summarize prior research, (2) critique prior research, and (3) present pertinent conclusions (Hart, 1998: 186–187).

Summarize prior research. Your summary of prior research must focus on the particular research questions that you will address, but you also may need to provide some more general background. Hoyle and Sanders (2000: 14) begin their *British Journal of Criminology* research article about mandatory arrest policies in domestic violence cases with what they term a "provocative" question: What is the point of making it a crime for men to assault their female partners and ex-partners? They then review the different theories and supporting research that has justified different police policies: the "victim choice" position, the "pro-arrest" position, and the "victim empowerment" position. Finally, they review the research on the "controlling behaviors" of men that frames the specific research question on which they focus: how victims view the value of criminal justice interventions in their own cases (p. 15).

Ask yourself three questions about your summary of the literature:

1. *Have you been selective?* If there have been more than a few prior investigations of your research question, you will need to narrow your focus to the most relevant and highest quality studies. Do not cite a large number of prior articles "just because they are there."
2. *Is the research up-to-date?* Be sure to include the most recent research, not just the "classic" studies.
3. *Have you used direct quotes sparingly?* To focus your literature review, you need to express the key points from prior research in your own words. Use direct quotes only when they are essential for making an important point (Pyrczak, 2005: 51–59).

Critique prior research. Evaluate the strengths and weaknesses of the prior research. In addition to all the points you develop as you answer the Article Review Questions in Appendix B, you should also select articles for review that reflect work published in peer-reviewed journals and written by credible authors who have been funded by reputable sources. Consider the following questions as you decide how much weight to give each article:

1. *How was the report reviewed prior to its publication or release?* Articles published in academic journals go through a rigorous review process, usually involving careful criticism and revision. Top refereed journals may accept only 10% of submitted articles, so they can be very selective. Dissertations go through a lengthy process of criticism and revision by a few members of the dissertation writer's home institution. A report released directly by a research organization is likely to have had only a limited review, although some research organizations maintain a rigorous internal review process. Papers presented at professional meetings may have had little prior review. Needless to say, more confidence can be placed in research results that have been subject to a more rigorous review.
2. *What is the author's reputation?* Reports by an author or team of authors who have published other work on the research question should be given somewhat greater credibility at the outset.
3. *Who funded and sponsored the research?* Major federal funding agencies and private foundations fund only research proposals that have been evaluated carefully and ranked highly by a panel of experts. They also often monitor closely the progress of the research. This does not guarantee that every such project report is good, but it goes a long way toward ensuring some worthwhile products. On the other hand, research that is funded by organizations that have a preference for a particular outcome should be given particularly close scrutiny (Locke, Silverman, & Spirduso, 1998: 37–44).

Present pertinent conclusions. Do not leave the reader guessing about the implications of the prior research for your own investigation. Present the conclusions you draw from the research you have reviewed. As you do so, follow several simple guidelines:

- Distinguish clearly your own opinion of prior research from conclusions of the authors of the articles you have reviewed.
- Make it clear when your own approach is based on the theoretical framework you are using rather than on the results of prior research.
- Acknowledge the potential limitations of any empirical research project. Do not emphasize problems in prior research that you cannot avoid either (Pyrczak, 2005: 53–56).
- Explain how the unanswered questions raised by prior research or the limitations of methods used in prior research make it important for you to conduct your own investigation (Fink, 2004: 190–192).

A good example of how to conclude an integrated literature review is provided by an article based on the replication in Milwaukee of the Minneapolis Domestic Violence Experiment. For this article, Paternoster et al. (1997) sought to determine whether police officers' use of fair procedures when arresting assault suspects would lessen the rate of subsequent domestic violence. Paternoster et al. conclude that there has been a major gap in the prior literature: "Even at the end of some seven experiments and millions of dollars, then, there is a great deal of ambiguity surrounding the question of how arrest impacts future spouse assault" (p. 164). Specifically, they note that each of the seven experiments focused on the effect of arrest itself but ignored the possibility that "particular kinds of police procedure might inhibit the recurrence of spouse assault" (p. 165).

So Paternoster et al. (1997) ground their new analysis in additional literature on procedural justice and conclude that their new analysis will be "the first study to examine the effect of fairness judgments regarding a punitive criminal sanction (arrest) on serious criminal behavior (assaulting one's partner)" (p. 172).

2 Searching the Web

To find useful information on the Web, you have to be even more vigilant than when you search the literature directly. With billions of webpages on the Internet, there is no limit to the amount of time you can squander and the volume of useless junk you can find as you conduct your research on the Web. However, we can share with you some good ways to avoid the biggest pitfalls.

Direct Addressing

Knowing the exact address (i.e., URL) of a useful Web site is the most efficient way to find a resource on the Web.

Professional Organizations

- American Society of Criminology (http://www.asc41.com)
- American Sociological Association (http://www.asanet.org)
- American Psychological Association (http://www.apa.org)

Government Sites

- U.S. Office of Justice Programs (http://www.ojp.usdoj.gov)
- U.S. Bureau of the Census (http://www.census.gov)

Journals and Newspapers

- *Annual Review of Sociology* (http://www.annualreviews.org)
- *The New York Times* (http://www.nytimes.com)

Bibliographic Formats for Citing Electronic Information

- Electronic reference formats suggested by the American Psychological Association (http://www.apastyle.org/elecref.html)
- Karla Tonella's Guide to Citation Style Guides (http://bailiwick.lib.uiowa.edu/journalism/cite.html) contains more than a dozen links to online style guides
- Style Sheets for Citing Resources (print and electronic) (http://www.lib.berkeley.edu/TeachingLib/Guides/Internet/Style.html)

When you find Web sites that you expect you will return to often, you can save their addresses as "bookmarks" or "favorites" in your Web browser. However, since these can very quickly multiply, you should try to be selective.

Browsing Subject Directories

Subject directories (also called guides, indexes, or clearinghouses) contain links to other Web resources that are organized by subject. They vary in quality and authoritativeness, but a good one can be invaluable to your research and save you much time. The main advantage to using subject directories is that they contain links to resources that have been selected, evaluated, and organized by human beings and thus present a much more manageable number of resources. If the person managing the guide is an expert in the field of concern, or just a careful and methodological evaluator of Web resources, the guide can help you to identify good sites that contain useful and trustworthy information, and you can avoid wading through thousands of "hits" and evaluating all the sites yourself.

There are general and specialized directories. The following are three examples of general directories:

- Yahoo! (http://www.yahoo.com) is often mistaken for a search engine, but it is actually a subject directory, and a monster one at that. It also functions as a portal or a gateway for a collection of resources that can be customized by the user. Unlike search engines, when you search Yahoo!, you are not searching across the Web but rather just within the Web pages that Yahoo! has cataloged. Yahoo! has a subject directory for the social sciences with more specific listings, including one for social work (http://dir.yahoo.com/social_science/social_work/). Yahoo! also links to versions of its site in about 20 countries, which would be good to go to when conducting extensive research on one of those countries (http:world.yahoo.com/).
- Open Directory (http://dmoz.org) is the largest Web directory with four million sites (Hock, 2007), and unlike Yahoo!, it is not a portal. In fact, other directories and search engines such as Yahoo! and Google use it. It has 16 top-level categories, including Social Sciences.
- Librarians' Index to the Internet (http://lii.org) is a small and highly selective Web directory produced by the Library of California.

Many other Internet subject directories are maintained by academic departments, professional organizations, and individuals. It is often hard to determine whether a particular subject directory such as this is up-to-date and reasonably comprehensive, but you can have some confidence in subject directories published by universities or government agencies. *The Internet Research Handbook* is an excellent source for more information on subject directories (Ó Dochartaigh, 2002).

Search Engines

Search engines are powerful Internet tools. It is already impossible to imagine life without them. The biggest problem is the huge number of results that come back to you. If the number of results is still unmanageable, you can try a title search. Exhibit A.6 shows the results of typing the following into the Google search box: ti: "informal social control." This search will retrieve those pages that have that phrase in their title as opposed to anywhere in the page. This practice usually results in a dramatically smaller yield of results. If you are looking for graphical information such as a graph or a chart, you can limit your search to those pages that contain an image. On Google, this just requires clicking on the "Images" link located above the search box.

There are many search engines, and none will give you identical results when you use them to search the Web. Different search engines use different strategies to find Web sites and offer somewhat different search options for users. Due to the enormous size of the Web and its constantly changing content, it simply is not possible to identify one search engine that will give you completely up-to-date and comprehensive results. You can find the latest information about search engines at http://searchenginewatch.com. Hock's (2004) *The Extreme Searcher's Internet Handbook* contains a wealth of information on specific search engines. Although there are many search engines, you may find the following to be particularly useful for general searching:

- Google (http://www.google.com) has become the leading search engine for many users in recent years. Its coverage is relatively comprehensive, and it does a good job of ranking search results by their relevancy (based on the terms in your search request). Google also allows you to focus your search just on images, discussions, or directories.
- AlltheWeb (http://www.alltheweb.com) is a more recent comprehensive search engine that also does a good job of relevancy ranking and allows searches restricted to images and so on.

Exhibit A.6 The Results of a Google Title Search

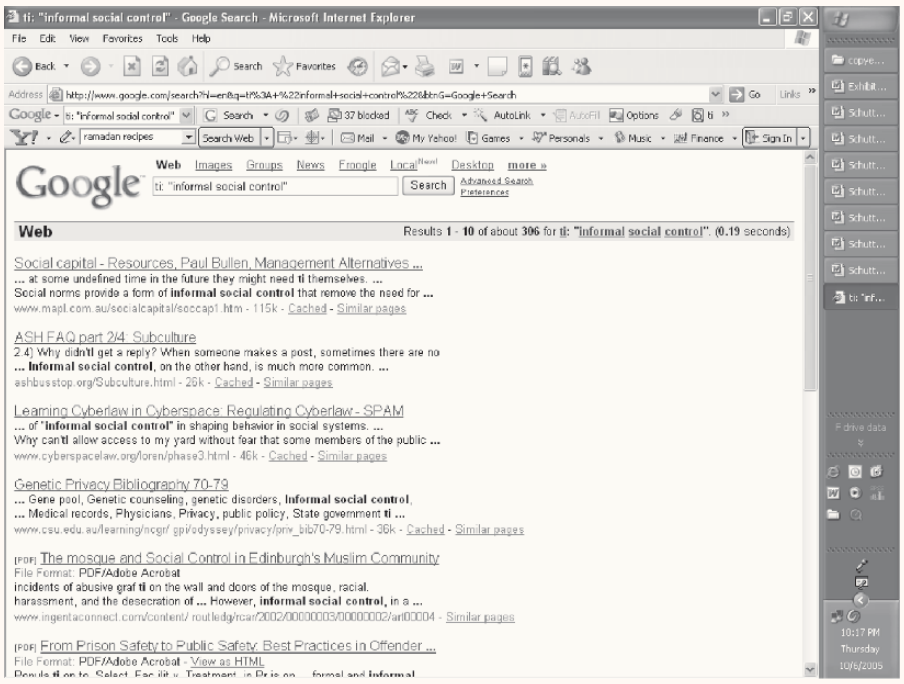

- Microsoft's search engine (http://search.msn.com) adds a unique feature: Editors review and pick the most popular sites. As a result, your search request may result in a Popular Topics list that can help you to focus your search.
- Teoma (http://teoma.com) is one of the newest search engines and has a unique Resources section that links users to specialized directories.

In conclusion, use the appropriate tool for your searches. Do not use a search engine in place of searching literature that is indexed in tools such as *Sociological Abstracts*. Bookmark the key sites that you find in your area of interest. Become familiar with subject directories that cover your areas of interest, and look there before going to a search engine. And when you do use a search engine, take a moment to learn about how it works and what steps you should take to get the best results in the least amount of time.

Appendix B

How to Read a Research Article

The discussions of research articles throughout the text may provide all the guidance you need to read and critique research on your own. But reading about an article in bits and pieces to learn about particular methodologies is not quite the same as reading an article in its entirety to learn what the research discovered. The goal of this appendix is to walk you through an entire research article, answering the review questions introduced in Appendix A. Of course, this is only one article, and our "walk" will take different turns from one taken by a review of other articles, but after this review, you should feel more confident when reading other research articles on your own.

For this example, we will use an article by Yi-Fen Lu, Yi-Chun Lu, Ling Ren, and Ineke Marshall that provides a test of self-control theory with a sample of Chinese adolescents. It contributes to our understanding of the power and efficacy of self-control theory. This theory specifies that the impact of self-control variables should not be affected by culture, but very few studies exist that have tested the theory in non-Western contexts generally or in China specifically. Moreover, the article is published in a reputable criminological journal, the *Journal of Contemporary Criminal Justice*, indicating the article makes an important contribution to what is known about the causes and correlates of delinquent behaviors.

The questions you need to answer for any article or research monograph are as follows:

1. What is the basic research question or problem? Try to state it in just one sentence. (Chapter 2)
2. Is the purpose of the study explanatory, evaluative, exploratory, or descriptive? Did the study have more than one purpose? (Chapter 1)
3. Was the theoretical framework presented? What was it? Did it seem appropriate for the research question addressed? Can you think of a different theoretical perspective that might have been used? (Chapter 2)
4. What prior literature was reviewed? Was it relevant to the research problem? To the theoretical framework? Does the literature review appear to be adequate? Are you aware of (or can you locate) any important omitted studies? (Chapter 2)

5. How well did the study live up to the guidelines for science? Do you need additional information in any areas to evaluate the study? To replicate it? (Chapter 2)
6. Did the study seem consistent with current ethical standards? Were any trade-offs made between different ethical guidelines? Was an appropriate balance struck between adherence to ethical standards and use of the most rigorous scientific practices? (Chapter 2 and in each methods chapter)
7. What were the major concepts in the research? How, and how clearly, were they defined? Were some concepts treated as unidimensional that you think might best be thought of as multidimensional? (Chapter 3)
8. Were any hypotheses stated? Were these hypotheses justified adequately in terms of the theoretical framework? In terms of prior research? (Chapter 2)
9. What were the independent and dependent variables in the hypothesis or hypotheses? Did these variables reflect the theoretical concepts as intended? What direction of association was hypothesized? Were any other variables identified as potentially important? (Chapter 2)
10. Did the instruments used—the measures of the variables—seem valid and reliable? How did the author attempt to establish this? Could any more have been done in the study to establish measurement validity? (Chapter 3)
11. What were the units of analysis? Were they appropriate for the research question? If some groups were the units of analysis, were any statements made at any point that are open to the ecological fallacy? If individuals were the units of analysis, were any statements made at any point that suggest reductionist reasoning? (Chapter 4)
12. Was the study design cross-sectional or longitudinal, or did it use both types of data? If the design was longitudinal, what type of longitudinal design was it? Could the longitudinal design have been improved in any way, as by collecting panel data rather than trend data, or by decreasing the dropout rate in a panel design? If cross-sectional data were used, could the research question have been addressed more effectively with the longitudinal data? (Chapter 5)
13. Were any causal assertions made or implied in the hypotheses or in subsequent discussion? What approach was used to demonstrate the existence of causal effects? Were all three criteria for establishing causal relationships addressed? What, if any, variables were controlled in the analysis to reduce the risk of spurious relationships? Should any other variables have been measured and controlled? How satisfied are you with the internal validity of the conclusions? (Chapters 5, 6)
14. Was a sample or the entire population of elements used in the study? What type of sample was selected? Was a probability sampling method used? Did the authors think the sample was generally representative of the population from which it was drawn? Do you? How would you evaluate the likely generalizability of the findings to other populations? (Chapter 4)
15. Was the response rate or participation rate reported? Does it appear likely that those who did not respond or participate were markedly different from those who did participate? Why or why not? Did the author(s) adequately discuss this issue? (Chapters 4, 6)
16. Was an experimental, survey, participant observation, or some other research design used? How well was this design suited to the research question posed and the specific hypotheses tested, if any? Why do you suppose the author(s) chose this particular design? How was the design modified in response to research constraints? How was it modified in order to take advantage of research opportunities? (Chapters 6–10)
17. Was an evaluation research design used? Which type was it? What was the primary purpose of the evaluation? (Chapter 11)

18. Were multiple methods used? Were findings obtained with different methods complementary? (Chapter 11)
19. Was any attention given to social context? To biological processes? If so, what did this add? If not, would it have improved the study? Explain. (Chapter 5)
20. Summarize the findings. How clearly were statistical and/or qualitative data presented and discussed? Were the results substantively important? (Chapters 9, 13)
21. Did the author(s) adequately represent the findings in the discussion and/or conclusions sections? Were conclusions well grounded in the findings? Are any other interpretations possible? (Chapter 14)
22. Compare the study to others addressing the same research question. Did the study yield additional insights? In what ways was the study design more or less adequate than the design of previous research? (Chapters 2, 10, 13)
23. What additional research questions and hypotheses are suggested by the study's results? What light did the study shed on the theoretical framework used? On social policy questions? (Chapters 2, 10, 13)

Below we have reproduced these questions followed by our answers from the review of the Lu et al. (2013) article. You can also follow our review by reading through the article itself and noting our comments.

1. *What is the basic research question or problem? Try to state it in just one sentence.* (Chapter 2)

   Is low self-control associated with misbehavior among juveniles in China? Is any impact of this variable independent of an effect of measures of social bond theory and individual demographic variables?

2. *Is the purpose of the study explanatory, evaluative, exploratory, or descriptive? Did the study have more than one purpose?* (Chapter 1)

   The study is explanatory. The authors wish to establish the potential impact of one (set of) variable(s) on another—measures of self-control and the social bond on the incidence of deviant behavior in juveniles in a specific cultural context.

3. *Was the theoretical framework presented? What was it? Did it seem appropriate for the research question addressed? Can you think of a different theoretical perspective that might have been used?* (Chapter 2)

   Two different theoretical frameworks are specifically being tested in this article—Gottfredson and Hirschi's self-control theory and Hirschi's social bond theory. This study is designed as a test of those theories as applied in an under-researched context (China). It would certainly be possible to examine the accuracy of other Western theories of delinquency with this population as well.

4. *What prior literature was reviewed? Was it relevant to the research problem? To the theoretical framework? Does the literature review appear to be adequate? Are you aware of (or can you locate) any important omitted studies?* (Chapter 2)

   In the section called "Theory and Prior Research," Lu et al. (2013) discuss the basics of self-control and social bond theories and cite a variety of other studies that have tested both of these theories. In addition, they provide a discussion of studies that have applied concepts from either theory in non-Western—specifically East Asian—contexts. Given the purpose of the study—to establish the efficacy of either or both of these theoretical frameworks in a non-Western context—the review of these areas of research appear to be appropriate and adequate. We leave it to you to decide if any important studies were omitted.

5. *How well did the study live up to the guidelines for science? Do you need additional information in any areas to evaluate the study? To replicate it?* (Chapter 2)

   The study clearly involves a test of ideas (two formal theories) against empirical reality (measures of behavior among Chinese adolescents). The "Methods" section of the article clearly tells us how the investigation was systematically carried out—there's a careful (and well-specified) research design. This design is well

documented and clear (and obviously publicly disclosed as the article has been published). They clarify their assumptions in the "Theory and Prior Research" section of the paper. There is a full section in the paper devoted to discussion of "Measures" that is, the ways in which key concepts in the study were defined and measured. The authors are building on other empirical research and attempting to replicate these studies but in a very different cultural context. They clearly maintain an interest in theory—this paper is a deductive approach to knowledge that presents a test of specific theories. They do not make any assumptions about what they will find—for example, they write that the analysis examines what effect of self-control measures—"if any"—remains after considering the effect of measures of social bond theory and demographic controls. They clearly have no assumptions about the potential impact of those social control measures. Their goal is to search for patterns of regularities in the data—"Are there predictable and discernible patterns that emerge in an examination of the delinquent behaviors of Chinese teenagers?" Thus, this study seems to exemplify adherence to basic scientific guidelines.

6. *Did the study seem consistent with current ethical standards? Were any trade-offs made between different ethical guidelines? Was an appropriate balance struck between adherence to ethical standards and use of the most rigorous scientific practices?* (Chapter 3)

   The authors make no specific references to adherence to ethical standards, nor is there a specific citation for approval by an internal review board for their study methodology. However, given that all four authors are employed at major American universities, it is likely safe to assume that such a review by a human subjects board did take place. The questionnaire used in the study asked adolescents about deviant behavior, and the authors do tell us that this questionnaire was anonymous. Although there is no reason to assume that ethical standards were not upheld, the authors might have been more specific in their description of the methodology in this respect.

7. *What were the major concepts in the research? How, and how clearly, were they defined? Were some concepts treated as unidimensional that you think might best be thought of as multidimensional?* (Chapter 4)

   The following concepts were used in the research: risky behavior, minor delinquency, self-control, attachment, school commitment, involvement, belief, age, gender, family structure, and delinquent peers. The definitions of the key concepts linked with the two theories (self-control, attachment, school commitment, involvement, and belief) receive special attention. Several of the variables used in the study are multidimensional. For example, the operationalization of the key concept of self-control utilizes an index designed to tap into various dimensions of this measure. Three of the four aspects of social bonds (attachment, involvement, and belief) were measured with multiple items. The fourth (school commitment) was operationalized as the response to a single item on the survey ("How well do you do in school compared with other students in your class?"). This measure of school commitment might have been more complex, although the strategy the authors use is not inconsistent with the way this concept has been operationalized in other research.

8. *Were any hypotheses stated? Were these hypotheses justified adequately in terms of the theoretical framework? In terms of prior research?* (Chapter 2)

   The authors offer no specifically stated set of hypotheses about the ways in which self-control theory and social bond theory will perform in predicting the delinquent behaviors in this population of Chinese youth. They certainly do identify the results of previous research with regard to the study of these measures in various contexts but do not express any expectations about the relationships that might be found in this cultural context.

9. *What were the independent and dependent variables in the hypothesis(es)? Did these variables reflect the theoretical concepts as intended? What direction of association was hypothesized? Were any other variables identified as potentially important?* (Chapter 2)

   There are two dependent variables—prevalence of risky behavior and of minor delinquency. Independent variables are self-control, attachment, school commitment, involvement, and belief. Demographic control

variables include age, gender, family structure, and delinquent peers. These variables are all directly related to the theories being tested.

10. *Did the instruments used—the measures of the variables—seem valid and reliable? How did the author attempt to establish this? Could any more have been done in the study to establish measurement validity?* (Chapter 4)

    Because both self-control and social bond theories have been widely tested in the field, there are certain agreed-upon operationalizations of concepts relevant for these theories that exist (all of which are discussed in the article). With regard to the survey instrument used to collect the data, the authors note that the "validity and reliability of the . . . core questionnaire have been examined and found to be quite satisfactory." Standardized measures (such as "Grasmick et al.'s . . . self-control scale, including 12 items on impulsivity, risk-seeking, self-centeredness, and temper") were also used. It appears that the authors made use of established measures, which have been previously subject to examination for reliability and validity, in their own work.

11. *What were the units of analysis? Were they appropriate for the research question? If some groups were the units of analysis, were any statements made at any point that are open to the ecological fallacy? If individuals were the units of analysis, were any statements made at any point that suggest reductionist reasoning?* (Chapter 5)

    The unit of analysis in this study was an individual—a student in a school in China. This unit of analysis is appropriate for this research question as the theories being tested here are ones that predict individual-level behavior. There are no statements that suggest reductionist reasoning.

12. *Was the study design cross-sectional or longitudinal, or did it use both types of data? If the design was longitudinal, what type of longitudinal design was it? Could the longitudinal design have been improved in any way, such as by collecting panel data rather than trend data or by decreasing the dropout rate in a panel design? If cross-sectional data were used, could the research question have been addressed more effectively with the longitudinal data?* (Chapter 6)

    This study was cross-sectional. Data were collected from the Chinese students at a single point in time.

13. *Were any causal assertions made or implied in the hypotheses or in subsequent discussion? What approach was used to demonstrate the existence of causal effects? Were all three criteria for establishing causal relationships addressed? What, if any, variables were controlled in the analysis to reduce the risk of spurious relationships? Should any other variables have been measured and controlled? How satisfied are you with the internal validity of the conclusions?* (Chapters 5, 6)

    Although the authors do not specifically say that they are in pursuit of causal relationships, there is some evidence that might be used to suggest the existence of a causal relationship. They establish association between the independent and dependent variables. They also address issues of spuriousness directly—the impact of self-control measures is examined both with and without the addition of measures of social bonds. The four control variables used (age, gender, family structure, and delinquent peers) are also appropriate in an effort to address spuriousness. These measures might be expected to be associated with both the independent and the dependent variable, so including them as controls is a wise move. There are potentially more problems with the time order element. Because data were only collected at a single point in time, it might be difficult to ascertain, for some measures, the direction of the causal relationship. For example, it is possible that a low level of school commitment (measured with a question that asked how well the respondent was doing in school) might have a causal association with deviant behavior (as predicted in social bond theory). It is also possible, however, that engaging in some risky behaviors (such as drinking alcohol) could have an effect on school performance (the measure of school commitment). Therefore, whereas association between the variables can be demonstrated, time order and nonspuriousness are not entirely established.

14. *Was a sample or the entire population of elements used in the study? What type of sample was selected? Was a probability sampling method used? Did the authors think the sample was generally representative of the population from which it was drawn? Do you? How would you evaluate the likely generalizability of the findings to other populations?* (Chapter 5)

    A probability sample is used in the study. The authors utilized a multistage cluster technique to randomly select middle schools in the city of Hangzhou and then randomly selected one class of seventh-, eighth-, and ninth-grade students in each school. All students in that randomly selected class were then asked to participate in the research. The site selected for the study was chosen because "[t]he city is a vivid reflection of the social and demographic changes in the coastal area in China where the economic boom has been the most noticeable." Consequently, the youth randomly selected from this population should be representative of a part of China that is of special interest. The random selection of participants should have ensured that the findings were generalizable.

15. *Was the response rate or participation rate reported? Does it appear likely that those who did not respond or participate were markedly different from those who did participate? Why or why not? Did the author(s) adequately discuss this issue?* (Chapters 5, 7)

    The response rate was quite high—96%. Lu et al. note that this very high response rate might be attributed in part to cultural factors. Response rates to self-administered surveys in China are routinely this high, whereas studies done in the United States with comparable methodologies have seen response rates in the area of 75%. The high response rate is a good thing, of course, with regard to the confidence a reader can have in the results of the study. It is also significant that Lu et al. provide context for this (perhaps seemingly inordinately) high response rate.

16. *Was an experimental, survey, participant observation, or some other research design used? How well was this design suited to the research question posed and the specific hypotheses tested, if any? Why do you suppose the author(s) chose this particular design? How was the design modified in response to research constraints? How was it modified in order to take advantage of research opportunities?* (Chapters 7, 8)

    The study employed a survey methodology. A great deal of other research (cited in the paper) that has tested both self-control theory and social bond theory with adolescents in the West has also used survey methodologies. The fact that survey research is established as a vehicle for examining the types of questions that are of interest here was likely a factor in the selection of this methodology for this study. The survey used in this study is based on a previously used instrument (the International Self-Report Delinquency Study). In writing about this survey, the authors state, "[t]he validity and reliability . . . have been examined and found to be quite satisfactory." The site of the study, China, did require a translation of this instrument. Lu et al. note that the translated survey was pretested with a group of Chinese exchange students "to make the questionnaire better fit the Chinese social, cultural, and language contexts."

17. *Was an evaluation research design used? Which type was it? What was the primary purpose of the evaluation?* (Chapter 10)

    No, this study is not an evaluation design.

18. *Were multiple methods used? Were findings obtained with different methods complementary?* (Chapter 11)

    This study used only survey methodology. Given the fact that the impact of cultural context was a variable of interest here, it is possible that this research question could also be effectively addressed with a more qualitative methodology in the future—one that would allow for the discovery of context and the application of meaning on the part of the subjects.

19. *Summarize the findings. How clearly were statistical and/or qualitative data presented and discussed? Were the results substantively important?* (Chapter 12)

The authors include a discussion and conclusions section where they clearly summarize the major findings of the analysis. They found that the results of the study in this Chinese context were comparable to those found with Western (primarily American) samples. Self-control was found to be inversely related to measures of both risky behavior and minor delinquency. The effects of the self-control measure remained even after including social bond measures and the demographic variables in the model, and the self-control variables had a stronger effect than did the social bond measures. Among the social bond measures, only belief and family attachment were found to have a significant effect on the dependent variables. Self-control was found to be more strongly related to minor delinquency than to risky behaviors.

20. *Did the author(s) adequately represent the findings in the discussion and/or conclusion sections? Were conclusions well grounded in the findings? Are any other interpretations possible?* (Chapter 12)

Lu et al. have extensive "Discussion" and "Conclusion" sections in which they summarize and review the major findings of their analysis and also offer some insight into the significance, and potential problems, with the research. First, it is a significant step to find confirmation of the predictions of self-control and social bond theories in a non-Western population of children. The authors made it clear that this was a primary goal of the study, and the fact that these theories perform as predicted in a very different cultural context adds to the evidence supporting the universality of these relationships. There are also some notable limitations in the study, which Lu et al. point out. For example, only a small amount of the total variance in risky behaviors or delinquency is accounted for by the variables used in this study. Thus, the authors recommend that other "theoretical explanations, such as general strain theory, may be considered for future investigations." They also suggest that more direct comparison of the data from this study with that obtained with the same (or similar) instruments in other countries would be useful—"the ISRD-2 survey has been implemented in 30 countries (and) future research may explore the comparative aspect of self-control across different cultures to examine the explanatory power of self-control theory."

21. *Compare the study to others addressing the same research question. Did the study yield additional insights? In what ways was the study design more or less adequate than the design of previous research?* (Chapters 2, 12)

The most significant thing about this study was its confirmation of the findings of many other studies regarding the impact of self-control and social bond variables on delinquency. The fact that these data came from a very different cultural context—China—was the main point of the research. In their explication of the theory, Gottfredson and Hirschi "pointed out (that) 'culture variability is not important in the causation of crime and we should look for constancy rather than variability in the definition of and causation of crime'" (1990, p. 33). The application and testing of the theory in other cultural contexts—as done in this article—can therefore be considered a necessary part of testing this popular and powerful approach to explaining delinquency and crime.

Lu et al. do call attention to some measurement issues with their study that might have been problematic. In measuring the key variable of self-control, they use an established index (Grasmick et al.'s self-control scale) but only make use of the short version (with 12 items) of this widely tested instrument.

22. *What additional research questions and hypotheses are suggested by the study's results? What light did the study shed on the theoretical framework used? On social policy questions?* (Chapters 2, 12)

The results of this study definitely contribute to the validation of self-control and social bond theories, offering a test of the theories in a non-Western context. To establish that the impact of self-control measures on the likelihood of the occurrence of criminal or delinquent behaviors is universal, the theory should be tested in more cultural contexts. This successful application in China also only involved students in one province of this very large country—more data from within China, as well as from other Asian countries, would also be useful. Lu et al. also argue that additional theoretical perspectives (e.g., strain theory) should also be tested in the cultural context of China.

Exploring the Utility of Self-Control Theory for Risky Behavior and Minor Delinquency Among Chinese Adolescents

Journal of Contemporary Criminal Justice
29(1) B-8–52

DOI: 10.1177/1043986212471240
ccj.sagepub.com
SAGE

Yi-Fen Lu[1], Yi-Chun Yu[1], Ling Ren[1], and Ineke Haen Marshall[2]

Abstract

Although scholarly interest in empirically assessing Gottfredson and Hirschi's (1990) self-control theory continues to grow, there is little research available on evaluating its utility in China. The current study examines the power of self-control theory as an explanation of juvenile minor risky behavior as well as minor delinquent behaviors in China, while simultaneously considering the role of social bonds. The data were collected from a probability sample of 7th, 8th, and 9th graders in Hangzhou, China (N = 1,043) using the second International Self-Report Delinquency (ISRD-2) survey instrument. The results from the logistic regression analyses show that self-control is associated with both measures of Chinese youth deviance in the expected direction, suggesting that the youth with low self-control have a higher likelihood to report minor risky and delinquent behavior. The effects of self-control remain when controlling for social bonding measures. Among the social bonding factors, beliefs (pro-violence attitude) and family bonding were found to be the significant predictors of risky behavior and minor delinquency, respectively.

Keywords

self-control, social bonding, Chinese adolescents, International Self-Report Delinquency survey

Introduction

Gottfredson and Hirschi's self-control theory (1990), known as "the general theory of crime," has been one of the most influential and most empirically scrutinized criminological theories in recent decades (Cheung & Cheung, 2008; Cretacci, Rivera, & Ding, 2009). An abundance of research has been generated to test the proposition that low self-control is the main cause of crime, delinquency, and numerous analogous behaviors (Arneklev, Grasmick, Tittle, & Bursik, 1993; Cochran, Wood, Sellers, Wilkerson, & Chamlin, 1998; Gibbs, Giever, & Martin, 1998; Gibson, Wright, & Tibbetts, 2000; Junger, West, & Timman, 2001; Keane, Maxim, & Teevan, 1993; Marshall & Enzmann, 2012; Wood, Pfefferbaum, & Arneklev, 1993). Its robustness is evidenced in the meta-analytic review of 21 empirical studies by Pratt and Cullen (2000), who found that low self-control is an important predictor of crime and of analogous behaviors across different measurements and types of samples. An important feature of self-control theory is that it departs from the culture-difference approach in the analysis of comparative criminology (Cheung & Cheung, 2008). As

[1]College of Criminal Justice, Sam Houston State University, Huntsville, TX, USA
[2]Department of Sociology & Anthropology and School of Criminology and Criminal Justice, Northeastern University, Boston, MA, USA

Corresponding Author:
Yi-Fen Lu, College of Criminal Justice, Sam Houston State University, Huntsville, TX 77341, USA.
Email: yxl007@shsu.edu

Source: Lu, Yi-Fen, Yi-Chun Yu, Ling Ren, and Ineke Haen Marshall. 2013. "Exploring the utility of self-control theory for risky behavior and minor delinquency among Chinese adolescents." *Journal of Contemporary Criminal Justice* 29(1): 32–52.

Gottfredson and Hirschi (1990) pointed out, "culture variability is not important in the causation of crime and we should look for constancy rather than variability in the definition of and causation of crime . . . " (p. 175). In other words, the predictive power of self-control on delinquency can be manifested across different cultural settings.

The empirical assessments of self-control theory, however, have been disproportionally conducted in American samples (Pratt & Cullen, 2000), with only a modest number in other Western nations (e.g., Caspi et al., 1994; LaGrange & Silverman, 1999; Marshall & Enzmann, 2012; Romero, Gomez-Fraguela, Luengo, & Sobral, 2003; Tittle & Botchkovar, 2005; Vazsonyi, Clifford Wittekind, Belliston, & Van Loh, 2004; Vazsonyi, Pickering, Junger, & Hessing, 2001), and even fewer in Asian countries (Cheung & Cheung, 2008; Cretacci, Ding, & Rivera, 2010; Cretacci et al., 2009; Hwang & Akers, 2003; Vazsonyi et al., 2004; Wang, Qiao, Hong, & Zhang, 2002). Researchers have highlighted the paucity of research on self-control in China, a nation where self-control appears to be more emphasized due to its collectivistic feature of traditional culture (e.g., Cheung & Cheung, 2008; Cretacci et al., 2009). As opposed to the emphasis on individualism in the West, the philosophy of Confucianism has been the prominent belief in China. The major Confucian concepts related to personal traits particularly stress self-discipline and suppression of individual autonomy for the purpose of facilitating collective interests (Chan & Lee, 1995; Yang, 1995). Children are taught to be disciplined from a very early age in the family and school. As a result, we would hypothesize that self-control may be more likely to develop among Chinese youth which in turn restrains them from delinquency.

To date, only three quantitative studies published in English could be found to inform this study on the utility of self-control theory in mainland China. They are Wang et al. (2002) on the relationship between two dimensions of self-control (impulsivity and persistency) and Chinese adolescences' substance abuse and delinquency; Cretacci et al. (2009) investigation of effects of the traditional Grasmick et al.'s self-control scale and Hirschi's revised self-control measure on deviance among Chinese college students; and Cretacci et al.'s (2010) expanded work based on the previous piece on Hirschi's social-bond-type measure of self-control. One similar finding across these three studies is that self-control variables measured partially or fully by Grasmick et al.'s scale did not demonstrate significant effects on various delinquent behaviors. At the same time, it is worth noting that social bond measures (e.g., family attachment, parental supervision, and educational commitment) included in these three studies showed noticeable predictive power in explaining Chinese deviance. Wang et al. (2002) and Cretacci and his associates (2009, 2010) made a strong argument for the importance of the inclusion of social bonding while testing self-control theory in the Chinese context. In the current article, we respond to this need.

The recent socioeconomic change in China has been astronomical by any standard. After 30 years of nonabating economic growth, China is officially the world's second largest economy, next to the United States. Behind these dazzling economic miracles are the unbalanced distribution of quickly accumulated wealth and disintegration of traditional values (Cao, 2007). Not surprisingly, Wang (2006) reported that the crime rate among juveniles was on the rise from 1980 to 2004 and identified a primary cause of this increase as associated with problems such as broken families, social disorder, and unemployment that "threaten the normal socialization of juveniles" (p. 5). In addition, China's rapid social changes have obvious implications for increasing adolescent problem behavior, including cigarette smoking and drinking (Jessor et al., 2003). Therefore, it seems imperative to investigate if self-control holds strong among adolescents in the large Chinese cities at the time when major social changes are taking place. Equally important, we are interested in examining the explanatory power of self-control theory for various forms of adolescent delinquency and analogous behaviors in comparison to the social bonding perspective.

Theory & Prior Research

Self-Control Theory & Social Bond Theory

Grounded on the perspective of classical school, both self-control and social bond theories assume "that humans had free will and that behavior was guided by hedonism" (Lilly, Cullen, & Ball, 2007, p. 15). Gottfredson and Hirschi (1990) posit that the propensity to engage in any crime, delinquency, and "analogous" behaviors (i.e., drinking, smoking, and substance use) is a result of low self-control in conjunction with the presence of opportunity. Low self-control is manifested by impulsivity, a preference for simple tasks and physical activity, risk-seeking, self-centeredness, and volatile temper (pp. 89-90). These elements, either constructing a unidimensional (Piquero & Rosay, 1998) or multidimensional latent trait (Longshore, Turner, & Stein, 1996), have been found significantly

associated with all forms of crime as well as other types of behavior among adolescents and adults (Arneklev et al., 1993; Baron, 2003; Brownfield & Sorenson, 1993; Burton, Cullen, Evans, Alarid, & Dunaway, 1998; Cochran et al., 1998; De Li, 2004; Evans, Cullen, Burton, Dunaway, & Benson, 1997; Gibbs et al., 1998; Gibson et al., 2000; Grasmick, Tittle, Bursik, & Arneklev, 1993; Polakowski, 1994; Vazsonyi et al., 2001; Wood et al., 1993).

In contrast, Hirschi's earlier (1969) social bond theory premises that people are naturally criminal and bonding to conventional society (i.e., family, school, and peers) acts to restrain these natural tendencies to be criminal. Delinquency results "when an individual's bond to society is weak or broken" (Hirschi, 1969, p. 16). Arguing in support of the primary importance of low self-control, Gottfredson and Hirschi (1990) argued that the weakness or absence of social bonds is not the cause of crime; rather, it is the *consequence* of low self-control. Individuals with low self-control are more likely to have weak bonds with families and teachers than those with higher self-control. Empirical studies that have attempted to disentangle the link between social bonding, self-control, and delinquency have provided mixed findings in this regard (Baron, 2003; Brownfield & Sorenson, 1993; Burton et al., 1998; Evans et al., 1997; Grasmick et al., 1993; Manson & Windle, 2002; Polakowski, 1994; Pratt & Cullen, 2000; Wright, Caspi, Moffitt, & Silva, 1999). Wright et al. (1999), for example, tested three theoretical models: a social-selection model, a social-causation model and a mixed selection-causation model. They found that in support of the social-selection model, the respondents with low self-control tend to possess weaker bonds to family and school as well as demonstrated through lower work achievement. Meanwhile, in support of the social-causation model, social bonds significantly predict criminal offending later in life. They also found the effects of self-control on crime are largely mediated by social bonds. Finally, in support of the mixed selection-causation model, despites partial attenuation, the correlation between the social bonds and the measure of delinquent behaviors remained statistically significant while controlling for childhood and adolescent self-control.

Empirical Studies From East Asia

One of the most daring theoretical propositions of self-control theory is that the low self-control-deviance relationship persists across cultural and national boundaries. Gottfredson and Hirschi devoted one entire chapter of their influential book in 1990 to culture and crime, where they explained why self-control should be applicable across cultures. Gottfredson (2006) reemphasized this notion by providing empirical support and explicating that self-control "should predict rate differences everywhere, for all crimes, delinquencies, and related behaviors, for all times, among all groups and countries" (p. 83). Empirical research conducted in Western countries support the conclusion that the effect of self-control on deviance is invariant across national boundaries (Caspi et al., 1994; Keane et al., 1993; LaGrange & Silverman, 1999; Romero et al., 2003; Tittle & Botchkovar, 2005; Wright et al., 1999). It is, however, worth noting that there has been scant and inconsistent empirical evidence derived from East Asian samples to support their claim that the self-control-deviance relationship is persistent across cultures.

Recently, researchers have extended the scope to East Asia to examine the culture-free propositions of self-control theory (e.g., Cheung & Cheung, 2008; Cretacci et al., 2009, 2010; Hwang & Akers, 2003; Vazsonyi et al., 2004; Wang et al., 2002). Yet the number of studies conducted in East Asia is rather small, including one in Japan by Vazsonyi and his colleagues (2004), Hwang and Akers's (2003) study in South Korea, one in Hong Kong by Cheung and Cheung (2008), and three in mainland China by Wang et al. (2002) and Cretacci et al. (2009, 2010). These studies not only assessed the effect of self-control but also gauged the relative importance of various social factors (i.e., social bonding, social learning, strain, and labeling).

For example, Hwang and Akers (2003) gathered data using a self-report questionnaire from a sample of 1,012 adolescents in Pusan, South Korea, in 1999. They found that the effects of self-control and social bonding on alcohol and tobacco use among South Korean youth disappeared when social-learning variables were considered. In addition, Wang et al. (2002), utilizing a 1997 sample of 527 adolescents, residing in the capital city of a province in Southern China, investigated the predictive power of self-control on illicit substance use (tranquilizer, opium, and heroine) and on deviant behavior (fighting, stealing, and telling lies). The results revealed that adolescent impulsivity (one dimension of self-control) was not directly related to either substance abuse or deviant behavior. The second dimension of self-control, persistency, was found to be positively related to substance abuse, which contradicts self-control predictions although it was indirectly associated with substance abuse and deviance, mediated by social bonding. They further argued that self-control theory might not be applicable in China. In Wang et al.'s (2002) study, however, only six items forming two dimensions (impulsivity and persistency) of self-control measure were included in the analysis. Similarly, Cheung and

Cheung (2008) failed to find evidence that self-control had an impact on delinquency, except violent delinquency, in a sample of 1,015 adolescents in Hong Kong after social bonds and peer influence were weighed in. Drawing upon multidimensional self-control measures, Vazsonyi et al. (2004) found that low self-control was consistently associated with diverse measures of Japanese late adolescent deviance, ranging from trivial to more serious norm-violating behaviors such as school misconduct and assault. Their results also suggested that there was no statistically significant relationship between low self-control and alcohol use among Japanese youth in the sample. More recently, recognizing the paucity of self-control research in mainland China, Cretacci and his associates (2009, 2010) collected data in a large Chinese university located in Beijing in the fall of 2007. In their first study published in 2009, Cretacci et al. employed both Grasmick et al.'s (1993) self-control scale and Hirschi's (2004) revision that reflects a one-dimensional, social bond type measure. They found that the significant effect of the Grasmick et al.'s scale disappeared when Hirschi's (2004) revised scale was added to the logistic regression model. They further suggested that self-control theory simply is not an important contributor to Chinese deviance models if Grasmick et al.'s scale is determined to be the better measure of self-control. Using the same data set, Cretacci et al. (2010) focused on Hirschi's revised scale and formulated three separate factors (maternal relationship, school attachment, and school authority). Their findings were largely in line with their previous study (2009). That is, the Grasmick model may have little impact on Chinese deviance and the Hirschi's revised "bond type" measure of self-control was significantly predictive of deviance in their sample. It is important to note that the external validity of Cretacci et al. (2009, 2010) might be weakened due to the convenience sample with only 150 university students.

The limited body of research on self-control in East Asian countries brings into question the applicability of self-control theory to appreciably different cultures. Of these studies discussed above, Vazsonyi et al. (2004) is the only study that demonstrated a significant relationship between self-control and various delinquent behaviors, but they did not investigate if the effect of self-control would remain when other competing theories were taken into consideration. In contrast, the findings from mainland China, Hong Kong, and South Korea were at odds with the propositions of self-control theory that have been largely supported in multiple Western societies. It has been argued that the cultures of Japan, Hong Kong, and South Korea generally are viewed as a broad development of Chinese-oriented culture. Compared with the United States or other Western countries, the societies of Chinese-oriented culture place greater interests on the collectivistic context. That is, in Chinese society, individuals are encouraged to control themselves and put societal interests above individual desires. At the same time, because of the close-knit Chinese culture socialization can be very effective in reinforcing or deterring youth delinquency (Wang et al., 2002).

The Current Study

The current study represents an attempt to expand the body of research on the assessment of self-control theory through analysis of a large probability sample drawn from a city with 5.5-million residents on the east coast of China. In addition to the measure of self-control, social bond factors, including family, school, and neighborhood bonds, school commitment, involvement in conventional activities, and attitudes toward violence are included to assess the variations in Chinese adolescent misbehavior. The adolescent misbehaviors include minor risky behaviors (smoking and drinking) and minor delinquency (vandalism, weapon carrying, and group fighting).[1] The analysis begins by determining if low self-control itself is associated with Chinese juvenile misbehavior. Then, it proceeds to examine whether or not the effect of self-control, if any, remains significant while the social conditions measured by the social bonding perspective and individual demographic variables are taken into consideration.

Method

Survey Instrument

The current study used the second International Self-Report Delinquency survey instrument (ISRD-2). This instrument is based on the one used in first large-scale International Self-Report Delinquency Study (ISRD-1) (For more information, please see Introduction to this Special Issue). The validity and reliability of the ISRD core questionnaire have been examined and found to be quite satisfactory (see Bruinsma, 1994; Marshall & Webb, 1990, 1994; Zhang, Benson, & Deng, 2000).

Sample & Data Collection

The research site for the Chinese study is Hangzhou, the capital city of Zhejiang province, which is about 150 miles southwest of Shanghai. Hangzhou has been a rapidly growing city and according to the city official website (http://eng.hangzhou.gov.cn/), the population of long-term residents (not including the migrant population) in urban districts is 5.5 million based on survey results from 2008. The city is a vivid reflection of the social and demographic changes in the coastal area in China where the economic boom has been the most noticeable.

Due to the large student population and complexity of its demographics in Hangzhou, a multistage cluster sampling technique was employed for the sample selection. This type of sampling approach is suitable when it is impractical to compile a complete sampling frame. As a result, nine middle schools located in five core urban districts in the city were selected. In each of the selected schools, one class each was randomly selected from the 7th to 9th grades. In collaboration with the Zhejiang Provincial Juvenile Delinquency Institute (ZPJDI), we gained access to these nine schools selected for the sample. The translated ISRD-2 instrument was pretested using 16 Chinese exchange students at Sam Houston State University (SHSU) to make the questionnaire better fit the Chinese social, cultural, and language contexts. Data were gathered in late December 2009 and early January 2010 by means of anonymous, self-report questionnaires (paper-and-pencil) administered during a class period. The Chinese school year is different from the United States; the winter break usually starts in the middle of January before the Chinese Lunar New Year. Researchers from SHSU worked closely with the trained members of the research staff from ZPJDI for data collection. At least two researchers were present in one classroom to administer the surveys and teachers and school administrators were asked to leave before the questionnaires were distributed. Questionnaires were filled out by 96% of the sample, resulting in 1,043 useable surveys.[2]

Measures

Dependent Variables. Two dependent variables were employed in the current study, the lifetime prevalence of risky behavior and of minor delinquency. Risky behavior was measured through three separate questions: (a) Did you ever drink beer, breezers, or wine? (b) Did you ever drink strong spirits (gin, rum, vodka, whisky)? and (c) Did you ever smoke? All items were coded as 1 if the answer was "yes" and 0 otherwise. An additive scale was calculated so the total score ranged from 0 to 3. A higher score represents a higher number of minor risky behaviors reported. This variable was subsequently recoded into a dichotomous measure of having ever engaged in any minor risky behavior (0 = *No*; 1 = *Yes*). To measure the second dependent variable, minor delinquency, respondents were asked three questions: (a) Did you ever participate in a group fight on the school playground, a football stadium, the streets or in any public place? (b) Did you ever carry a weapon, such as a stick, knife, or chain (not a pocket-knife)? and (c) Did you ever damage something on purpose, such as a bus shelter, a window, a car, or a seat in the bus or train? The variable transformation was identical to the variable of risky behavior.

Independent Variables

Self-Control. The current study used a shortened version of the Grasmick et al.'s (1993) self-control scale, including 12 items on impulsivity, risk-seeking, self-centeredness, and temper (see appendix). The results of both exploratory factor analysis and confirmatory factor analysis suggest that the 12 items form a single latent trait.[3] The unidimensional model of the self-control measure is consistent with prior studies (Baron, 2003; Cochran et al., 1998; Gibson et al., 2000; Marshall & Enzmann, 2012, p. 319; Piquero & Rosay, 1998). Scores of the items were used to form a composite self-control scale that is calculated by dividing the sum of total scores by 12 and multiplying the quotient by 100. A lower score of the scale indicated a lower level of self-control. The Cronbach's alpha of this scale was 0.852 with an Eigenvalue of 4.776, which is compatible with Grasmick et al.'s scale (approximate .80; Grasmick et al., 1993; Piquero, MacInosh, & Hickman, 2000).

Attachment. Similar to prior research on testing social bond theory (Cheung & Cheung, 2008; De Li, 2004), we followed Hirschi's (1969) original classification of four components of the bond: attachment, commitment, involvement, and belief. Attachment was measured based on family bonding, school bonding, and neighborhood bonding. Family bonding scale was captured by the three items asking "How do you usually get along with the

man you live with (father, stepfather . . .)?," "How do you usually get along with the woman you live with (mother, stepmother . . .)?," and "Do your parents (or the adults you live with) know whom you are going out with when you go out every time?" The scale was calculated as the sum of scores of three items divided by three and multiplied the product by 100. A higher number indicated a higher level of family bonding (α = .590, Eigenvalue = 1.682). School bonding scale was measured using a 4-point Likert-type scale (1 = *not at all true*, 2 = *not true*, 3 = *true*, 4 = *very true*) on the questions: "If I had to move, I would miss my school," "Teachers do notice when I am doing well and let me know," "I like my school," and "There are other activities in school besides lessons." The scale was formed by dividing the sum of scores of the four items by four and multiplied the quotient by 100. A higher score suggested a higher level of school bonding (α = .751, Eigenvalue = 2.314). Neighborhood bonding scale was measured by six items led by the questions: "If I had to move, I would miss the neighborhood," "I like my neighborhood," "There is a lot of space for children to play," "People around here are willing to help their neighborhood," "There is a close-knit neighborhood," and "People in this neighborhood can be trusted." It was calculated as the sum of scores of the six items divided by six and multiplied the quotient by 100. A higher score reflected that respondents live in a neighborhood with a higher level of bonding (α = .881, Eigenvalue = 3.854).

School Commitment. Due to the limitation of the data, self-evaluated school achievement is employed as a proxy measure for school commitment. Some previous studies have tapped into this concept by using a similar measure. For example, Gibson et al. (2000) employed a student's average grade (one of three items) to measure the adolescents' commitment to school. In addition, family commitment to education as well as the intensely competitive education system in China puts greater *emphasis* on *academic success* than any other aspect of school life (Chen, Lee, & Stevenson, 1996). It is, therefore, plausible to argue that the youth with higher school achievement are more likely to put forward an effort to commit to education. School commitment was assessed by using a 3-point Likert-type scale (1 = *below average*, 2 = *average*, 3 = *above average*) on the question: "How well do you do in school compared with other students in your class?" A higher number indicated a higher level of school commitment.

Involvement. Involvement in conventional activities was gauged by asking respondents how much time they spend on an average school day on each of the activities: reading a book, reading magazines or comic strips, playing sports, playing a music instrument, and doing housework. Each item was measured by an ordinal variable and coded into 6 groups: (1) none; (2) ½ hr; (3) 1 hr; (4) 2 hr; (5) 3 hr; (6) 4 hr or more. The involvement scale was constructed as the sum of the five items divided by five and multiplied the quotient by 100. A higher score indicated that respondents spent more time on the conventional activities on an average school day (α = .637, Eigenvalue = 2.065).

Beliefs. Because the data set does not have a direct measure of "beliefs" as defined by Hirschi, in this study, belief (that one should obey the rules of society) was measured inversely by asking respondents' attitudes toward the use of violence. Thus, youth who responded that they have a positive attitude toward violence are considered to have a lower level of belief in traditional moral values. Respondents were asked how strongly they agree or disagree with the following statements of violent behavior done by young people: "A bit of violence is part of the fun," "One needs to make use of force to be respected," "If somebody attacks me, I will hit him/her back," "Without violence, everything would be much more boring," and "It is completely normal that boys want to prove themselves in physical fights with others." The responses range from 1 = *fully disagree* to 4 = *fully agree*. The belief scale was constructed as the sum of scores of five items divided by five and multiplied the quotient by 100. A higher score of the scale reflected a more pro-violence attitude (α = .719, Eigenvalue = 2.445).

Control Variables. Four additional variables are considered as control variables, including age, gender, intact family, and delinquent friends. Age was measured by actual years, ranging from 12 to 17. Gender was coded as 0 for female and 1 for male. Family structure was a dichotomous measure indicating 0 for broken family and 1 for intact family. Finally, a key control variable is delinquent friends. To measure this, we created a measure of delinquent friends adding four items into a index based on the following questions whether they have friends who stole something from a shop or department store, who entered a building with the purpose to steal something, who threatened somebody with a weapon or to beat him up just to get money or other things from him/her, and who beat someone up or hurt someone badly with something like a stick or a knife. A higher score indicated that respondents had more delinquent friends.

Results

The current study includes two dependent variables, the lifetime prevalence of self-reported minor risky behaviors and minor delinquency. Table 1 presents the descriptive statistics for all the variables employed in this study. For the prevalence of minor risky behavior, more than half of respondents (52.5%) reported having ever had beer, breezers or wine, strong spirits, or smoked in their lifetime. Further breakdowns for each category show that slightly more than half of respondents (51.8%) reported ever drinking beer/wine and 10.7% reported ever drinking strong spirits, which echoes the finding of the alcohol consumption pattern in China by Li, Fang, Stanton, Feigelman, and Dong (1996). Li et al. (1996) found that 63% and 54% of the 6th, 8th, and 10th graders in a sample of 1,040 reported that at least one time in their lifetime they have had beer and wine, respectively, and 11% have had hard liquor. The prevalence of teenage smoking reported in this study (6.3%), however, is lower than those in prior studies conducted in China (Li, Fang, & Stanton, 1999; Unger et al., 2001). Li and colleagues (1999) found that about 15% of the 7th to 9th grade students ($N = 323$) in Beijing reported ever having smoked, and Unger et al. (2001) found more

Table 1. Descriptive Statistics ($N = 1,043$).

Variables	Percentage (%)	Min	Max	Mean	Standard deviation
Dependent variables					
Minor risky behavior					
0 = *No*	47.5				
1 = *Yes*	52.5				
Minor delinquency					
0 = *No*	93.1				
1 = *Yes*	6.9				
Independent variables					
Self-control		0.00	100.00	82.77	15.63
Family bonding		16.67	100.00	88.96	16.42
School bonding		0.00	100.00	77.40	21.76
Neighborhood bonding		0.00	100.00	72.18	25.77
School commitment		1.00	3.00	2.15	0.52
Involvement		0.00	68.00	18.84	13.48
Belief		0.00	100.00	19.17	16.78
Control variables					
Age group					
12	8.3				
13	27.5				
14	29.4				
15	27.1				
16	7.3				
17	0.5				
Gender					
0 = *Female*	47.5				
1 = *Male*	52.5				
Intact family					
0 = *No*	14.1				
1 = *Yes*	85.9				
Delinquent friends					
0 = *0 Delinquency friend*	93.7				
1 = *1 Delinquent friend*	3.9				
2 = *2 Delinquent friends*	2.1				
3 = *3 Delinquent friends*	0.3				

Table 2. Logistic Regressions on the Prevalence of Minor Risky Behavior.

	Model 1				Model 2			
Variables	*b*	*b**	*SE*	*Odds Ratio*	*b*	*b**	*SE*	*Odds Ratio*
Intercept	3.615***	—	0.418	37.165	3.124***	—	0.764	22.744
Independent variables								
Self-control	–0.042***	–.295	0.005	0.959	–0.026***	–0.171	0.006	0.974
Family bonding					–0.007	–0.048	0.005	0.993
School bonding					–0.002	–0.018	0.004	0.998
Neighborhood bonding					–0.004	–0.043	0.003	0.996
School commitment					–0.072	–0.016	0.137	0.930
Involvement					–0.006	–0.034	0.005	0.994
Belief					0.018***	0.128	0.005	1.019
Control variables								
Age					0.108	0.050	0.065	1.114
Gender					0.118	0.025	0.142	1.125
Intact family					–0.386	–0.056	0.210	0.680
Delinquent friends					0.534*	0.086	0.238	1.706
–2 Log Likelihood	1,335.545				1,197.015			
Overall fit (Chi-Square)	88.423***				128.231***			
R_L^2	0.0621				0.0968			

Note: Logistic regression coefficients (*b*), standardized coefficient (*b**), standard errors (*SE*), and odds ratios are presented. R_L^2 is the likelihood ratio R^2, indicating the proportional reduction in the *-2LL* statistic. Asterisks represent statistically significant effects at the following levels: $*p < .05$. $**p < .01$. $***p < .001$.

than 30% of the same grade range students ($N = 6{,}992$) in Wuhan reported having tried smoking ever. In contrast to the overall high level of drinking and smoking, a much lower prevalence of self-reported minor delinquency was observed in the sample. That is, only about 7% of respondents reported having participated in a group fight (2.4%), carried a weapon (2.7%), or damaged something on purpose (3.9%) in their lifetime. This seems consistent with other reports of comparatively low levels of juvenile delinquency in China by using self-report data collection method (Greenberger, Chen, Beam, Whang, & Dong, 2000; Jessor et al., 2003; Wei, Homel, Prichard, & Xu, 2004).

For the independent variables, the mean value of the self-control scale was 82.77 with a standard deviation 15.63. As to the attachment measures, the respondents reported a high level of family bonding as being 88.96 with a standard deviation 16.42, while the average scores for school bonding and neighborhood bonding were comparatively lower, being 77.40 and 72.18, respectively. With respect to school commitment, the mean value was 2.15 with 1 indicating below average and 3 above average in the self-evaluated school performance. In addition, the average time the respondents spent on various conventional activities was about 18.84 with a standard deviation 13.48. Finally, the respondents as a whole reported a low level of positive attitudes toward violence ($M = 19.17$).

In terms of the demographic variables, the vast majority of the respondents (84%) fell within the age range 13 to 15 and slightly more than half of them identified themselves as male (52.5%). About 86% of the survey participants came from an intact family, and approximately 94% of the sample subjects reported they did not have friends who involved in shoplifting, burglary, extortion, and assault.

The results of the bivariate logistic regressions on the two dependent variables are partitioned into two tables (see Model 1 in Table 2 & Model 3 in Table 3). In the bivariate logistic regressions, the self-control scale was entered as the only independent variable to examine its effect on both risky behavior and minor delinquent behavior when other social factors were not considered. Then, multivariate logistic regressions were performed to assess the conditional

Table 3. Logistic Regressions on the Prevalence of Minor Delinquency.

	Model 3				Model 4			
Variables	*b*	*b**	*SE*	*Odds Ratio*	*b*	*b**	*SE*	*Odds Ratio*
Intercept	1.299**	—	0.506	3.667	0.716	—	1.173	2.046
Independent variables								
Self-control	–0.050***	–0.275	0.007	0.951	–0.030**	–0.160	0.009	0.970
Family bonding					–0.017*	–0.095	0.007	0.983
School bonding					–0.009	–0.066	0.007	0.991
Neighborhood bonding					–0.003	–0.026	0.006	0.997
School commitment					0.008	.001	0.240	1.008
Involvement					–0.016	–0.074	0.010	0.984
Belief					0.012	0.069	0.008	1.012
Control variables								
Age					0.160	0.060	0.130	1.174
Gender					0.960**	0.164	0.304	2.612
Intact family					0.213	0.025	0.405	1.237
Delinquent friends					0.588*	0.077	0.243	1.801
–2 Log Likelihood	461.599				396.007			
Overall fit (Chi-Square)	54.473***				89.219***			
R_L^2	0.1056				0.1839			

Note: Logistic regression coefficients (*b*), standardized coefficient (*b**), standard errors (*SE*), and odds ratios are presented. R_L^2 is the likelihood ratio R^2, indicating the proportional reduction in the *-2LL* statistic. Asterisks represent statistically significant effects at the following levels: *$p < .05$. **$p < .01$. ***$p < .001$.

effect of self-control after taking into account the various measures of social bonding (controlling for the demographic variables as well as delinquent friends). The outcomes from the multivariate logistic regressions are reported in the second portion of Table 2 and Table 3.

As can be seen in Model 1 (Table 2), the self-control scale was a statistically significant predictor of minor risky behavior involvement ($b = –.042$, $p < .001$), and it alone explained 6.21% of variation in the dependent variable. More specifically, a one-unit increase in the self-control measure was associated with a reduction of .042 in the log odds of ever having minor risky behaviors. Respondents who scored lower in the self-control scale were approximately 1.04 times more likely than those who scored higher to report having ever used alcohol and/or tobacco. When the social bonding variables were included in Model 2 (Table 2), self-control remained statistically significant ($b = –.026$, $p < .001$) and had the strongest effect ($b^* = –.171$) on youths' minor risky behavior. Among the social bond variables added to the model, the only significant predictor was the belief scale ($b = .018$, $p < .001$), suggesting that a one-unit increase in the belief scale was associated with an increase of .018 in the log odds of ever committing the minor risky behaviors. Respondents who reported having had pro-violence attitudes were 1.019 times more likely to have reported the use of alcohol and/or tobacco. It is interesting to note that having more delinquent friends was positively associated with minor risky behavior. Specifically, those who reported having more delinquent friends were 1.706 times more likely to involve in drinking and/or smoking ($b = .534$, $p < .05$). Among these three significant predictors, the self-control scale carried the most weight in explaining the dependent variable. Model 2 explained 9.68% of variation in the youth minor risky behaviors, a slight increase of 3.47% from Model 1.

With the same model specification, Table 3 displays the effects of self-control alone and with the social bonding measures on the second dependent variable, adolescent minor delinquency measured by group fight, carrying a weapon, and vandalism. As shown in Model 3, self-control was associated with youth minor delinquency ($b = –.050$, $p < .001$), and it alone explained 10.56% of variation. Specifically, the youth who had a lower score in the self-control scale were 1.052 times more likely to be involved in minor delinquent behavior than those with a higher self-control score. When including the

social bonding factors (Model 4), the effects of self-control did not disappear; instead, it remained statistically significant ($b = -.030$, $p < .01$) and was found as the second strongest predictor in terms of the relative predictive power ($b^* = -.160$), next to gender ($b^* = .164$). Among the social bonding measures, family bonding stood out as the only significant predictor of the involvement of minor delinquency ($b = -.017$, $p < .05$). Specifically, in comparison with the respondents who were closely attached to their parents, the youth with the lower family bonding were 1.017 times more likely to admit having committed the minor delinquency. As to the control variables, gender ($b = .960$, $p < .01$) and delinquent friends ($b = .588$, $p < .05$) manifested sizable effects on the dependent variable. Male respondents were 2.612 times more likely than their female counterparts to be delinquent. At the same time, the youth who reported having more delinquent friends were 1.801 times more likely to report having committed the minor delinquency such as group fight, carrying a weapon and vandalism. Finally, Model 4 explained 18.39% of variation in the Chinese youth minor delinquency, increasing by 7.83% compared with the self-control-alone model (Model 3).

Discussion & Conclusion

Although scholarly interest in empirically assessing self-control theory continues to grow, one shortcoming of the literature is that it has rarely been tested utilizing international data (Cretacci et al., 2009; Teasdale & Silver, 2009; Vazsonyi & Huang, 2010). To the best of our knowledge, there has been little literature available on evaluating the utility of the perspective in the distinctive Chinese culture. In this regard, only three empirical studies were conducted in mainland China. Based on a large school-based sample collected in Hangzhou, China, the current study joins this line of research to explore the generalizability of self-control in explaining adolescent risky behavior and minor delinquency. Specifically, the present study focuses on the explanatory power of respectively self-control theory (Gottfredson & Hirschi, 1990) and social bonding (or social control) theory (cf. Hirschi, 1969). Although we attempted to operationalize Hirschi's original (1969) four dimensions of social control (social bonding) theory, using measures for, respectively, attachment (family, school and neighborhood), involvement (in conventional activities), commitment (to school), and beliefs (support of pro-violent values), it is possible that the operationalization of these concepts was less than optimal for the Chinese context. Nonetheless, our results appear sufficiently robust to warrant a number of observations.

Our first noteworthy observation pertains to the empirical findings regarding the measure of self-control. The dimensionality of Grasmick et al.'s self-control measure has been an ongoing debate that continues to generate a fair amount of discussion in the literature. The existing empirical evidence has supported both unidimensional and multidimensional models. For example, Longshore, Turner and Stein (1996) found support for multidimensionality—six dimensions or factors as originally hypothesized and specified by Gottfredson and Hirschi. Similarly, using a sample of 335 late adolescents from a medium city in Japan, Vazsonyi et al. (2004) tested both unidimensional and multidimensional models in their investigation and found that self-control was best represented as a multidimensional construct. In contrast, the results reported by Piquero and Rosay (1998) on survey data provided evidence supporting unidimensionality. In addition, Longshore, Chang and Messina (2005) argued that self-control can defensibly be analyzed as a unidimensional construct. Similar results were reported by Marshall and Enzmann (2012) in their analysis of the dimensionality of the Grasmick et al. self-control scale for the 30-country ISRD-2 sample.

In the current study, the short version of Grasmick et al.'s self-control measure was utilized. More specifically, 12 items tapping into impulsivity, risk-seeking, self-centeredness, and temper were included in the survey questionnaire. The results from both exploratory factor analysis (data-driven) and confirmatory factor analysis (theory-driven) suggest that the one-factor model was a better fit to the data than the four-factor model. Stated differently, respondents in the sample did not make a distinction between impulsiveness, risk-seeking, self-centeredness, and short temper; it seems to them that all 12 items collectively reflect a unitary measure of self-control.

Our second noteworthy observation concerns the hypothesized relationship between self-control measure and various forms of adolescent problem behaviors. According to Gottfredson and Hirschi (1990), the self-control construct holds a central place of importance in explaining crime and deviance across different cultural contexts. Our findings suggest that self-control has a significant inverse relation with both adolescent minor risky behavior (smoking and drinking) and minor delinquency (weapon carrying, group fighting, and vandalism) in China, a nation that is sharply different from the Western countries in terms of culture, tradition, and social settings. It is important to note that the finding here is consistent with prior studies conducted in Western countries, including the analysis of ISRD-2 data

(Junger-Tas, Enzmann, Steketee, & Marshall, 2012). Furthermore, the self-control-deviance relationship detected in our sample is in line with Vazsonyi et al.'s (2004) findings derived from a sample in Japan, whereas it is a contradiction with the three studies conducted in mainland China (e.g., Cretacci et al., 2009, 2010; Wang et al., 2002). In addition, the effects of self-control remain strong, even after adding the social bonding factors and controlling for the demographics. It is interesting to note that only two social bond measures, namely beliefs and family attachment, stood out as significant predictors of risky behavior and minor delinquency, respectively. The comparison of the standardized coefficients between self-control and social bond factors (beliefs and family attachment) indicates that self-control scale carried more weight than social bond measures in its relative predictive power.

Finally, self-control is more powerful in explaining delinquency than risky behavior such as drinking and smoking. The variance explained was almost doubled in the minor delinquency model compared with the risky behavior model (18.39% vs. 9.68%). We speculate that this might be explained by culturally related factors. It has been observed that small-to-moderate amounts of alcohol consumption by male teens on social and ceremonial occasions are generally not considered as risky behavior in China. Parents often hold permissive attitudes toward underage drinking in family gatherings. The widespread cultural acceptance of underage drinking might have weakened the predictive power of self-control on alcohol consumption among teens. In spite of the fact that the Chinese government banned underage drinking in 2006, the influence of this cultural acceptance has not yet faded away.

This study contributes to our understanding of the generalizability of self-control theory in China, but it does have some noteworthy limitations. First, as previously mentioned, the self-control items employed in the current study are from the shortened version of Grasmick et al.'s (1993) scale. It will be interesting to have the full version of 24 items to examine if the unidimensional model remains valid by using the technique of confirmatory factor analysis. Second, in addition to the dimensionality issue, a second topic related to self-control theory that generates a fair amount of debates in the literature is the gender issue (e.g., Vazsonyi et al., 2004). In the current study, the dimensionality issue by gender was not examined. In other words, it is unclear if there is any difference between male and female students in the dimensionality of self-control measure. Third, the results derived from our sample indicated that only moderate amount of variation in the dependent variables is explained by self-control and social bond perspectives. More theoretical explanations, such as general strain theory, may be considered for future investigation. Finally, the current study is based on a school sample from Hangzhou, China and no comparative tasks were carried out with the data collected in the Western countries. Given the fact that the ISRD-2 survey has been implemented in 30 countries, future research may explore the comparative aspect of self-control across different cultures to examine the explanatory power of self-control theory.

APPENDIX: Measurement of Self-Control

Self-Control Items

- I act on the spur of the moment without stopping to think
- I do whatever brings me pleasure here and now, even at the cost of some distant goal
- I'm more concerned with what happens to me in the short run than in the long run
- I like to test myself every now and then by doing something a little risky
- Sometimes I will take a risk just for the fun of it
- Excitement and adventure are more important to me than security
- I try to look out for myself first, even if it means making things difficult for other people
- If things I do upset people, it's their problem not mine
- I will try to get the things I want even when I know it's causing problems for other people
- I lose my temper pretty easily
- When I am really angry, other people better stay away from me
- When I have a serious disagreement with someone, it's usually hard for me to talk calmly about it without getting upset

Acknowledgment

The authors would like to thank the Zhejiang Provincial Juvenile Delinquency Institute for its generous support and assistance with data collection in Hangzhou, China.

Declaration of Conflicting Interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Notes

1. The minor delinquent conducts are chosen as the dependent variables because there is a lack of sufficient variations among the more serious delinquency (e.g., illegal drug use, drug dealing, car theft, and robbery).
2. The high return rates of self-report surveys are commonly seen in school-based studies on juvenile delinquency conducted in China. For example, in their comparative study on adolescent problem behavior in China and the United States, Jessor et al. (2003) found that the response rate of 98% obtained from the Chinese sample was much higher than that of the U.S. sample (74%). Similar numbers were also reported by Greenberger, Chen, Beam, Whang and Dong (2000). We do not know the extent to which this difference in the response rates leads to biased responses.
3. In response to the debate of unidimensional versus multidimensional latent trait of self-control, additional analyses were performed by using confirmatory factor analysis to examine if a single-factor model or a 4-factor model would fit the data of 12-item measure of self-control. The results indicated that a single-factor model fit better than a 4-factor model based on the model fit indices (i.e., TLI [the Tucker Lewis Index], CFI [the comparative fit index], and RMSEA [the root means square error of approximation]). In addition, according to the model modification indices in the 4-factor model, some indicators of self-centeredness can be loaded on impulsiveness, risk-seeking, and temper, suggesting that the 12 indicators in this study virtually represent a single, unidimensional latent trait of self-control. Confirmatory factor analysis was conducted by using the Mplus version 6.

References

Arneklev, B. J., Grasmick, H. G., Tittle, C. R., & Bursik, R. J. Jr. (1993). Low self-control and imprudent behavior. *Journal of Quantitative Criminology*, *9*, 225–247.

Baron, S. W. (2003). Self-control, social consequences, and criminal behavior: Street youth and the general theory of crime. *Journal of Research in Crime and Delinquency*, *40*, 403–425.

Brownfield, D., & Sorenson, A. M. (1993). Self-control and juvenile delinquency: Theoretical issues and an empirical assessment of selected elements of a general theory of crime. *Deviant Behavior*, *4*, 243–264.

Bruinsma, G. I. N. (1994). *The test-retest reliability of the ISRD self-report instrument*. Paper presented at the Annual Conference of the American Society of Criminology, Miami, FL.

Burton, V. S. Jr., Cullen, F. T., Evans, T. D., Alarid, L. F., & Dunaway, R. G. (1998). Gender, self-control, and crime. *Journal of Research in Crime & Delinquency*, *35*, 123–147.

Cao, L. (2007). Returning to normality: Anomie and crime in China. *International Journal of Offender Therapy and Comparative Criminology*, *51*, 40–51.

Caspi, A., Moffitt, T. E., Silva, P., Stouhamer-Loeber, M., Krueger, R. F., & Schmutte, P. S. (1994). Are some people crime prone? Replications of the personality-crime relationship across countries, genders, races, and methods. *Criminology*, *32*, 163–195.

Chan, H. M., & Lee, R. P. L. (1995). Hong Kong families: At the crossroads of modernism and traditionalism. *Journal of Comparative Family Studies*, *26*, 83–99.

Chen, C., Lee, S., & Stevenson, H. W. (1996). Academic achievement and motivation of Chinese students: A cross-national perspective. In S. Lau (Ed.), *Growing up the Chinese way: Chinese child and adolescent development* (pp. 69–91). Hong Kong: The Chinese University.

Cheung, N. W. T., & Cheung, Y. W. (2008). Self-control, social factors, and delinquency: A test of the general theory of crime among adolescents in Hong Kong. *Journal of Youth Adolescence*, *37*, 412–430.

Cochran, J. K., Wood, P. B., Sellers, C. S., Wilkerson, W., & Chamlin, M. B. (1998). Academic dishonesty and low self-control: An empirical test of a general theory of crime. *Deviant Behavior*, *19*, 227–255.

Cretacci, M. A., Ding, F., & Rivera, C. J. (2010). Traditional and bond measures

of self-control and their impact on deviance among Chinese university students. *International Journal of Criminal Justice Sciences*, *5*, 220–238.

Cretacci, M. A., Rivera, C. J., & Ding, F. (2009). Self-control and Chinese deviance: A look behind the bamboo curtain. *International Journal of Criminal Justice Sciences*, *4*, 131–143.

De Li, S. (2004). The impacts of self-control and social bonds on juvenile delinquency in a national sample of midadolescents. *Deviant Behavior*, *25*, 351–373.

Evans, T. D., Cullen, F. T., Burton, V. S. Jr., Dunaway, R. G., & Benson, M. L. (1997). The social consequences of self-control: Testing the general theory of crime. *Criminology*, *35*, 475–501.

Gibbs, J. J., Giever, D., & Martin, J. (1998). Parental-management and self-control: An empirical test of Gottfredson and Hirschi's general theory. *Journal of Research in Crime and Delinquency*, *35*, 42–72.

Gibson, C. L., Wright, J. P., & Tibbetts, S. G. (2000). An empirical assessment of the generality of the general theory of crime: The effects of low self-control on social development. *Journal of Crime and Justice*, *23*, 109–134.

Gottfredson, M. R. (2006). The empirical status of control theory in criminology. In F. T. Cullen, J. P. Wright & K. R. Blevins (Eds.), *Taking stock: The status of criminological theory* (pp. 77–100). New Brunswick, NJ: Transaction.

Gottfredson, M. R., & Hirschi, T. (1990). *A general theory of crime*. Stanford, CA: Stanford University Press.

Grasmick, H. G., Tittle, C. R., Bursik, R. J. Jr., & Arneklev, B. J. (1993). Testing the core empirical implications of Gottfredson and Hirschi's general theory of crime. *Journal of Research in Crime and Delinquency*, *30*, 5–29.

Greenberger, E., Chen, C., Beam, M., Whang, S., & Dong, Q. (2000). The perceived social contexts of adolescents' misconduct: A comparative study of youths in three cultures. *Journal of Research on Adolescence*, *10*, 365–388.

Hirschi, T. (1969). *Causes of delinquency*. Berkeley: University of California Press.

Hirschi, T. (2004). Self-control and crime. In R. F. Baumeister & K. D. Vohs (Eds.), *Handbook of self-Regulation: Research, theory, and applications* (pp. 537–552). New York, NY: Guilford Press.

Hwang, S., & Akers, R. L. (2003). Substance use and Korean adolescents: A cross-cultural test of social learning, social bonding, and self-control theories. In R. L. Akers & G. F. Jensen (Eds.), *Social learning theory and the explanation of crime* (pp. 39–63). New Brunswick, NJ: Transaction.

Jessor, R., Turbin, M. S., Costa, F. M., Dong, Q., Zhang, H., & Wang, C. (2003). Adolescent problem behavior in China and the United States: A cross-national study of psychosocial protective factors. *Journal of Research on Adolescence*, *13*, 329–360.

Junger, M., West, R., & Timman, R. (2001). Crime and risky behavior in traffic: An example of cross-situational consistency. *Journal of Research in Crime and Delinquency*, *38*, 439–459.

Junger-Tas, J., Enzmann, D., Steketee, M., & Marshall, I. Haen (2012). Concluding observations: The big picture. In J. Junger-Tas, I. Haen Marshall, D. Enzmann, M. Killias, M. Steketeee & B. Gruszczynska, *The many faces of youth crime. Contrasting theoretical perspectives on juvenile delinquency across countries and cultures* (pp.329–353). New York, NY: Springer.

Keane, C., Maxim, P. S., & Teevan, J. J. (1993). Drinking and driving, self-control, and gender: Testing a general theory of crime. *Journal of Research in Crime and Delinquency*, *30*, 30–46.

LaGrange, T., & Silverman, R. A. (1999). Low self-control and opportunity: Testing the general theory of crime as an explanation for gender differences in delinquency. *Criminology*, *37*, 41–72.

Li, X., Fang, X., & Stanton, B. (1999). Cigarette smoking among schoolboys in Beijing, China. *Journal of Adolescence*, *22*, 621–625.

Li, X., Fang, X., Stanton, B., Feigelman, S., & Dong, Q. (1996). The rate and pattern of alcohol consumption among Chinese adolescents. *Journal of Adolescent Health*, *19*, 353–361.

Lilly, J. R., Cullen, F. T., & Ball, R. A. (2007). *Criminological theory: Context and consequences* (4th ed.).Thousand Oaks, CA: SAGE.

Longshore, D., Chang, E., & Messina, N. (2005). Self-control and social bonds: A combined control perspective on juvenile offending. *Journal of Quantitative Criminology*, *21*, 419–437.

Longshore, D., Turner, S., & Stein, J. (1996). Self-control in a criminal sample: An examination of construct validity. *Criminology*, *3*, 209–228.

Mason, W. A., & Windle, M. (2002). Gender, self-control, and informal social control in adolescence: A test of three models of the continuity of delinquent behavior. *Youth and Society, 33*, 497–514.

Marshall, I. H., & Enzmann, D. (2012). The generalizability of self-control theory. In J. Junger-Tas, I. H. Marshall, D. Enzmann, M. Killias, M. Steketee, & B. Gruszczynska, *The many faces of youth crime: Contrasting theoretical perspectives on juvenile delinquency across countries and cultures* (pp. 285–326). New York, NY: Springer.

Marshall, I. H., & Webb, V. J. (1990). *Omaha IRSD Pilot Study*. Paper presented at the meeting of International Self-report Delinquency Project, The Hague, The Netherlands.

Marshall, I. H., & Webb, V. J. (1994). Self-reported delinquency in a midwestern American city. In J. Junger-Tas, G. Terlouw & M. W. Klein (Eds.), *Delinquent behavior among young people in the Western world: First results of the International Self-report Delinquency Study* (pp. 319–342). Amsterdam, The Netherlands: Kugler Publications.

Piquero, A. R., MacIntosh, R., & Hickman, M. (2000). Does self-control affect survey response? Applying exploratory, confirmatory, and item response theory

analysis to Grasmick et al.'s self-control scale. *Criminology*, *38*, 897–929.

Piquero, A. R., & Rosay, A. B. (1998). The reliability and validity of Grasmick et al.'s self-control scale: A comment on Longshore et al. *Criminology*, *36*, 157–173.

Polakowski, M. (1994). Linking self- and social control with deviance: Illuminating the structure underlying a general theory of crime and its relation to deviant activity. *Journal of Quantitative Criminology*, *10*, 41–78.

Pratt, T. C., & Cullen, F. T. (2000). The empirical status of Gottfredson and Hirschi's general theory of crime: A meta-analysis. *Criminology*, *38*, 931–964.

Romero, E., Gomez-Fraguela, A. J., Luengo, M. A., & Sobral, J. (2003). The self-control construct in the general theory of crime: An investigation in terms of personality psychology. *Psychology Crime & Law*, *9*, 61–86.

Teasdale, B., & Silver, E. (2009). Neighborhoods and self-control: Toward an expanded view of socialization. *Social Problems*, *56*, 205–222.

Tittle, C. R., & Botchkovar, E. V. (2005). Self-control, criminal motivation and deterrence: An investigation using Russian respondents. *Criminology*, *43*, 307–353.

Unger, J. B., Yan, L., Chen, X., Jiang, X., Azen, S., Qian, G., & Johnson, A. (2001). Adolescent smoking in Wuhan, China, Baseline data from the Wuhan smoking prevention trial. *American Journal of Preventive Medicine*, *21*, 162–169.

Vazsonyi, A. T., Clifford Wittekind, J. E., Belliston, L. M., & Van Loh, T. D. (2004). Extending the general theory of crime to "The East:" Low self-control in Japanese late adolescents. *Journal of Quantitative Criminology*, *20*, 189–216.

Vazsonyi, A. T., & Huang, L. (2010). Where self-control comes from: On the development of self-control and its relationship to deviance over time. *Developmental Psychology*, *46*, 245–257.

Vazsonyi, A. T., Pickering, L. E., Junger, M., & Hessing, D. (2001). An empirical test of a general theory of crime: A four-nation comparative study of self-control and the prediction of deviance. *Journal of Research in Crime and Delinquency*, *38*, 91–131.

Wang, D. (2006). The study of juvenile delinquency and juvenile protection in the People's Republic of China. *Crime and Justice International*, *22*, 4–14.

Wang, G. T., Qiao, H., Hong, S., & Zhang, J. (2002). Adolescent social bond, self-control, and deviant behavior in China. *International Journal of Contemporary Sociology*, *39*, 52–68.

Wei, Z., Homel, R., Prichard, J., & Xu, J. (2004). Patterns of juvenile offending in Shanghai and Brisbane. *Australian and New Zealand Journal of Criminology*, *37*, 32–51.

Wood, P. B., Pfefferbaum, B., & Arneklev, B. J. (1993). Risk-taking and self-control: Psychological correlates of delinquency. *Journal of Criminal Justice*, *16*, 111–130.

Wright, B. R. E., Caspi, A., Moffitt, T. E., & Silva, P. A. (1999). Low self-control, social bonds, and crime: Social causation, social selection, or both? *Criminology*, *37*, 479–514.

Yang, K. S. (1995). Chinese social orientation: An integrative analysis. In T. Y. Lin, W. S. Tseng & E. K. Yeh (Eds.), *Chinese societies and mental health* (pp. 19-39). Hong Kong, Oxford University Press.

Zhang, S., Benson, T., & Deng, X. (2000). A test-retest reliability assessment of the International Self-report Delinquency Instrument. *Journal of Criminal Justice*, *28*, 283–295.

Author Biographies

Yi-Fen Lu is a doctoral student in the College of Criminal Justice at Sam Houston State University. She earned her MA in Criminal Justice from Sam Houston State University in 2010. Her research interests include public perceptions of crime and criminal justice, biosocial criminology, gene-environmental interactions, and quantitative methodology.

Yi-Chun Yu is a doctoral student in the College of Criminal Justice at Sam Houston State University. She earned her MA in Criminal Justice from Sam Houston State University in 2010. Her areas of research interests are criminological theory, crime mapping, policing, and quantitative methodology.

Ling Ren is an associate professor in the College of Criminal Justice at Sam Houston State University. She earned her PhD from University of Nebraska at Omaha in 2006. Her research interests include policing, cross-national research on crime and delinquency, and the applications of quantitative methodology.

Ineke Haen Marshall is a professor at the School of Criminology and Criminal Justice and the Department of Sociology and Anthropology at Northeastern University. Her research interests are comparative criminology, self-report survey research, criminal careers, and comparative penal policy.

Appendix C

How to Use a Statistical Package

With the assistance of Lisa M. Gilman and with contributions by Joan Saxton Weber

Computers and statistical software such as the Statistical Package for the Social Sciences (SPSS) make complex statistical computations simple and fast. SPSS is one of the most popular comprehensive statistical software packages used in the social sciences. You can use it to calculate a great many statistics and to create charts and tables for presentations with just a few clicks of the mouse.

This appendix provides a basic introduction to SPSS. Even if you are unfamiliar with computers or apprehensive about statistics, you will find that SPSS for Windows is very user friendly. Please bear in mind that all of the examples use version 13 or 14 of SPSS, with data from the 2004 version of the General Social Survey (GSS); if you are using a different SPSS version or another year of the GSS, you may find some slight differences in procedures or answers.

C.2 Basic Procedures

To start SPSS for Windows, click or double-click on the SPSS icon using the left mouse button. If you are unable to locate the SPSS icon, click the Start button, then click on Programs, and then click on SPSS.

A new screen will open with a "What would you like to do?" box superimposed on the screen. For now, click on the Cancel button in the "What would you like to do?" box, to get it out of the way. You are now looking at the SPSS Data Editor window. This screen is where data to be analyzed are entered or data files that have already been created (datasets) are loaded.

To access the data for this appendix, click:

File—Open—Data (or just click on the open folder icon).

Select (highlight) the GSS2004mini file. If the GSS2004mini file hasn't been transferred to your hard disk, you can find it on the CD-ROM that came with this text.

Exhibit C.1 SPSS Data Editor

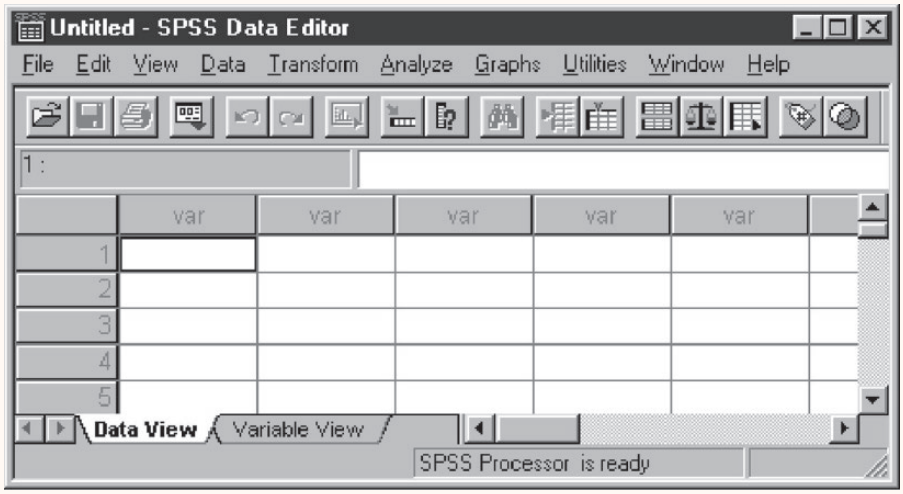

The Data View in the Data Editor consists of columns and rows, with each column representing a different variable—their names are at the top—and each row representing one case or "observation" (Exhibit C.1). If you are using the student version of SPSS, you are limited to no more than 50 variables and 1,500 observations. (The GSS2004mini dataset has 1,404 cases, or observations, and 48 variables. You can confirm this by moving the scroll bars on the bottom and right of the Data View). There are no variable/observation limits in the commercial version.

Another screen you should be familiar with is the Variable View screen. To access the variable view screen, click on the Variable View tab at the bottom left of the Data Editor (not available in early versions of SPSS for Windows). The Variable View screen contains a list of all the variables included in the dataset and their characteristics. Each row corresponds to a single variable. Variable characteristics are indicated at the top of each column. When you are in the Variable View mode, you can create, edit, or view variable information. Now click on the Data View tab to return to the data view mode.

Looking at the Data

There are several ways to learn more about the variables and numbers that you see in the Data View screen. The names of the variables are listed across the top of the screen. You can also see the list of variables in the GSS2004mini file by clicking on the variable list icon at the top of the Data Editor screen. When you click on any variable in the resulting window, you can see the details that show how it was coded. But the easiest way to learn about the variables in the dataset is to switch from the Data View to the Variable View screen by clicking on the tab in the lower left.

You can tell what some variables are just by their SPSS variable names, like SEX. However, many variables are responses to specific questions and we can't tell just what they are on the basis of the shorthand variable name. Instead, we can inspect the variable label in the "Label" column corresponding to that variable in the Variable View screen (you may need to widen the Label column by dragging the separator line at the top of the column with your mouse).

SPSS understands numbers much better than text, so all the answers to each GSS question have been coded. This means that, for every variable, each response category has been assigned a numerical value. For some variables, such

as AGE and EDUC, the number you see is simply the number of years. For other variables, such as SEX and RACE, each number corresponds to a particular response category. For example, go to the Variable View screen and click on the "Values" cell in the row corresponding to the variable SEX. Now click on the gray box in that cell and you can see the labels for specific codes, which show that men are coded "1" and women are coded "2."

Values that stand in for missing data are indicated in the "Missing" column. Special labels identifying the reason for missing data, such as DK for Don't Know or NA for No Answer, also appear in the Value Labels column. For ABRAPE, the variables dialog box shows that 8 corresponds with DK (Don't Know) and 9 corresponds with NA (Not Available). Therefore, for the variable ABRAPE, both 8 and 9 are missing values.

For most statistical calculations, SPSS ignores missing values and calculates the statistic based on the responses of the respondents who answered the question. In general, however, you should make sure missing values are not inadvertently included in your analysis. Common values used to identify missing data are –1, 0, 8, and 9. If the variable is two digits wide, such as AGE, the missing values are usually 97, 98, and/or 99.

What if you have data for new cases to enter, or you have new variables that you need to define? You can simply enter the data in the Data View and enter the labels and missing value codes in the Variable View. Click on File—Save after you are sure that everything is correct.

Before you proceed to the next step, make one minor adjustment in the display options. Click:

Edit—Options

The Options dialog box will open. You will see a series of "tabs" along the top of the window. You should be looking at the General options (if not, click on the General tab). On the upper left side of the General options screen, you will see the Variable Lists options.

Under Variable Lists, click the radio button next to Display Names (to mark it with a small dot). This will display the convenient short variable names rather than the descriptive variable label in some of the selection boxes.

2 Univariate Statistics

Now that you have an open data file and are somewhat familiar with the data itself, you can begin exploring the data statistically (statistics are discussed in Chapter 14).

Frequencies

A *frequency distribution* is a table that displays how many and what percentage of observations fall into given categories for a variable of interest. In other words, a frequency distribution tells you how many people said yes, how many people said no, and so on. The purpose of obtaining a frequency distribution is to summarize the data so that they are easy to understand. The variable ABRAPE (pregnant as a result of rape) will serve as a good example for a frequency distribution. Click:

Analyze—Descriptive Statistics—Frequencies

The frequencies dialog box will open, as shown in Exhibit C.2. Choose ABRAPE by clicking on it from the list on the left. Click the arrow in the center of the dialog box to move the ABRAPE to the "variable(s)" box on the right. Then click OK.

After SPSS has processed the command, the results appear in a new window titled "Output1–SPSS Viewer." The Output window has two panes. The left pane is referred to as the "output navigator." The output navigator contains an outline of everything you ask SPSS to do from the beginning of the session. It allows you to easily refer back to any given table or graph. To go to a specific table, you find it in the output navigator, click on the table you want, and it will be displayed in the right pane. As indicated, the right pane contains the actual output for the commands you gave SPSS. It is often referred to as the "output" or the "results."

In the example above, SPSS created a frequency distribution for the variable ABRAPE. SPSS produces two boxes (Exhibit C.3). The first box just displays the valid number of cases (*N*) and the total number of missing cases for the variable of interest. The second box contains the actual frequency distribution. If the entire table is not visible, use the

Exhibit C.2 The Frequencies Dialogue Box

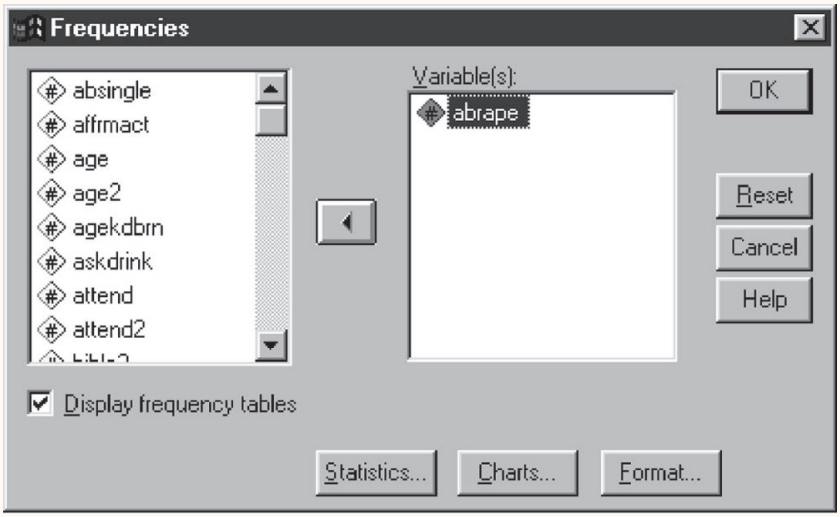

Exhibit C.3 SPSS Frequencies Output

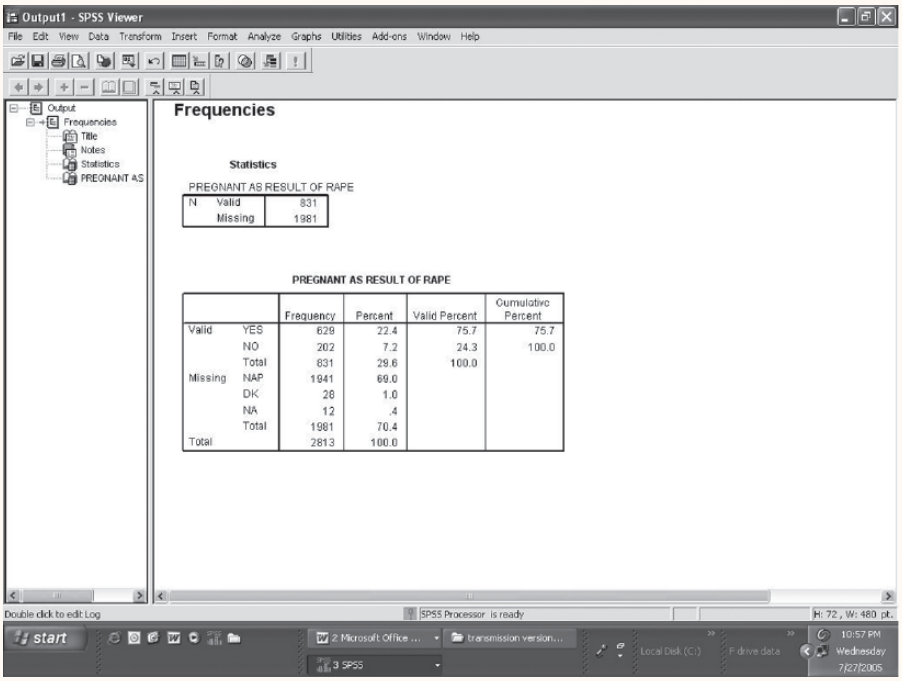

scroll bar on the right to move down the output window. The distribution shows that about 76% of respondents believe abortion is permissible when a pregnancy is the result of rape, whereas 24% of respondents do not believe that abortion is permissible in this circumstance.

Descriptive Statistics

If you need to calculate the mean, median, or standard deviation of an interval or ratio level variable, you can do so simply with the Descriptive Statistics procedure. Click:

Analyze—Descriptive Statistics—Descriptives

and click the EDUC variable over into the Variables window before clicking OK.

Graphing

Another method of describing the distribution of a single variable is through the use of a histogram. A histogram is useful for graphically displaying the distribution of a given variable. Histograms are useful when the variable(s) you are interested in are continuous in nature (interval/ratio) and have a large number of categories. For example, if you want to look at the number of respondents within certain age groups, it is much simpler to visually display the data than to look at counts. Let's look at how age is distributed in the GSS200mini data included with this book. Click:

Graphs—Histogram

Click on the AGE variable in the variable list and move it to the Variable box and click OK (Exhibit C.4).

Exhibit C.4 Selecting a Variable for a Histogram

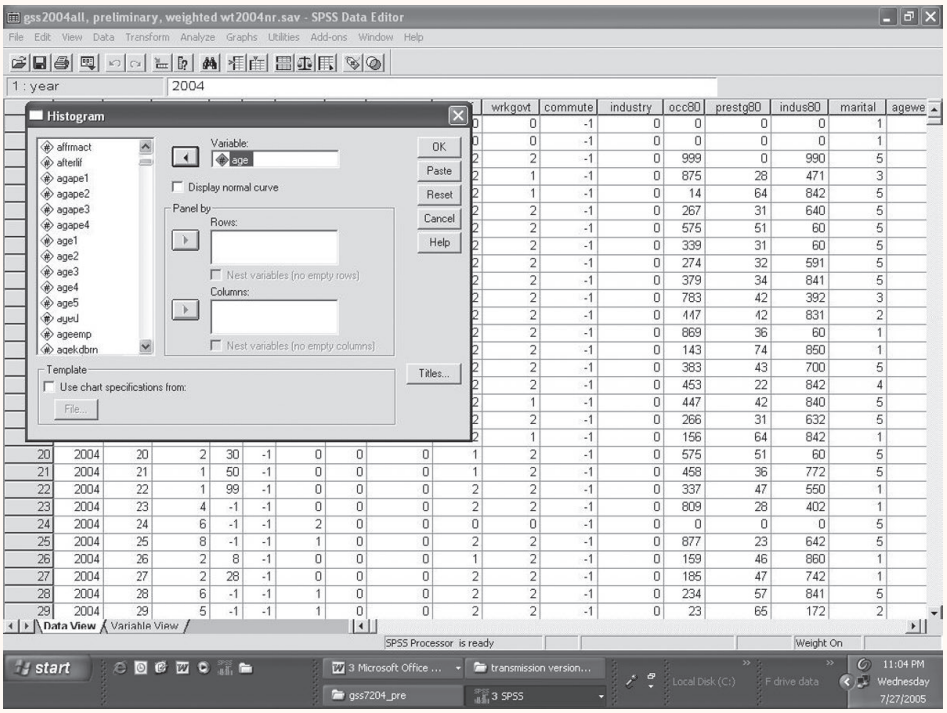

When you look at the histogram (Exhibit C.5), you can see that the age distribution has a slight positive skew. The majority of respondents are under the age of 60. The statistics to the right of the histogram indicate that the mean age of the respondents for this sample is approximately 45 years old with a standard deviation of about 16 years.

Exhibit C.5 **Histogram of Age**

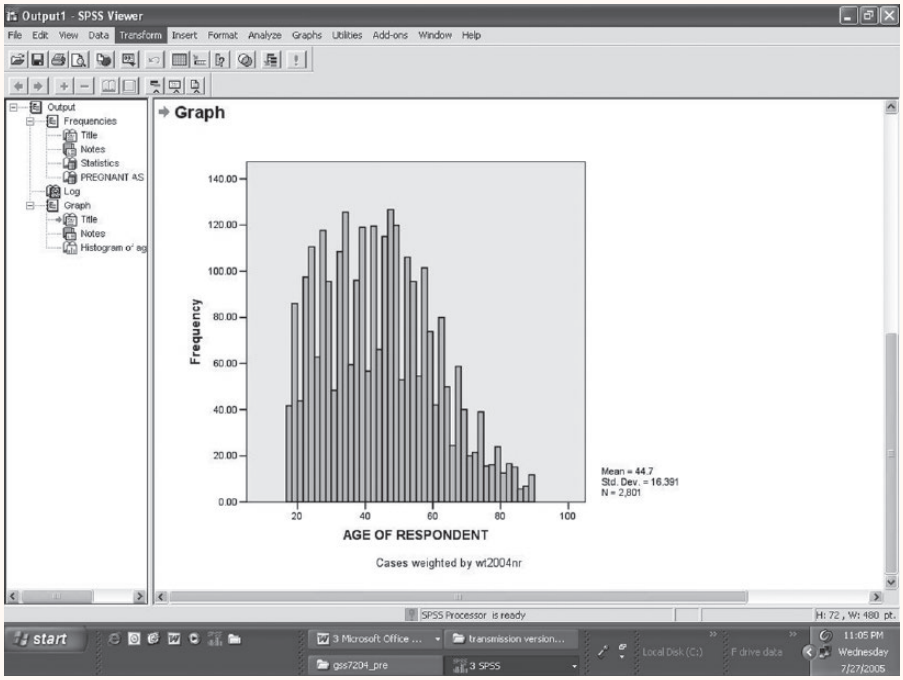

A similar graphic device for nominal or ordinal variables is the bar chart. You can create a bar chart for Race by clicking:

Graphs—Bar

and then highlighting the options Simple and "Summaries for groups of cases" before clicking Define. Now scroll through the variable list to find race, click this into the Category Axis box, and select the Bars Represent . . . % of cases option. Now, click OK. And there you have it.

2 Recoding Variables

In some instances, variables with many categories may be confusing to use and/or interpret. At other times, your research question may make it necessary to limit your analysis to certain categories. The variable EDUC (highest year of education completed) is a good example. You seldom want to know if people with 11 years of education are different from people who have 10 years of education. It might be better to examine differences of opinion among high school graduates versus college graduates.

To create these groups, it will be necessary to recode the EDUC variable. When recoding a variable, be sure to check the missing values and to look at the minimum and maximum values for valid responses with a frequency distribution.

Let's recode EDUC into the following four categories:

Those who have 0 to 11 years of education are put into category 1.

Those who have exactly 12 years of education are put into category 2.

Those who have 13 to 15 years of education are put into category 3.

Those who have 16 or more years of education are put into category 4.

In this example, you will collapse the original 21 categories (years of education from 0 to 20) into the 4 categories indicated above. To do this, you will create a new variable called ED4CAT. (You should always recode variables into "different variables" to retain the original variable.) Click:

Transform Recode—Into Different Variables

This takes you to the first recode dialog box, in which you tell SPSS the name of the old variable you are recoding from (EDUC) and the name of your new recoded variable (ED4CAT) (see Exhibit C.6). Using the scroll bar, find EDUC in the variable list on the left and move it to the "Numeric Variable—Output Variable" box. Under the "Output Variable" section, type the name of the new variable (ED4CAT) and a variable label, such as "Recoded Years of Education," and then click Change.

Exhibit C.6 **Selecting Variable for Recoding**

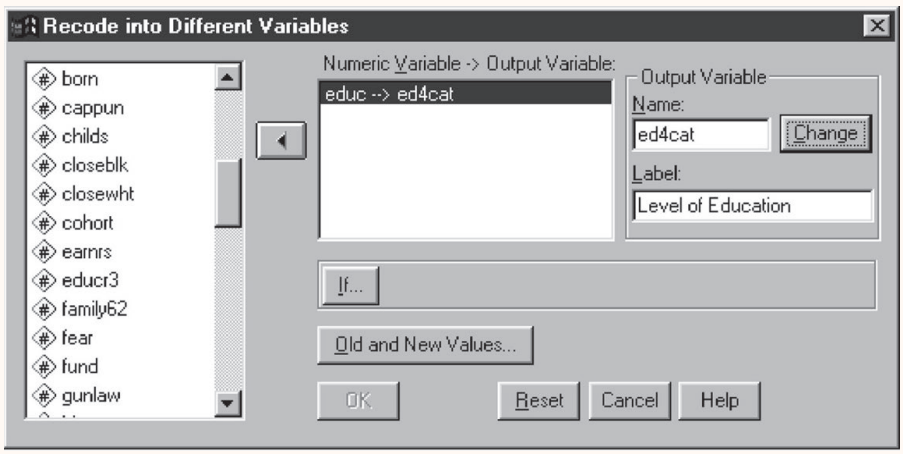

After doing this, click the "Old and New Values . . ." button. A new dialog box will open (Exhibit C.7). This second box is where you tell SPSS into what categories you want to recode the original data. To do this, you identify the original ("old") values of the variable on the left side of the box, and you put the values for your new variable on the right side of the box. When you recode a single value, use the "Value" box under the "Old Value" heading; when you recode a range of values to a single new value, specify the range under one of the "Range" options under "Old Value." Click the Add button to proceed with each conversion. Be sure to recode any missing values (you could just recode "System- or user-missing" under "Old Value" to "System-missing" under "New Value").

After you have created the desired categories, click OK. You will be returned to the first recode box and will need to click OK again. The newly created variable (ED4CAT) will be located at the end of the data matrix. Obtain a frequency distribution for ED4CAT. You will notice that there are no value labels. Therefore, to make this output easier to read, it will be necessary to attach value labels to the numeric codes. Go back to the Data View window. Scroll to the far right of the data

Exhibit C.7 Recoding Instructions

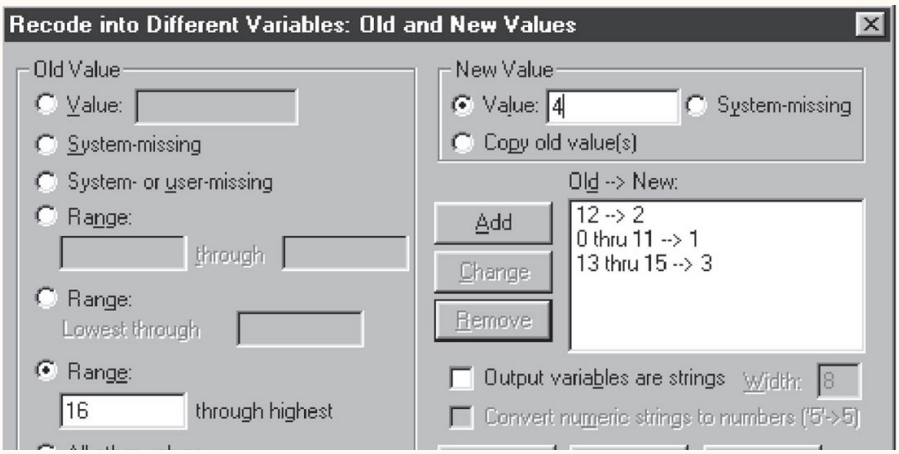

matrix. Double-click the name of the new variable ED4CAT. This will put you in Variable View mode, which will allow you to edit the characteristics of the variable.

Using the right arrow key on the keyboard, go to the "Values" cell. Click on the button located next to the word "None." A box entitled "Value Labels" will appear. Enter each category value and the corresponding value label where indicated (Exhibit C.8). After all value labels are entered, click OK.

Computing a New Variable

Recoding is very useful when you want to alter the format of a given variable. However, sometimes it may be necessary to combine multiple variables. For example, you may have a survey measuring aggressive behavior. To best measure levels of aggressiveness it would make sense to build an additive index from the data so that each respondent has

Exhibit C.8 Defining Value Labels

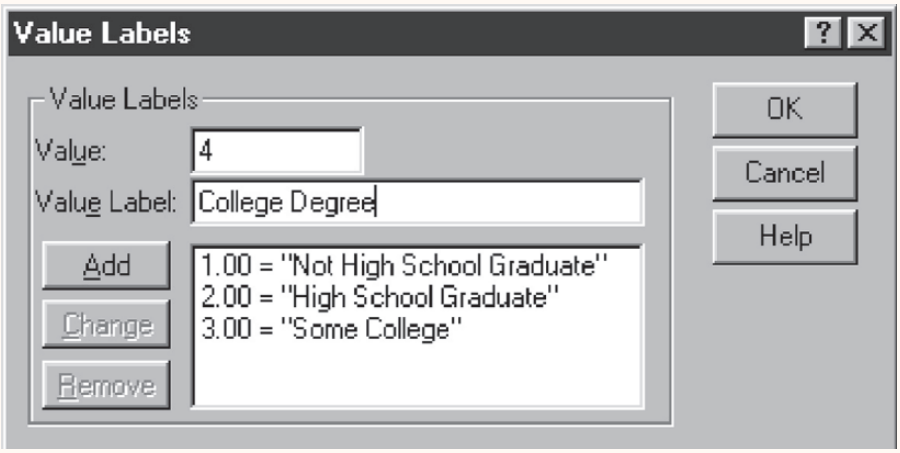

Exhibit C.9 Computing a New Variable

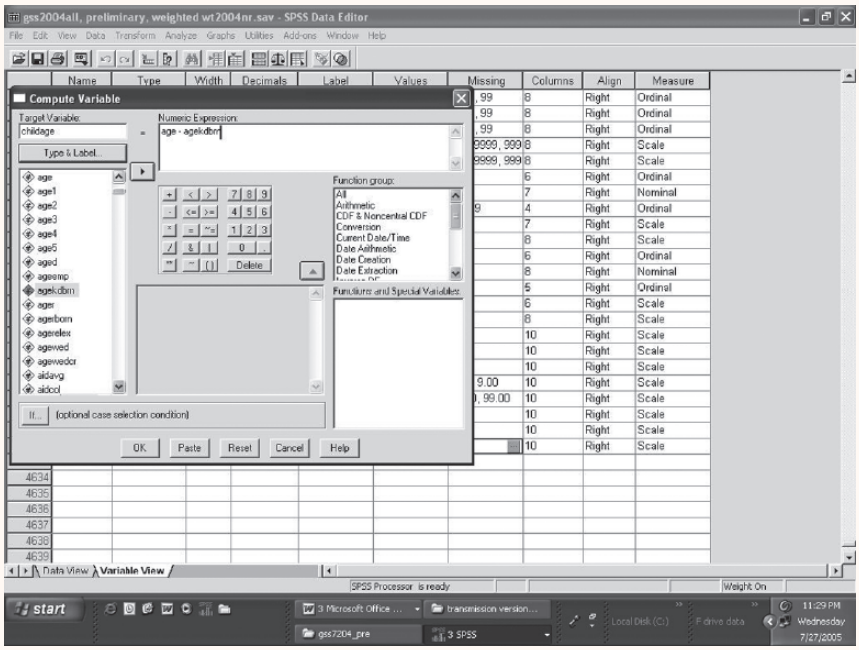

an aggressiveness score (see Chapter 4). To do this, it is necessary to modify the data through the use of the Compute command. This command works well for combining two or more variables and for performing other types of mathematical transformations of the data.

Using the Compute command, let's calculate the age of each respondent's eldest child. Click:

Transform—Compute

The Compute command will create a new variable based on the mathematical functions you select—the target variable. Name the target variable CHLDAGE (Exhibit C.9). To compute this, we need to calculate the difference between the respondents' age and the age of the respondent when his or her first child was born. Therefore, move the AGE variable to the "Numeric Expression" box. Using your mouse or keypad, click on the subtraction symbol. Now, move the AGEKDBRN variable into the "Numeric Expression" box, and click OK. You will find the new variable to the far right of the existing variables in the "Data View." If you generate a histogram for CHLDAGE, you will see that 24 is the mean age at which the first child was born.

If you scroll down the CHLDAGE column in the Data View, you will notice a number of cells with a "." in them. This is SPSS's default for "missing" data. The age of the eldest child could not be calculated for those individuals who did not have any children. Therefore, SPSS uses the "." to indicate that there were no data (or the data were missing) for those cases.

2 Bivariate and Multivariate Statistics

Frequency distributions are good for describing the distribution of the variables in the dataset. However, to examine the relationship between two or more variables, a different set of statistical techniques are necessary. These are referred to as bivariate (two-variable) and multivariate (multiple-variable) statistics.

Crosstabulation

One way of exploring the relationship between two variables is through the use of a contingency table, more commonly referred to as a crosstab. Crosstabs are useful for exploring relationships between categorical variables. For example, we can use a crosstab to explore the relationship between level of education as measured by DEGREE and attitudes toward abortion for any reason (ABANY). Click:

Analyze—Descriptive Statistics—Crosstabs

For this analysis, DEGREE is the independent variable and ABANY is the dependent variable. Therefore, you will want DEGREE in the columns and ABANY in the rows. Locate these variables in the variable list (Exhibit C.10). Using the arrow buttons, place DEGREE and ABANY in their respective boxes.

SPSS automatically computes cell counts (the number of respondents that are "observed" within each category). However, you should also have SPSS calculate percentages for the independent variable to interpret differences. Therefore, click the Cells . . . button, choose the column percentages, and then click Continue to return to the first dialog box. If you have previous knowledge of statistics, you may want SPSS to also calculate a specific statistical test, such as chi-square. To do this, click the Statistics . . . button, make the appropriate choices, then click Continue. Click OK to run the crosstab command. Take a minute to inspect the resulting output (Exhibit C.11) and interpret the table (review Chapter 14 if necessary).

Now repeat this procedure, substituting INCOM98R for DEGREE. This will generate a crosstabulation of support for abortion by respondent's family income.

The results indicate that support for abortion if the woman wants one for any reason increases with education (rising from about 18% to over 65% across the five education categories) and also with family income (an increase of about 24% from the lowest to the highest recoded income categories).

Exhibit C.10 Requesting Crosstabs

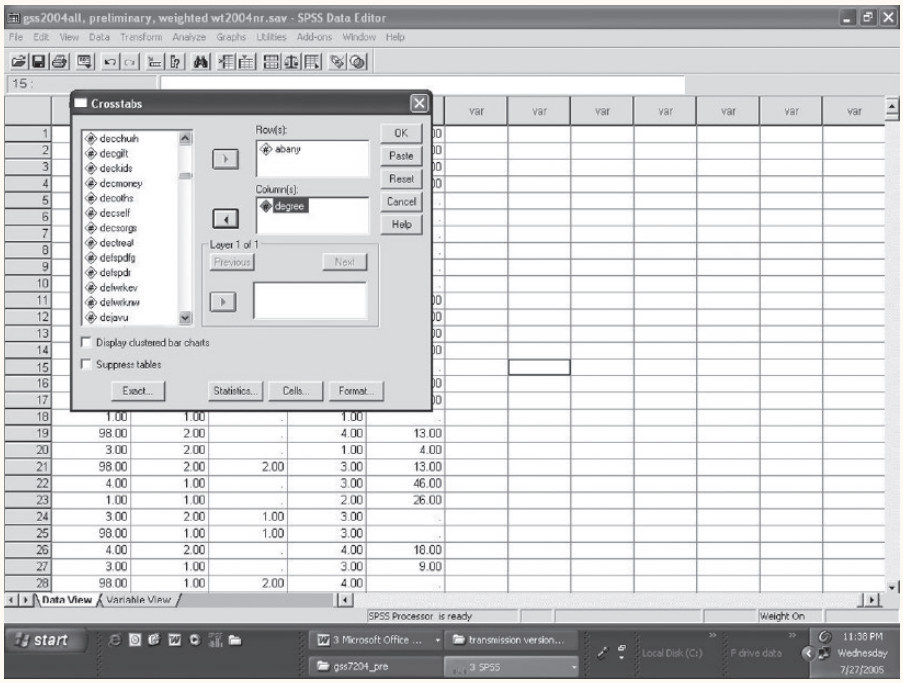

Exhibit C.11 Crosstabulation of ABANY by DEGREE

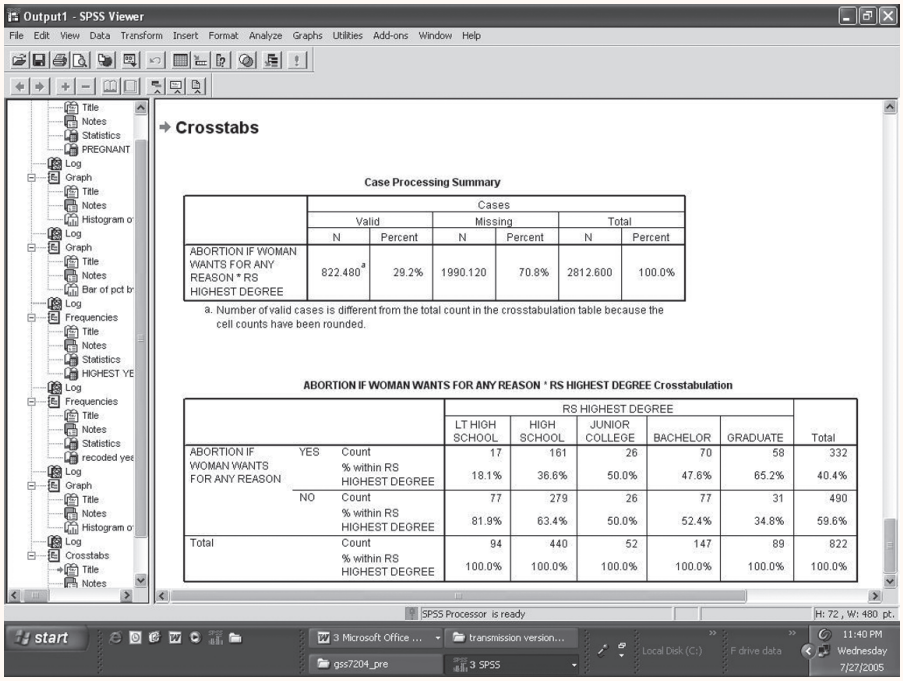

Three-Variable Crosstabs

In the last example, it was determined that there was a relationship between income and support for abortion. Could this relationship be due to an extraneous factor? A three-variable crosstab can be created to assess whether controlling for another variable eliminates the apparent influence of income on support for abortion.

You have already seen that support for abortion increases with education (DEGREE). Because income also increases with degree, and we know that maximum educational attainment most often precedes employment as an adult, it is possible that support for abortion varies with income because both of these variables are, in turn, influenced directly by level of education. If this is the case, the relationship between income and support for abortion would be spurious—that is, it would be due to the extraneous factor of education.

We can control for degree in order to assess the effects of education on the relationship between income and support for abortion. To do this, click:

Analyze—Descriptive Statistics—Crosstabs

You should click ABANY into the rows box and INCOM98R into the columns box in the Crosstabs dialog box (unless they are still there after your last procedure). Select the variable DEGREE from the variables list and move it to the empty box on the bottom titled "Layer 1 of 1." Select the column percentages and the chi-square statistic, if desired, then click OK.

This is a much bigger table to interpret, as it is really four subtables of ABANY by INCOM98R for the four different values of DEGREE. If you compare the percentages of YES responses from the lowest to the highest income level, first for the subtable representing respondents with less than a high school education and then for the other four subtables, you'll no longer find the clear pattern we found of increasing support for abortion as family income

increases. Instead, the support for abortion increases a bit across some, but not all, of the income categories. The results indicate that education largely explains the effect of family income on support for abortion for any reason. You might take a minute to speculate about what might account for this pattern (and if you have had a course in statistics, and you check the chi-square statistic for these tables, you may realize that we cannot reject the possibility that the weak relationship between support for abortion and income is simply due to chance).

Comparing Means

How would you compare mean years of education of whites and minorities? This type of question requires a comparison of means. A comparison of means is useful when you want to look at differences between two or more groups, such as by race or gender.

Click:

Analyze—Compare Means—Means

You must choose two variables: a dependent variable, which must be an interval/ratio level variable (or an ordinal variable that you are treating as interval), and an independent variable, which should be nominal or ordinal. For this example, EDUC is the dependent variable and RACE is the independent variable (Exhibit C.12).

After you identify your dependent and independent variables, click OK. SPSS will produce a table that shows the mean years of education for each category within the RACE variable.

The results indicate that the mean years of education for whites is about one year higher than the mean number of years of education for blacks, but it is about the same as for individuals in the "other" category.

Exhibit C.12 Requesting a Comparison of Means

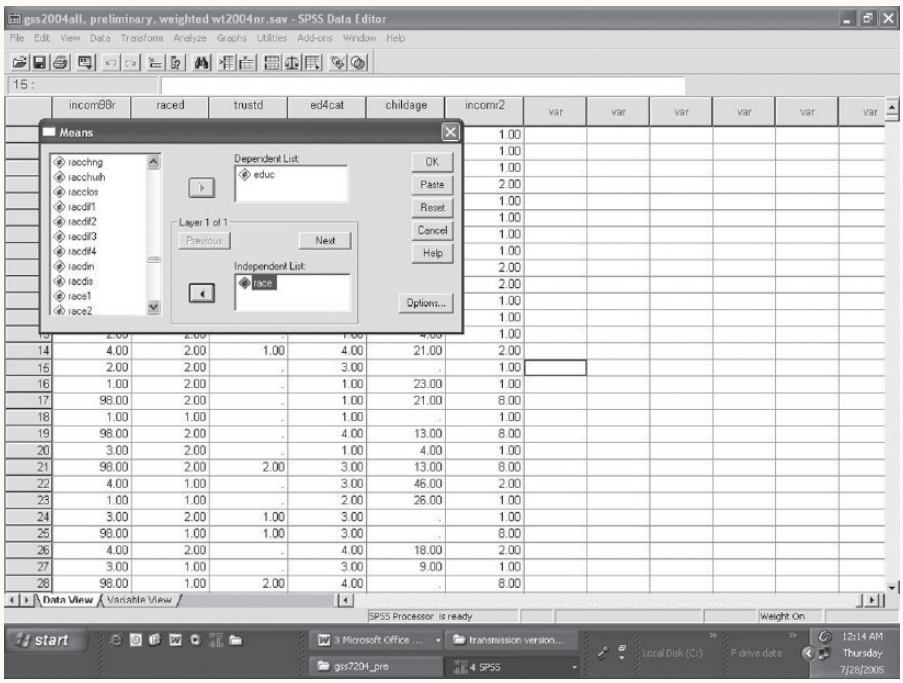

2Saving and Retrieving Files in SPSS

A number of the exercises in this book require you to recode variables. To eliminate the need to re-create variables you created via the Recode command, it is a good idea to save the dataset as a new file—with a new name—on either your hard drive or a floppy disk or a memory stick. In the Student Version of SPSS, the maximum number of variables allowed in any given dataset is 50. If you exceed the 50-variable limit, SPSS will save only the first 50 variables. You can delete variables you do not need by highlighting the corresponding column in Data View and pressing Delete on your keyboard. In any case, do *not* save the dataset with the same name. You might unintentionally make a permanent change in the values of variables in the dataset and later come to regret it.

Saving SPSS Data Files

To save the dataset as a new file, click:

File—Save As

The Save As dialog box will open. It should be similar to the File_Open dialog box shown at the beginning of this appendix.

If you are saving the file to your hard drive, you should save it in the SPSS folder or My Documents folder to find it easily later. After you type in the new name, click the Save As button.

Be forewarned: Output files can quickly take up a *lot* of space on your hard disk, or on a floppy. In most cases, you should just print the output instead of saving it in a file.

Opening Output Files

To open the output file(s) you saved, simply click:

File—Open—Output

The Open File dialog box will open. From the "Look in" drop-down text box, locate the folder where you saved the output file. Click on "ABANYFreq.spo" (or any other Viewer output file you want to open) to highlight it, and then click Open. The file will open in the Output window.

2Printing Your Output

There are two ways to print from SPSS: from the menu bar and from the tool bar. First, let's look at how to print from the menu bar. Click:

File—Print

From the print dialog box, you can choose to print all of your output or only selected items, the type of printer, and then number of copies to print. After you have made your selections, click OK. If you want to print only part of the output, you can select multiple portions by depressing the Ctrl key while clicking on each portion you would like to print (this must be done prior to selecting the print option). Then, from the print dialog box, click on the radio button for Selection in the "Print Range" section (to mark it with a small dot), and then click OK. Exhibit C.13 shows an illustration of some of these features in the print dialog box.

In addition to printing from the menu bar, you can also print by clicking the "Print" button on the tool bar (the printer icon). This brings up the same print dialog box.

Now that you have been introduced to a variety of different ways to communicate with SPSS, you may be feeling a bit overwhelmed. While you are still familiarizing yourself with SPSS, you may want to stick to the menu bar options as they guide you on how to proceed. When you become more comfortable using SPSS, you may like the convenience of the tool bar shortcuts. Keep in mind that you don't have to learn everything all at once.

Exhibit C.13 **Print Dialog Box**

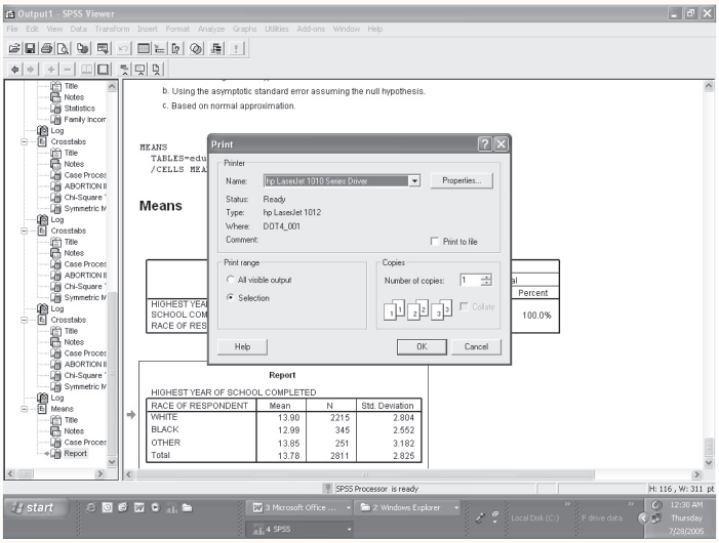

2Text Analysis

You can also use a supplementary SPSS module to code textual data systematically, if you or your university has purchased it. SPSS Text Analysis for Surveys combines a text search procedure with an automated linguistic technique that allows you to classify text according to key words and adjectives. Exhibit C.14 displays a Text Analysis screen that

Exhibit C.14 **Data Categorized with SPSS Text Analysis Software**

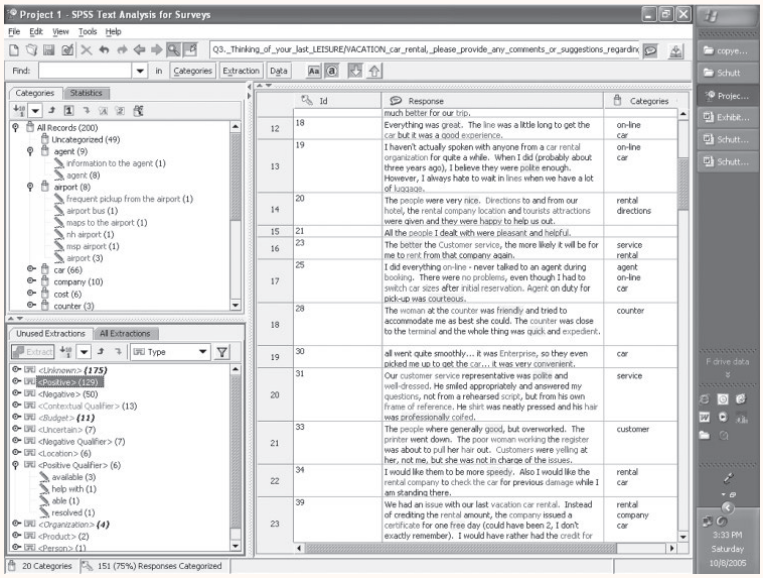

illustrates key features using an SPSS-supplied database. On the right, you can see some of the text that was entered into the case records in the original SPSS file and then imported into the Text Analysis program. In the upper left window, you can see the categories that were identified for the analysis, as well as the phrases that are grouped within the categories of "agent" and "airport." In the lower left window, you can see the text extracts that are used by the linguistic processing feature to identify statements about the categories as positive or negative. The words used in the categories and extracts are highlighted in the text, on the right. The categories also appear in a separate column on the right, for each case.

Development of the SPSS Text Analysis system reflects the way in which quantitative and qualitative data analysis procedures are converging in some areas, such as in the analysis of survey data that contains some textual data along with the quantitative responses. If you have access to this program, you will find many opportunities for using its capabilities.

2 Conclusion

At this point, you should be familiar with SPSS and able to complete all the exercises included in this book. Remember, the more you practice, the more comfortable you will be with using SPSS and statistics. If you have trouble, don't be afraid to refer to the SPSS Help file or Statistical Tutor, which are available from the Help drop-down menu, or ask your instructor or the computer lab consultant. Everyone runs into problems, but you can solve problems more quickly if you do the following:

- Write down the error message you are getting.
- Try to determine whether your problem is an SPSS problem. If so, look for help from the many sources of assistance out there: the Help drop-down menu in SPSS, the SPSS manual, some campus computer lab consultants, your instructor, and quite possibly your classmates.
- Try to determine whether your problem is a Windows problem or a computer hardware problem. If so, your college probably offers computer lab technical support. Introductory books for Windows are an excellent resource because they typically describe problems common to Windows and to many software programs.

Box C.1 Summary of SPSS Drop-Down Menus

The **File** menu enables you to open and save data files, import data created by other software, and print the contents of the data editor (or print the output, when you are in the SPSS Viewer window).
Use the **Edit** menu to cut, copy, and paste data; find data within the open data set; insert a new variable or case; and change options settings. You can also enter data directly in the "Data View" spreadsheet.
The **View** menu allows you to change how the data editor looks by changing fonts, turning toolbars on and off, turning grid lines on and off, and turning value labels on and off.
Use the **Data** menu to sort, select, or weight cases.
The **Transform** menu lets you make changes to selected variables, compute new variables, recode the values of existing ones, and replace missing values.
Use the **Analyze** menu to select the statistical procedure you want to use.
The **Graphs** menu allows you to produce a variety of two- and three-dimensional graphical displays of your data for purposes of data analysis and presentation.
Use the **Utilities** menu to obtain information about the variables in the open data set and select a list of variables to appear in dialog boxes.
The **Window** menu lets you switch between SPSS windows or minimize all open SPSS windows.
Use the **Help** menu to access SPSS help topics and run the tutorial.

Appendix D

How to Use a Qualitative Analysis Package

Ann Dupuis

You may have noticed the "HyperRESEARCH"™ option on the computer CD bundled with this book, and wondered what it's about. The short answer is: This is the demo version of the HyperRESEARCH software (for both Macintosh and Windows), including an electronic manual and a full set of tutorials with all necessary support and sample files.

For the long answer, please read on.

What is HyperRESEARCH

HyperRESEARCH is a software tool for qualitative data analysis, developed by ResearchWare, Inc. (www.researchware.com). It is one of several CAQDAS packages available. (CAQDAS is an acronym for computer-assisted qualitative data analysis software.) Like many CAQDAS programs, HyperRESEARCH's essential capabilities are for qualitative analysis—code-and-retrieve data analysis features, report-generating capabilities, multimedia support (for data including graphics, video, and audio as well as text), and theory-building tools. These features are packaged within an easy-to-use package that helps you, the user, take control.

A demo version of HyperRESEARCH is available on the accompanying CD. (You also can download HyperRESEARCH from www.researchware.com.) This version of the software is fully functional; however, you are limited to using "demo mode," which places limits on the size of your study. (Your Master Code List is limited to 75 code names, you may have only seven cases, and you may apply only 50 code instances to each case.) You

NOTE: Ann Dupuis is a Technical Support and Documentation Specialist at ResearchWare, Inc. (support@research ware.com).

Exhibit D.1 **HyperRESEARCH Basic Interface: Study Window, Code List Editor, Source Window, and Annotation Window; Not Shown: Hypothesis Tester, Report Generator, Code Map, and Other Specialized Functions**

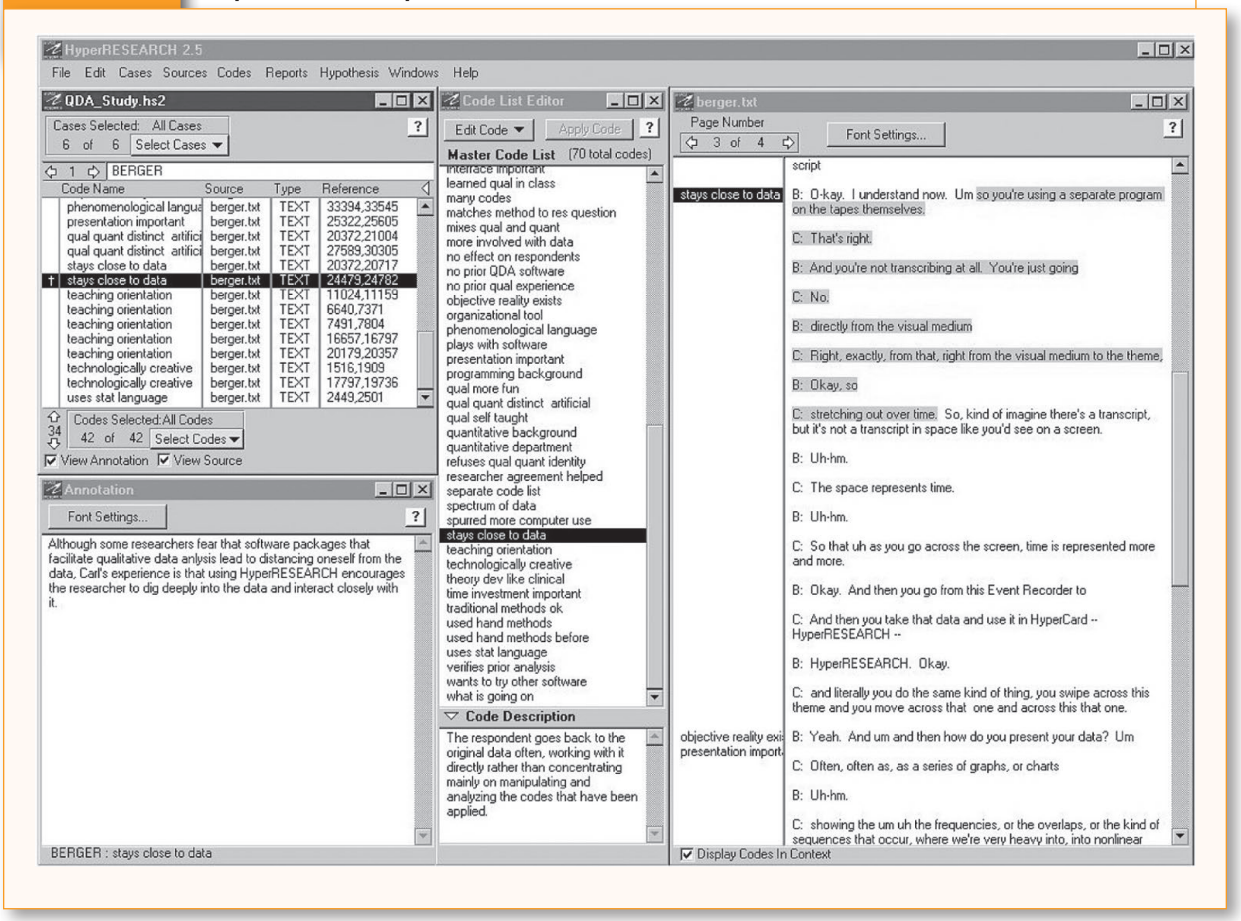

can purchase a "license key" from ResearchWare, Inc. or any of its reseller partners to unlock the software (lift the restrictions imposed in demo mode).

If you are a student, you may find that the free demo version of the software is adequate to your needs while you learn how to conduct qualitative research. If you are an instructor, you should note that the demo version of HyperRESEARCH is free for you and your students to use. ResearchWare, Inc. also has a promotional program in which qualifying educators will receive a free license key to unlock the unrestricted version of HyperRESEARCH for their own use. Visit its Web site at www.researchware.com for more information.

2A Quick Tour of HyperRESEARCH

This introduction to HyperRESEARCH briefly shows the major features of the software. For a more in-depth, step-by-step look at how HyperRESEARCH facilitates qualitative data analysis, it is recommended that you use the tutorials on the CD. You can install HyperRESEARCH on your computer and follow the step-by-step tutorials to learn first hand how to use the software. The CD includes sample research materials and HyperRESEARCH studies you can use in conjunction with the tutorials, or to explore the software's capabilities on your own.

Exhibit D.2 Graphic Representation of the Structure of a HyperRESEARCH Study

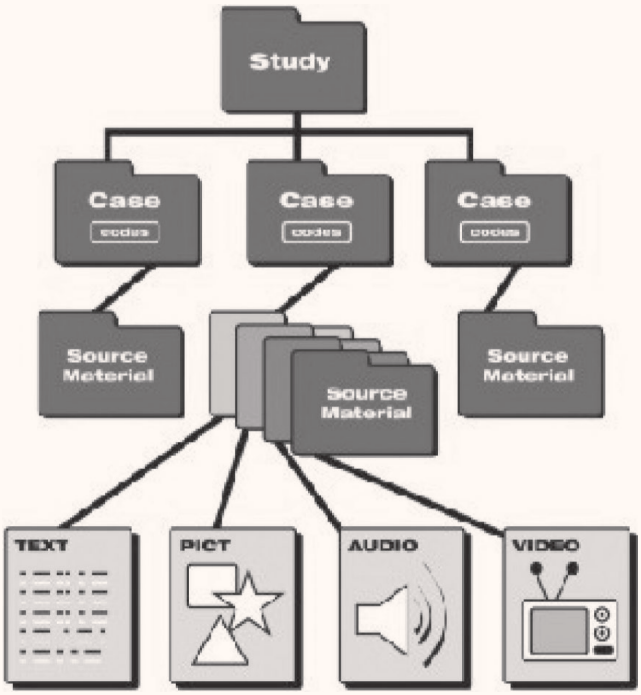

HyperRESEARCH's Flexible Structure and Point-and-Click Interface

HyperRESEARCH allows you to organize your data in many ways. A study consists of one or more cases (a case is the unit of analysis in a HyperRESEARCH study) (Exhibit D.2). You decide what a case will represent, such as an individual, a time period, or a focus group. HyperRESEARCH allows you to choose your codes and code relationships, the depth of your analysis, and the source of your data (i.e., text, graphic, audio, and video sources). HyperRESEARCH allows you to apply codes for multiple sourcefiles to a single case (or a single source file to multiple cases). You also can assign multiple codes to any chunk of source material.

HyperRESEARCH's point-and-click interface features pull-down menus, click-and-drag selection, and keyboard shortcuts for those who prefer keyboard commands to mouse clicks.

The Study Window

The Study window is the main HyperRESEARCH window, where your cases and code references are displayed. The Study window also shows you how many cases are currently in your study, and how many code instances have been applied to the current case (Exhibit D.3).

You can view the code of references of one case at a time. Each code reference consists of a code name (from your Master Code List, entered into the Code List Editor), the source file name, the source type (text, audio, video, graphic), and a code reference (HyperRESEARCH's reference points for recalling the source material).

With the View Source option selected, clicking on any code reference will recall the underlying source material. HyperRESEARCH will open the file in a Source window, with the underlying source material highlighted.

Exhibit E.3 Study and Source Windows Showing Highlighted Code and Source Material

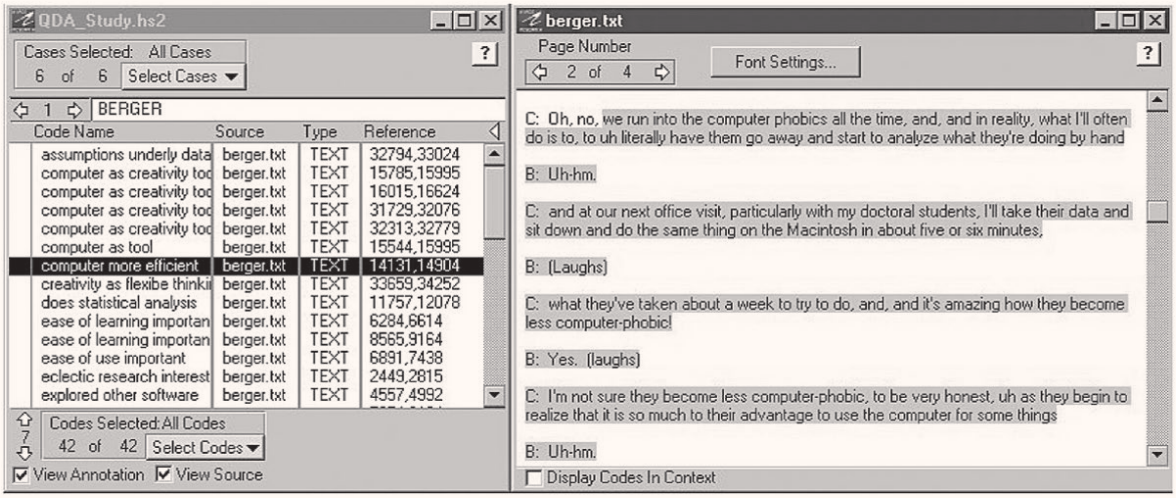

While working in your Study window, you also can manipulate your code instances. The Codes menu offers a variety of commands you can use with your codes, including duplicating code instances (applying additional codes to the same source material referenced by the original code), recoding (changing the code name applied to the referenced source material), and deleting (removing one or more specific code instances from the case) (Exhibit D.4).

You can also sort the code instances in your study by name, by reference, by type, by source file, or by any combination of these criteria.

The Select Cases and Select Codes commands allow you to concentrate on subsets of your cases and codes. These powerful commands facilitate quick review of themes and patterns in your data and coding. Used in

Exhibit D.4 Codes Menu Showing Options Available While Working With Code Instances in the Study Window

conjunction with the report generator and the hypothesis tester, code and case selections allow you to temporarily ignore extraneous data when generating reports or testing hypotheses.

You may select codes by name, by type (text, audio, etc.), by criteria (including Code Proximity functions), or via the Code Map. The codes you select will appear in the Study window. Codes not included in the selection will be hidden from view. You can recall them at any time by altering your code selection parameters (Exhibit D.5).

The Code List Editor

The Code List Editor allows you to create, view, and manipulate your Master Code List. You can add codes, edit codes (with changes being reflected in the individual code references throughout your study), and enter detailed definitions or descriptions for your master codes (Exhibit D.6).

You can use the Code List Editor to enter a Code Description for any of your master codes. It's great for quick reference when deciding exactly which code to apply to a source chunk (a block of text), especially when several researchers are working on the same study.

Any changes you make to the Master Code List will be reflected in the individual code instances applied to the cases in your study. Thus, if you wish to rename all instances of a master code throughout your study, you would use the global Rename command available from the Edit Code menu in the Code List Editor.

To affect specific code instances, rather than all codes throughout your study, use the main Codes menu commands in conjunction with individual code instances in the Study window.

The Source Windows

HyperRESEARCH has four Source window types, one for each of the four types of source material (text, graphics, audio, and video). Displayed reports include hyperlinks to underlying source material. HyperRESEARCH's ability to

Exhibit D.5 **Study Window Showing Codes Selected By Name (All Other Codes Are Hidden)**

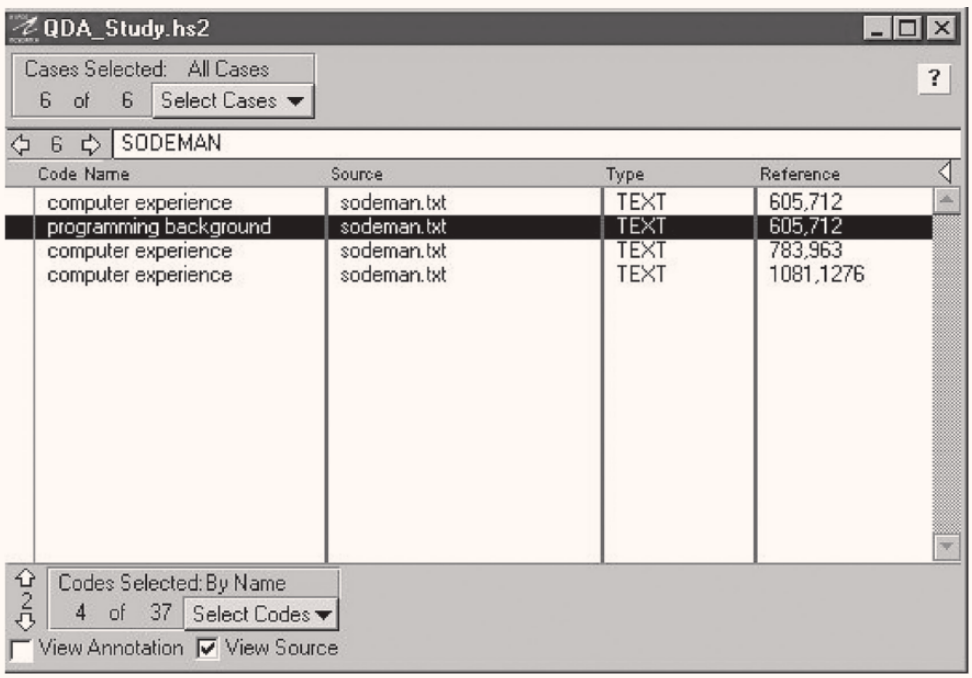

Exhibit D.6 The Code List Editor, Showing a Code Description

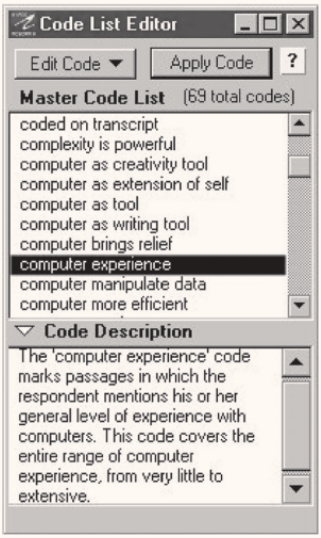

work with multiple data types, such as text, graphic, audio, and video sources, provides the flexibility to integrate all of the data necessary to conduct your research.

The Text Source Window

The text Source window displays text files. You can customize the Font Settings (typeface and size) and also choose whether to Display Codes in Context (code names appear in the left margin). This window is fully resizable and movable.

HyperRESEARCH 2.5 allows you to select any chunk of text (from one character to an entire file) and apply any number of codes to it. Text source files also can be split into multiple pages, if you wish (Exhibit D.7).

View your codes adjacent to the source material by using Display Codes In Context. The material can be sent to a printer with the codes appearing in the left margin. The Codes In Context feature also can be turned off, maximizing the space available for text (with no margin for viewing code names).

To select text for coding, simply click and drag over the desired chunk. Then use the Code List Editor to apply one or more codes to the source selection.

The Graphic Source Window

The graphic Source window allows you to display still images (.gif,. jpg, or similar graphic files) and assign codes to selected portions of the image (Exhibit D.8).

To select a portion of a graphic for coding, click and drag from one corner of a rectangular chunk to the opposite corner. Then use the Code List Editor to apply one or more codes to the graphic selection.

The Movie Source Window

The movie Source window displays movie files (with their audio tracks, if any) using Apple's QuickTime software. You can select and code any number of frames, which can be replayed when recalling the source material from the Study window or in a hyperlinked report (Exhibit D.9).

Exhibit D.7 Source Window (Text)

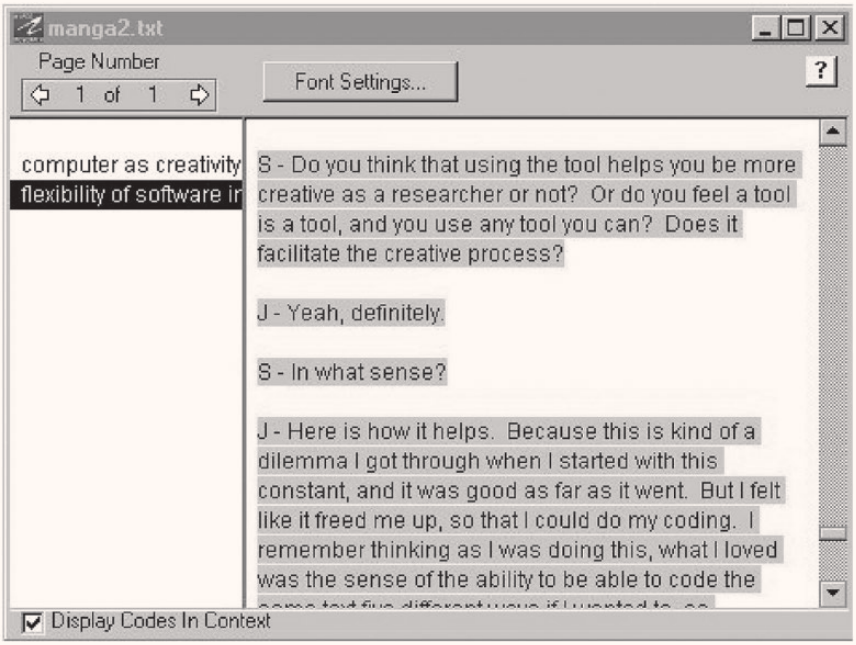

Exhibit D.8 Source Window (Graphic)

Exhibit D.9 Source Window (Video)

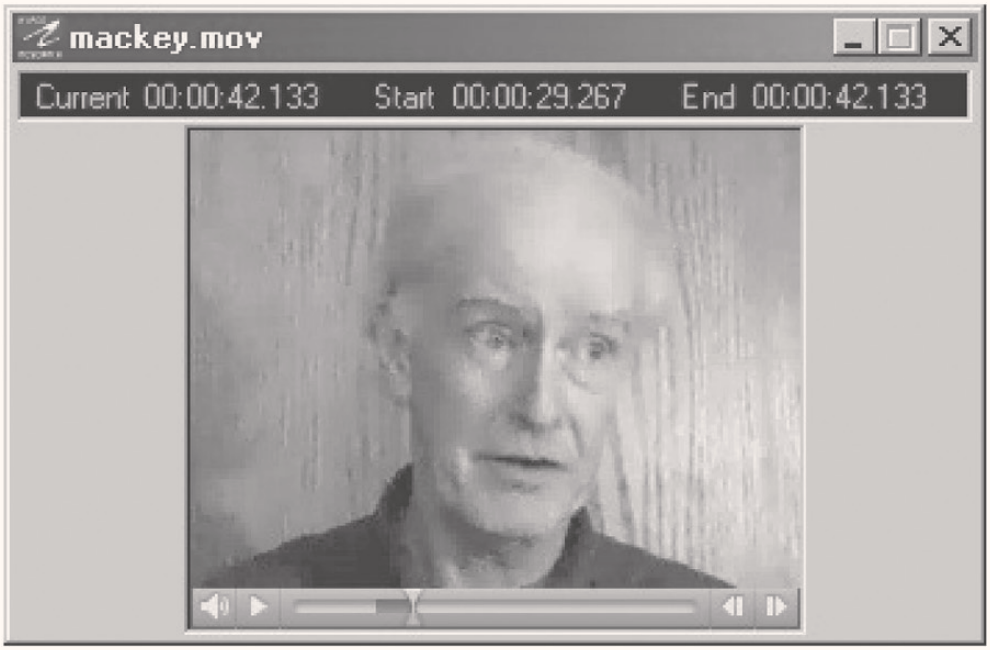

To select a video segment for coding, simply click and drag along the slide track while holding down the Shift key. Alternatively, position the slide bar at the beginning of the desired segment, then hold the Shift key down and click on the Play button. Release the Shift key when you reach the end of the desired segment.

After you have selected a video segment, you can fine-tune your selection. The video window controls offer frame-by-frame precision.

As with text selections, you can apply more than one code to a given chunk.

The Audio Source Window

The audio Source window uses Apple's QuickTime software to allow you to play back an audio file and select portions for coding. Viewing the coded source material of an audio file (either by selecting the code reference on the Study window with the View Source feature active, or clicking on a hyperlinked code reference in a report) recalls and replays the selected portion of the audio track (Exhibit D.10).

The audio Source window controls are identical to the video Source window controls.

Exhibit D.10 Source Window (Audio)

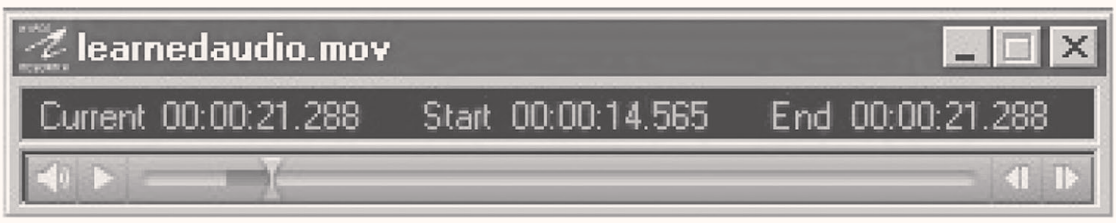

The Annotation Window

The Annotation window allows you to add a memo or annotation to any code reference in your study. Unlike a Code Definition, which applies to a master code, an Annotation is specific to an individual code reference and its underlying source material (Exhibit D.11).

To annotate a code, select a code reference on a case card, choose the Annotate command, and add up to 32,000 characters of information per reference. Annotations can be used as a built-in memo system. All annotations are fully editable and can be included in reports.

The Report Window

The Report window presents the report generation options. Use this window to customize the data you wish retrieved for a given report.

Generate custom reports and display them on your screen or save them as text to output to a word processor, spreadsheet, or statistical package. Hyperlinked reports allow you to click on any code reference to view the source material (Exhibit D.12).

In addition to specifying which elements you wish included in a report, you can use the Select Cases and Select Codes commands to specify precisely which cases and codes you wish included in the report. Generate a report that includes all codes across your entire study, or report on any subset of cases and codes.

If you choose any report elements from the right column (Master Case List through Master Source File List), the report generator will also display a header page with the specified information (Exhibit D.13).

The body of the report will present results based on the current selection of cases and code instances in your Study window. You may choose to include the source material and any annotations for reported codes as well as the code names themselves.

If the Hyperlinks elements are all included, the code reference lines will actually be hyperlinks. Clicking on a code reference (e.g., "MACKEY used hand methods before 1 TEXT 13109,13334 mackey.txt") will open the underlying source material in a Source window. This allows you to view the source material in the context of the rest of the file if you wish (Exhibit D.14).

Exhibit D.11 Use of Boolean Connectors in a Literature Search

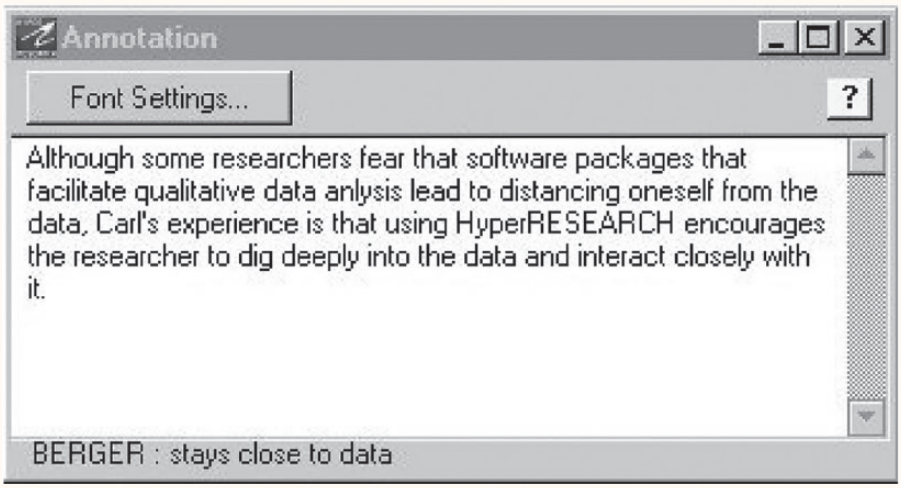

Exhibit D.12 Report Window Showing Report Element Options

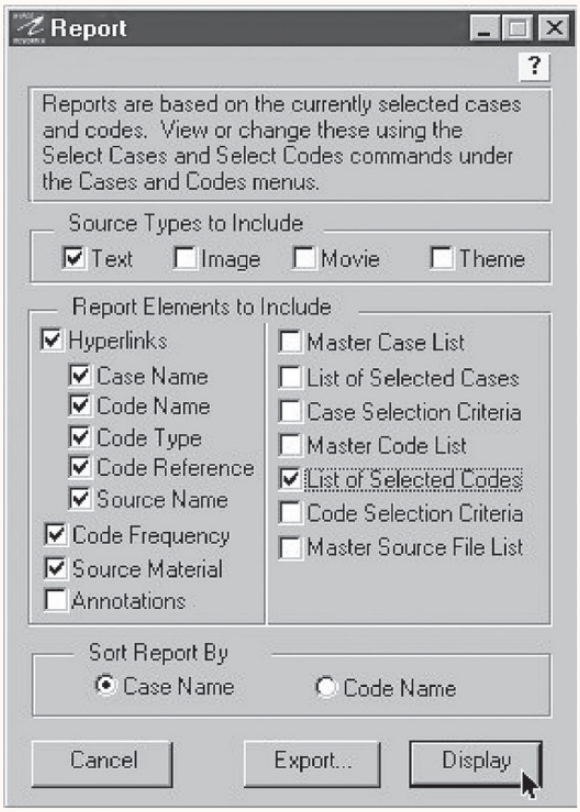

Exhibit D.13 Report Display Window with Header Information

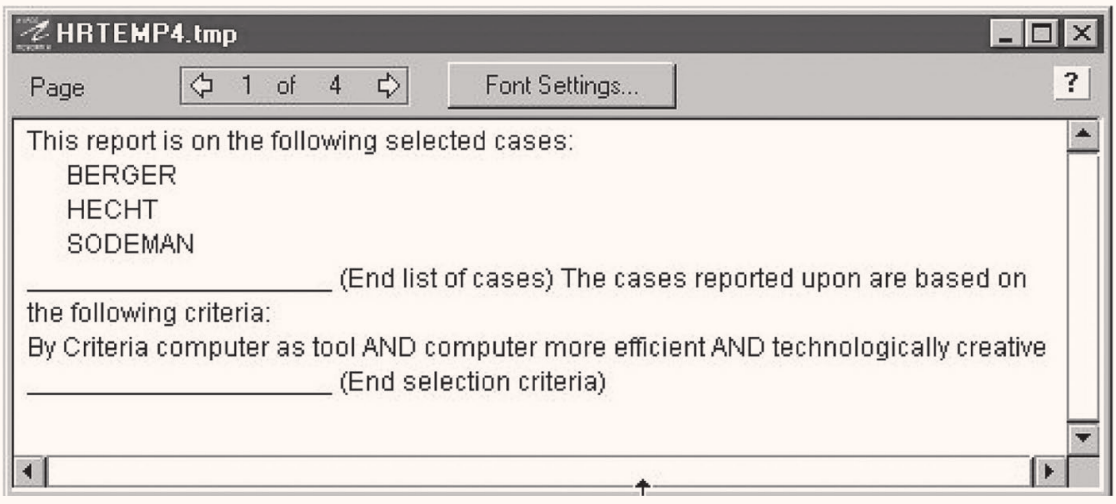

Exhibit D.14 **Report Display Window Showing Hyperlinked Code References and Source Material**

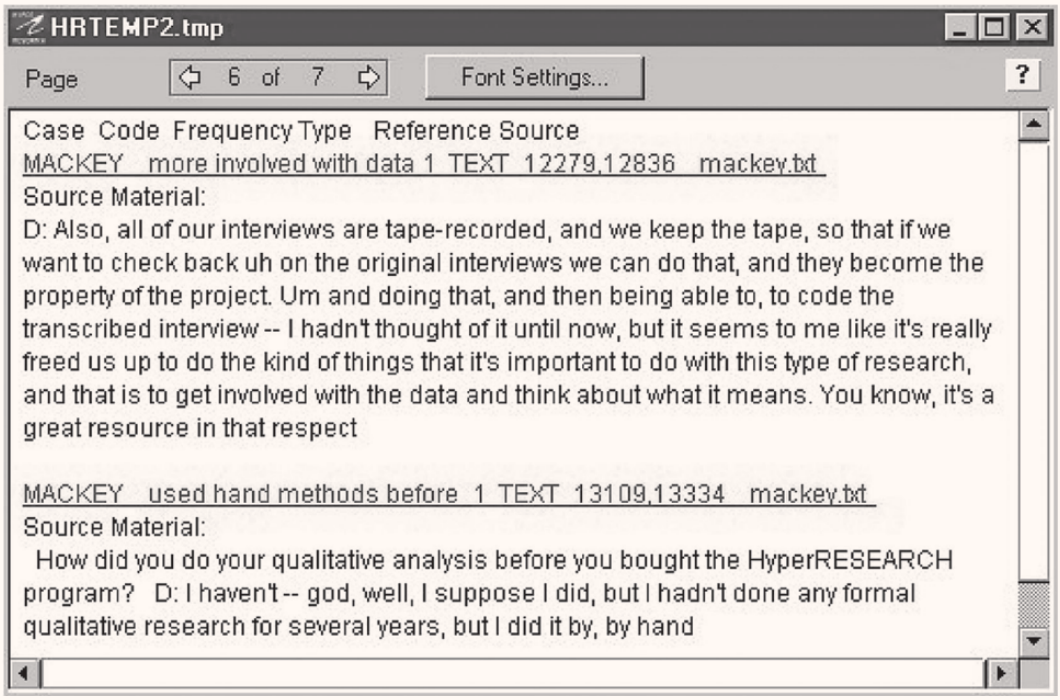

You can save report settings (including the current selection of cases and codes in your study) to run the same report again. You also can export the generated report to a text file, which you can work further with in a word processor or other program.

2Learning More About the Basics

We've covered the basics of HyperRESEARCH: coding and retrieving, code manipulation, and generating reports. You may want to stop reading now, and start experimenting with the software itself. Tutorials One through Four on the CD offer step-by-step "walkthroughs" for the procedures for starting a study (or opening an existing one), coding source material, manipulating code instances, and generating reports. Tutorial Five covers coding and retrieving graphic, video, and audio source material. After you've mastered these procedures, you'll be ready for the more advanced and specialized capabilities HyperRESEARCH offers.

For an introduction to HyperRESEARCH's advanced features (covered in depth in Tutorials Six and Seven), read on.

2Advanced Features

In addition to the basic code-and-retrieve features HyperRESEARCH offers, there are several more advanced features available.

Autocoding

With Autocode, you can automatically assign a code to multiple sources and multiple cases, looking for several phrases or words in a single pass. Specify a number of characters, words, or lines before and/or after the found phrases to be included in the chunk selected for autocoding.

First you assign source files to the proper cases. Then you specify the phrase (or phrases) to search for and how much surrounding source material to include (Exhibit D.15).

Finally, you select the codes to apply to the matching source materials (Exhibit D.16).

The Autocode button becomes active when all necessary parameters are supplied.

HyperRESEARCH will apply the selected code or codes to all instances of the selected phrase it finds, and place in the Study window for the specified cases.

Autocoding is best used as a first-pass tool, to earmark certain sections of your textual data for more precise coding. You may wish to use separate code names for autocoded passages (such as the "autocode phobic" code name in the example pictured in Exhibit D.15). These code names would then be replaced with final code names (such as "evidence of computer phobia") when you examine the actual source material and select a more precise segment of text to code.

Code Proximity Searches

HyperRESEARCH lets you conveniently reference overlapping code instances with Code Proximity searches. One of several code- and case-selection tools available, the Code Proximity functions will seek out specific relationships between two code names (Exhibit D.17).

Exhibit D.15 **Autocode (Phrases) Window Set to Select Entire Paragraphs Around Each Occurrence of the Specified Phrases, in Selected Sources**

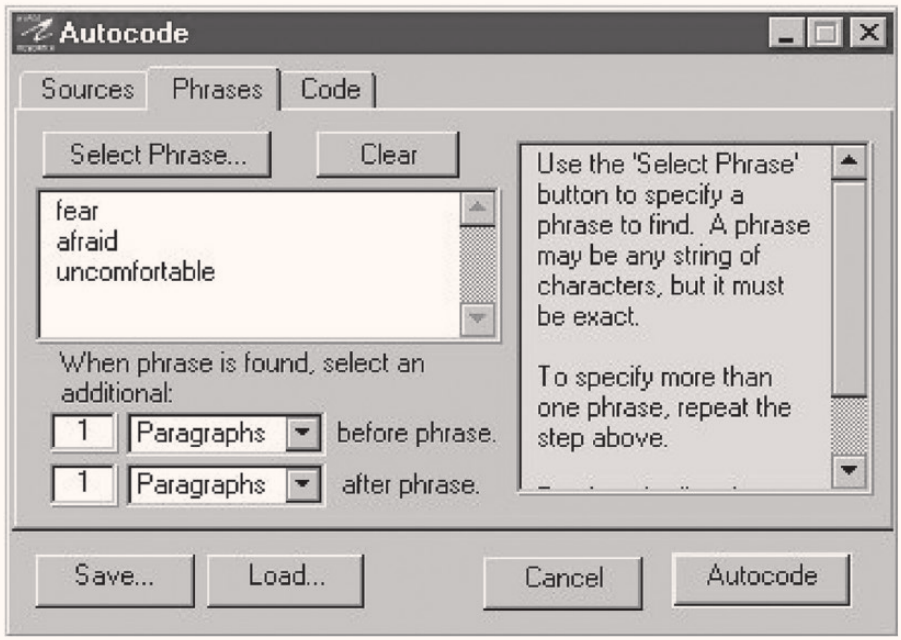

Exhibit D.16 Autocode (Code) Window With Code Names to Assign

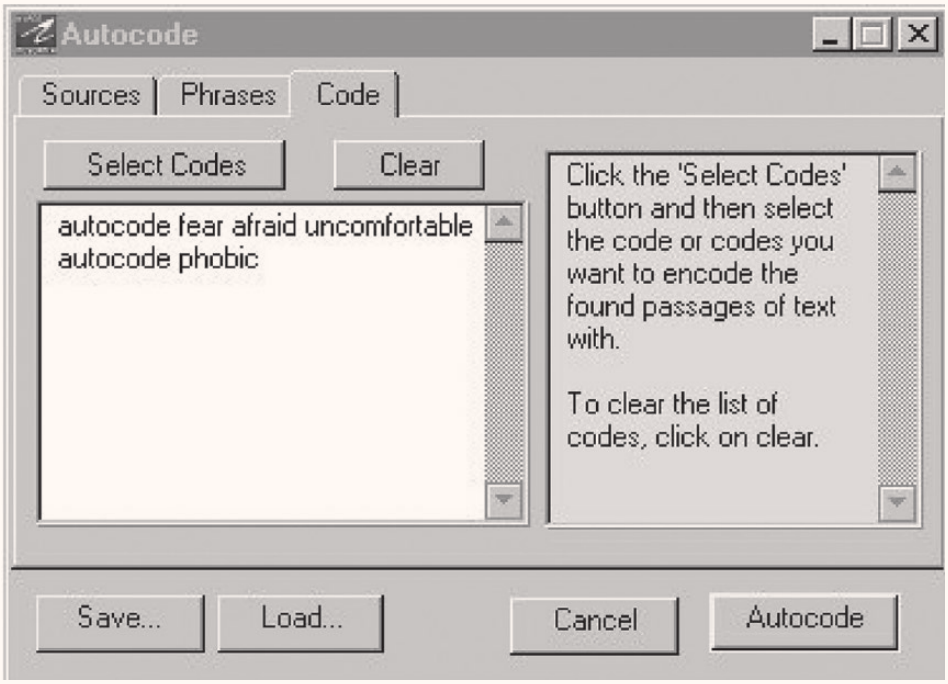

Exhibit D.17 Selection of Codes Based on Code Proximity Functions

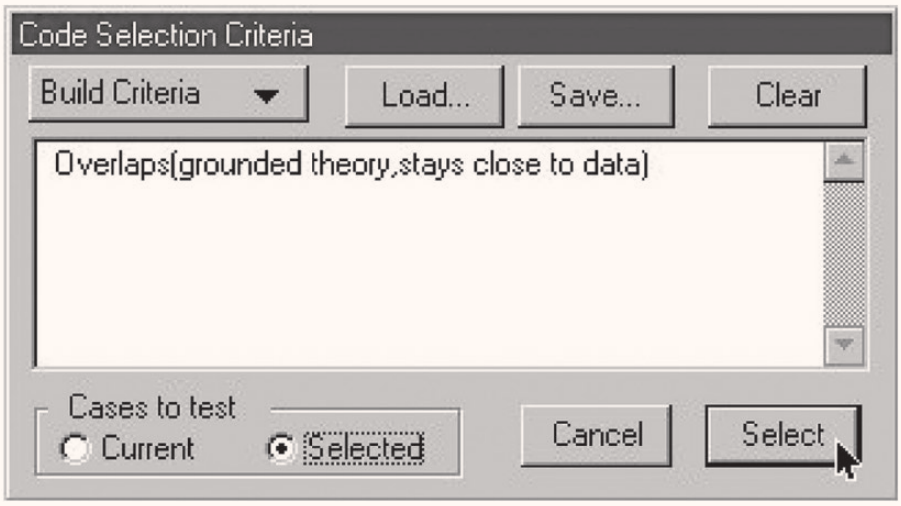

Available functions are:

- *Equals*— The source material for "code 1" and "code 2" match exactly. With text files, the starting and ending character placement for the source material selection will be exactly the same for being compared.
- *Excludes*—The source material for "code 1" completely excludes any source material coded with "code 2." There are no overlapping characters, pixels, or video or audio segments.
- *Includes*—The source material for "code 1" completely includes the source material for "code 2." For example, a textual selection of two paragraphs coded with "code 1" will include the source material for "code 2" if "code 2" has been applied to one of those paragraphs (and no other source material outside the selection for "code 1").
- *Overlaps*—The source material for "code 1" overlaps one or more characters coded with "code 2." Code references that qualify for the equals or includes functions will also qualify for the overlaps function. However, the matches for overlaps don't need to be as specific as those for includes or equals. Coded text segments that share even one character can qualify as overlapping. With graphics, video, and audio files, the overlapping of even one pixel or time segment is enough to qualify for the overlaps function.

You can use the Code Proximity functions to select subsets of your codes or subsets of your cases, based on the relative placement of the codes segments within your source files.

The Code Map Window

The Code Map window allows you to explore graphic representations of the relationships between your master codes. You can group codes in any way you wish, and visually link master code names to one another. Arrange your codes visually to indicate code families, trees, or networks (Exhibit D.18).

Exhibit D.18 The Code Map Window

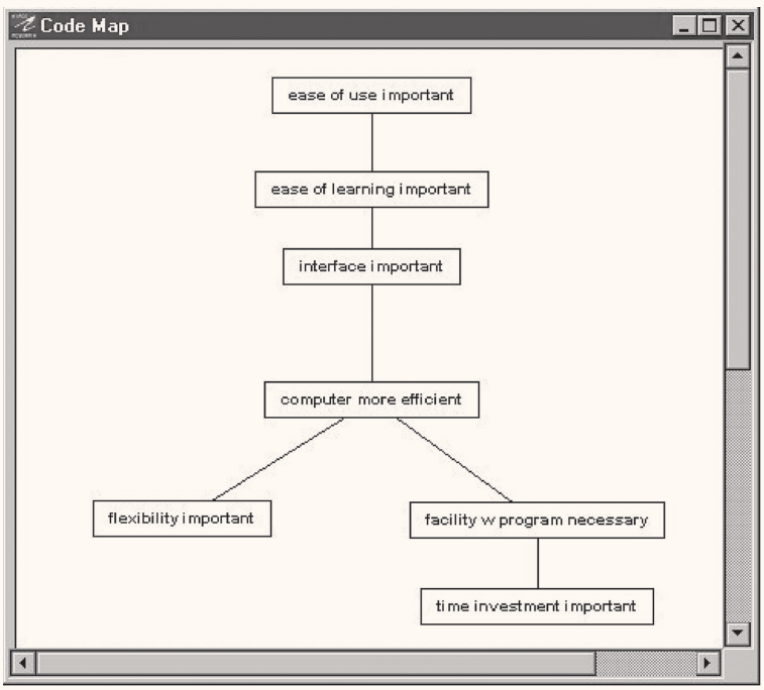

You can select codes using your Code Map and apply that selection to the Study window, which will then display only those code references corresponding to the selected master codes. Selections of mapped codes can be made by selecting each code individually, or by selecting one code and then expanding the selection based on code links. In Exhibit D.18, commanding HyperRESEARCH to select all codes within two links of the "computer more efficient" code would select everything but "ease of use important."

Thus, the Code Map can be used as a visually oriented code selection tool as well as a visualization tool.

The Hypothesis Window

The Hypothesis Tester is an "expert system" that helps you build theories and test them against the codes you've applied to your data. Like the Report Generator, the Hypothesis Tester consists of two windows: the Hypothesis Test window and a Report Display window.

The Hypothesis Test window includes a section that displays the current Hypothesis Test Rules, and a section that allows you to edit those rules (Exhibit D.19).

When the Hypothesis Rule List is complete, you can Export the hypothesis test report to a text file, or Display it to the Report Display window. You can also save the hypothesis for later use, or open an already-constructed hypothesis (Exhibit D.20).

Although the Hypothesis Tester may seem complicated at first, at its heart it's another way to examine your cases, looking for combinations of the presence and/or absence of code names. It utilizes Boolean expressions (delineating the code combinations to search for with the use of logical AND, OR, and NOT) and can also utilize the Code Proximity functions to look specifically for coded segments that overlap one another.

You can also use the Hypothesis Tester to add Theme codes to your case. Say you'd like to find every case that has been coded with both "gets married and stays married" and "wants kids." You'd like to apply the code "high family commitment" to each relevant case.

Exhibit D.19 **The Hypothesis Window Showing Part of a Hypothesis Rule List**

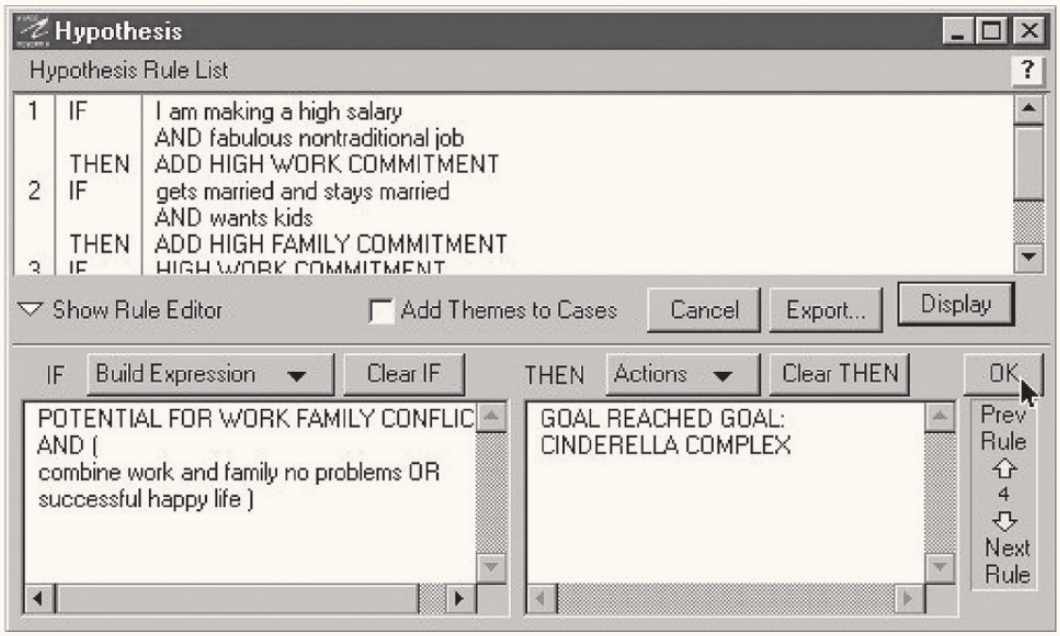

Exhibit D.20 The Hypothesis Menu

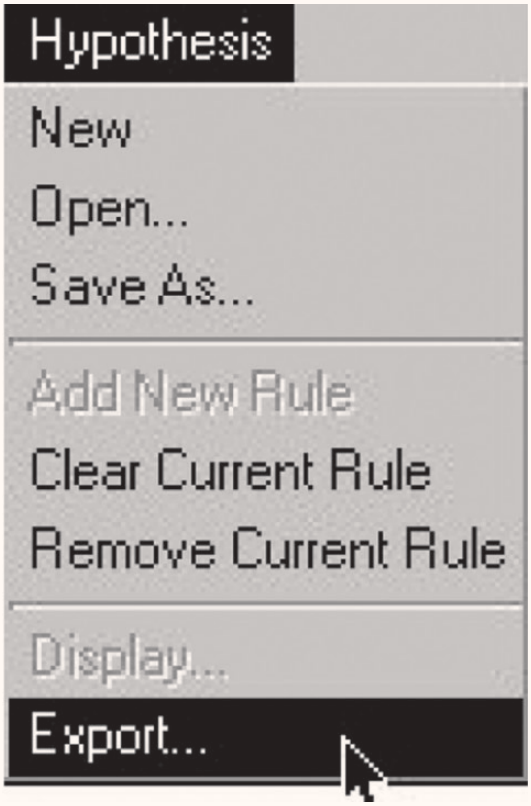

You could do this by selecting cases by criteria ("gets married and stays married" AND "wants kids"). You could then go to each selected case and either duplicate one or more "gets married and stays married" and "wants kids" with the "high family commitment" code, or you could open the relevant source file, select a passage (possibly one related to those already coded), and apply the "high family commitment" code directly. Alternatively, you can do this by creating a hypothesis test that does essentially the same thing. Such a test would have one rule: IF ("gets married and stays married" AND "wants kids") THEN ADD CODE "HIGH FAMILY COMMITMENT". (Using all capital letters for the code name helps distinguish it as a Theme code.) With the Add Themes to Cases option checked, running this hypothesis test would tell HyperRESEARCH to find any case that had one or more instances of "gets married and stays married" and also one or more instances of "wants kids" already coded to it. HyperRESEARCH would then add the "HIGH FAMILY COMMITMENT" code name as a Theme code. This code would not point to any specific source file or source material. It would, however, be considered in any Select Cases or Select Codes command, and would show up on your Reports (provided you chose to include Theme types as a Report Element) (Exhibit D.21).

Exhibit D.21 Study Window With Theme Code "HIGH FAMILY COMMITMENT" Applied

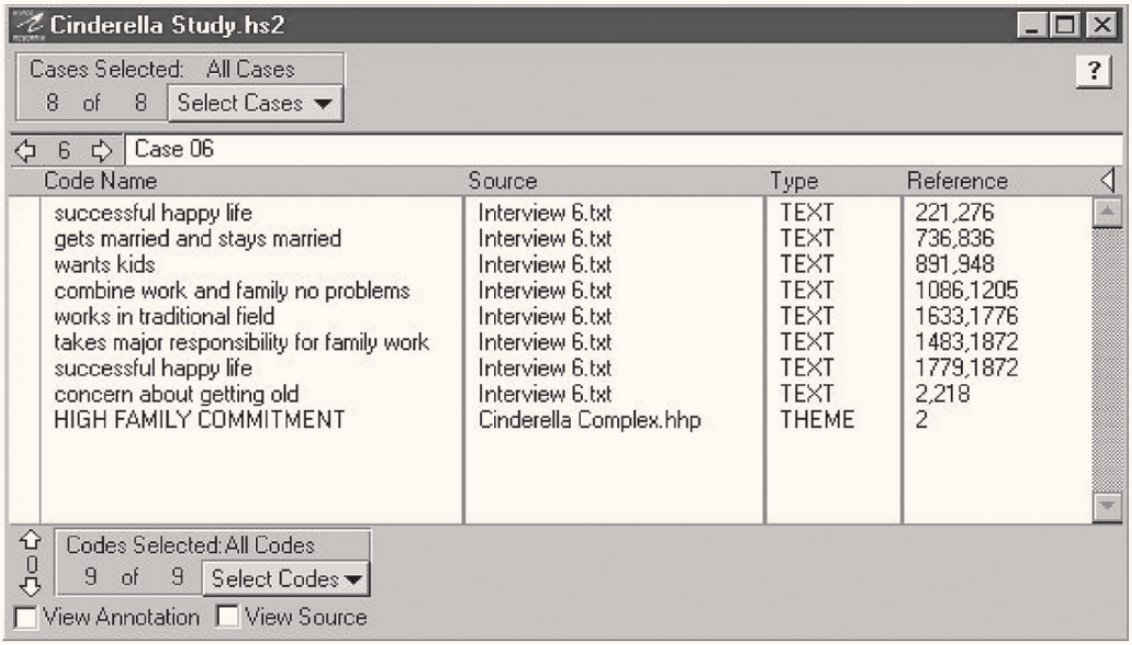

2 Learning more about HyperRESEARCH

The best way to learn more about HyperRESEARCH is to dive right into the materials on the CD. Install the software (both Macintosh and Windows versions are included), print out the tutorials (at least Tutorials One through Four, which cover the basics), and play around with the sample studies (the Cinderella Study and the Qualitative Data Analysis Study).

You also can visit ResearchWare's Web site at www.researchware.com. You'll find instructions there on how to join the HyperRESEARCH email discussion list. The Web site also lists events (workshops, conferences, and trade shows), links for online resources for qualitative data analysis, and more.

Welcome to the world of CAQDAS!

Appendix E

How to Use a Data Spreadsheet

Excel

One does not necessarily have special statistical software to perform statistical analyses. Microsoft Office Excel can be used to run statistical procedures. Although in some respects Excel is not as preferable for data analyses as IBM SPSS, it is very user-friendly with simpler statistical procedures. This appendix describes how to use Excel to execute basic statistical calculations. Data from the 2004 version of the General Social Survey (GSS) is used for examples. This Appendix is based on Excel 2007 version, which differs in certain aspects from Excel's previous versions. The most notable change that affects the exercises presented in this appendix concerns the pivot table feature.

2 Basic Procedures

Starting Excel

To start Excel using Windows, click on the **Start** button at the bottom left corner of the screen. Under **Programs** locate and click the **Microsoft Excel** icon.

The layout of the Excel program has changed substantially for the Microsoft Office 2007 edition compared to its predecessors. Commands are now grouped in ribbons that are accessed by clicking on a specific tab. Thus, the **Home** tab grants access to a ribbon of several command groups: **Clipboard**, **Font**, **Alignment**, **Number**, **Styles**, **Cells**, and **Editing**.

Once the program is started you will see a Worksheet Area that consists of cells forming columns and rows. Rows are identified by numbers, and columns are identified by letters. Consequently, each cell has its own unique address—a

combination of letters and numbers. For example, cell C6 is in column C, row 6. The dark rim around a cell means that the cell is highlighted or active. You can highlight a range of cells by clicking and dragging the cursor across several adjacent cells.

At the bottom left of the screen you will find worksheet tabs labeled Sheet 1, Sheet 2, etc. You can rename these worksheets, add additional ones, or delete ones, you do not need.

Opening data

The data you need to use might be saved in a format other than Excel (file.xls). The data file we will be using in this tutorial is saved in SPSS format (file.sav). This is not a problem. Open the data file in SPSS. Click **File—Save As**—choose *Excel 97 and later (.xls)* in the **Save as Type** drop down menu. The file is ready to be open in Excel. To do so—either double-click on the icon of the saved Excel file or, if Excel program is already open, click **Office** button at the top left of the screen—**Open**—locate the data file you need to work with.

Looking at the Data

Once the data is open you will discover that the view is somewhat different than what you saw in SPSS program. Variable names are displayed in the first row (row A). Unfortunately, there is no variable view worksheet in Excel, and transferring the file from SPSS to Excel results in a loss of variable labels and value labels. Keep the codebook for the data on hand!

Sorting Data

As you are exploring the data you might want to take advantage of the Sorting tool, which allows you to sort data by two or more variables in ascending or descending order. To sort your data, click the **Data** tab, find the **Sort and Filter** command group, and click on the **Sort** icon. A dialog window will open where you will specify the variables and order you want your data to be sorted by. If the names of variables (variable labels) are entered in the first row of your data make sure you check the "*My data has headers*" box at the top right of the dialog window. You can sort numeric data as well as text (in A to Z order). If you want to sort by more than one variable click **Add Level** button at the top of the dialog window. Click **OK**.

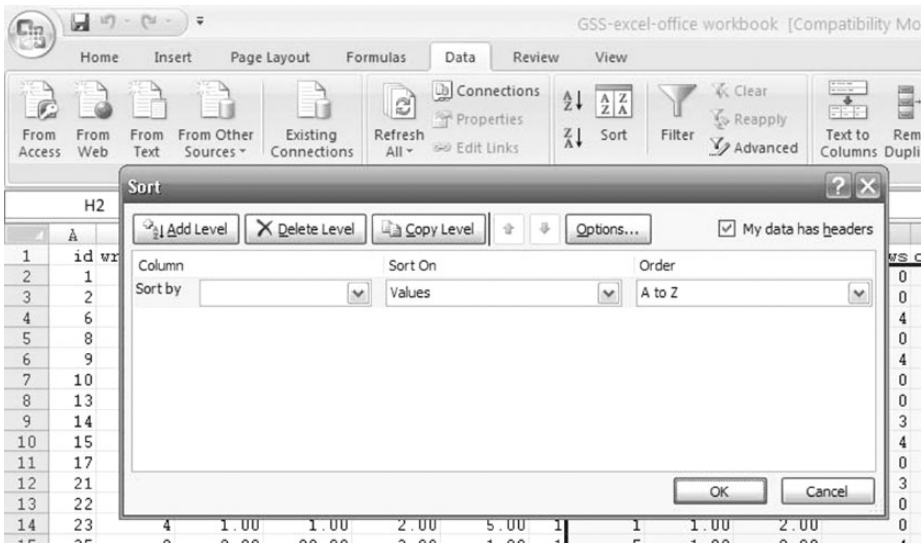

Missing Data

If you have converted an SPSS file to be used in Excel, be mindful of how missing data was coded in the original file. If missing data was entered as blanks, you have nothing to worry about. However, if missing data in the original data file was entered as a certain numerical expression (e.g. zeroes or 99) this might create problems for running certain statistical calculations in Excel (e.g. computing mean).

Loading Data Analysis Toolpak

Before we start with running some basic analysis make sure that Excel's Data Analysis Toolpak is loaded on your computer when you start the program. The Data Analysis Toolpak comes as an Add-In to your Excel program and is found on the installation disk. To see if the Toolpak is loaded in Excel 2007, check if there is an Analysis icon under the Data tab.

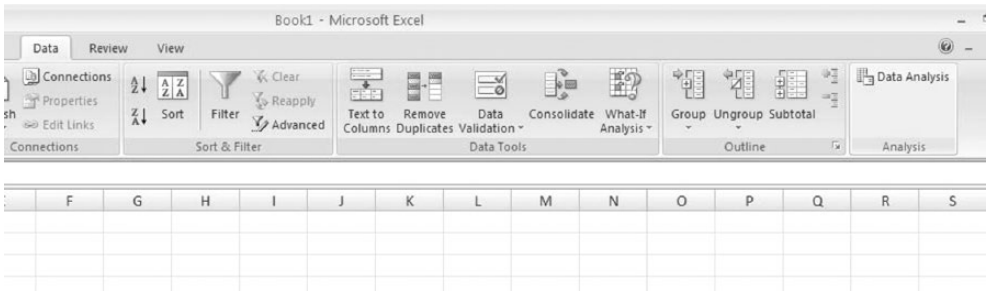

If you do not see one, click **Office** button—**Excel Options**—**Add-Ins**. In the dialog window that opens check if Analysis Toolpak and Analysis Toolpak-VBA are listed under *Active Application Add-ins*.

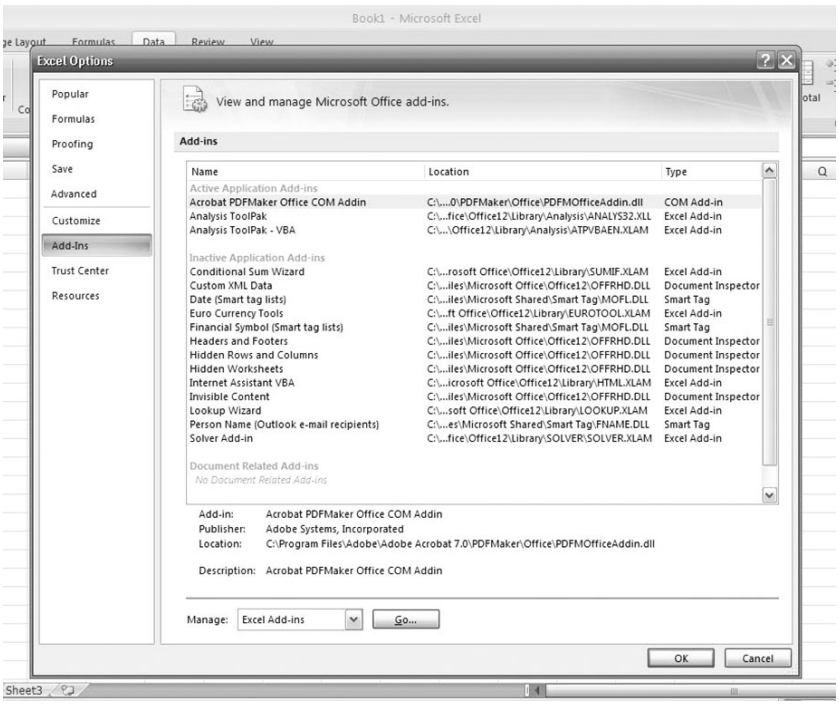

If they are listed as inactive select both of them and click **GO** at the bottom of the window. Next, check the boxes for both Add-Ins in the dialog box that opens. Click **OK**.

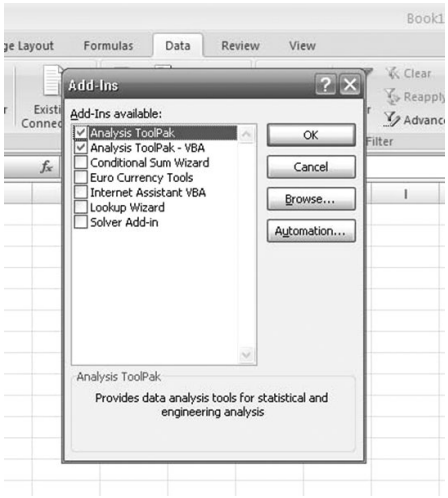

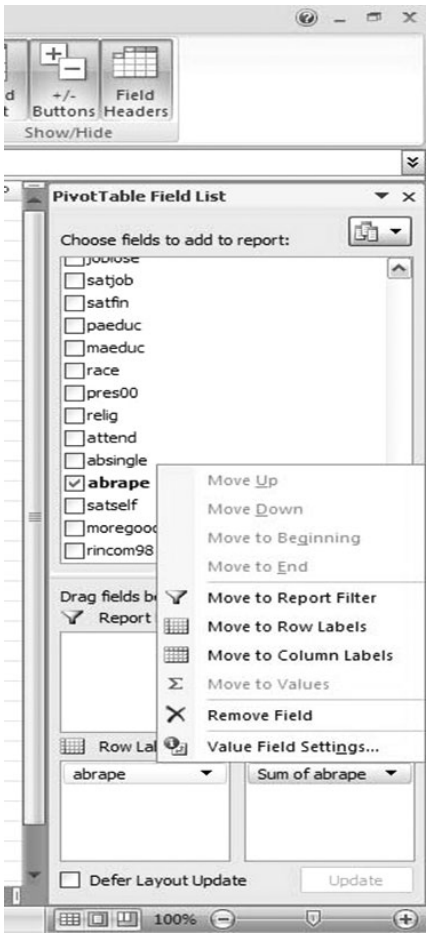

2 Univariate Statistics

Frequencies

The easiest way to create frequency distributions in Excel is by using the Pivot Table. To create a frequency distribution of the variable ABRAPE (pregnant as a result of rape) click on any cell in your data. Next click **Insert Tab** - **Pivot Table**. In the opened dialog window you will see a *Table/Range specified* that includes the call range for all of your data. It is also possible to specify the range manually by typing A1:AW1500. Next, lace your Pivot Table: Click **New Worksheet—OK**.

A new worksheet will open with newly created Pivot Table. In a **Pivot Table Field List** on the right side choose and click the variable we are interested in – ABRAPE. Do not be alarmed that your newly created Pivot table displays the sum of all the ABRAPE values in the data. Since you are interested in the count or each of the values of the variable and not the sum of the values, click a dropdown menu located under the **Pivot Table Field List** and make sure that the **Fields Section and Areas Section Stacked** is selected. Next click on the ABRAPvariable and drag it down tothe **Row Labels** box. Next click the dropdown menu to the right of Sum of Abrape—**Value Field Settings**—change **Sum** to **Count—OK**. Now you have a table of frequency distributions for the variable ABRAPE. Unfortunately, Excel does not dispvalue labels like SPSS does, so your table can look a little bit confusing with numbers in place of the actual values of the variable (e.g. "1" instead of "Yes"). To rectify this problem, you can either correct the codebook for the data or recode the variable prior to creatinga Pivot table (substituting numerical values with text).

Excel does not offer an option to simultaneously display both the count and percentage within the same Pivot table. However, you can view the percentage distriby clicking **Value Field Settings—Show values as**—choose **% of total**.

Descriptive Statistics

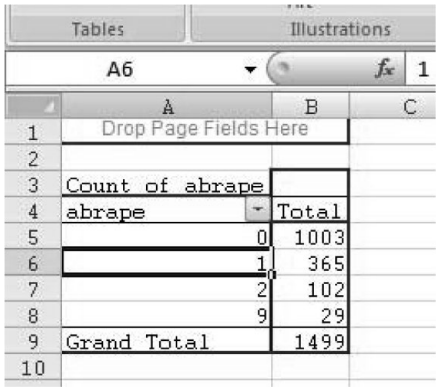

The easiest way to obtain descriptive measures on a variable in Excel is by using the Data Analysis Tool. Under the **Data** tab click **Data Analysis** - **Descriptive Statistics**—**OK**. In the descriptive statistics dialog window specify the range of the variable you are interested in. Variable EDUC occupies the range O1: O1500. You can either specify the range manually by typing it into the *Input Range* Dialog box or highlight the column that the variable occupies in the dataset.

Excel offers the option of grouping each variable in its own column or its own row, with former being the default option. Click *Labels in First Row* option, or you are risking receiving an error message "Input range contains non-numeric data" since the program will get confused with the name of the variable being the fist cell in the column.

Next you can choose to either place the output into a new worksheet (default option), new workbook, or in the same worksheet as the data (Output Range). In the later case you have to specify the upper left cell "address" where you want your output to be pasted. Click the *Summary Statistics* box—**OK**. You will be presented with information on a range of statistical measures: mean, standard error, median, mode, standard deviation, sample variance, kurtosis, skewness, range, minimum, maximum, sum, count.

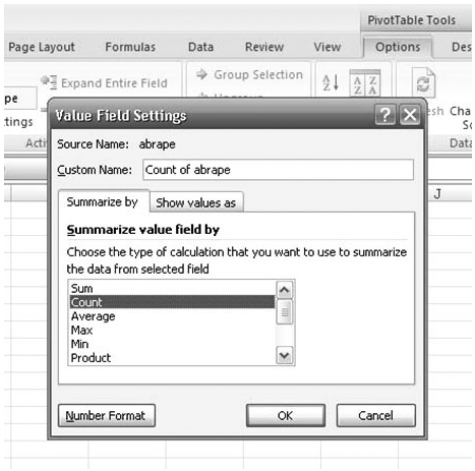

If you are interested in obtaining only a certain descriptive statistic you can take advantage of Excel's functions instead of using the Descriptive Statistics tool. For example, let's calculate the Mode of the EDUC variable. Either below or to the side of you data find some empty cells. In one of the cells type EDUC Mode. Then activate the cell immediately to the right, click **Formulas**—**Insert Function**. In the dialog window specify *Statistical* and then choose **MODE**—**OK**. In the *Function Arguments* window specify the range of the variable of interest (EDUC occupies O2:O1500). Click **OK.**

	A	B	C
1	Frequency distribution		
2	Count of abrape		
3	abrape	Total	
4	YES	365	
5	NO	102	
6	N/A	29	
7	NAP	1003	
8	Grand Total	1499	
9			
10			
11	Count of abrape		
12	abrape	Total	
13	YES	24.35%	
14	NO	6.80%	
15	N/A	1.93%	
16	NAP	66.91%	
17	Grand Total	100.00%	
18			

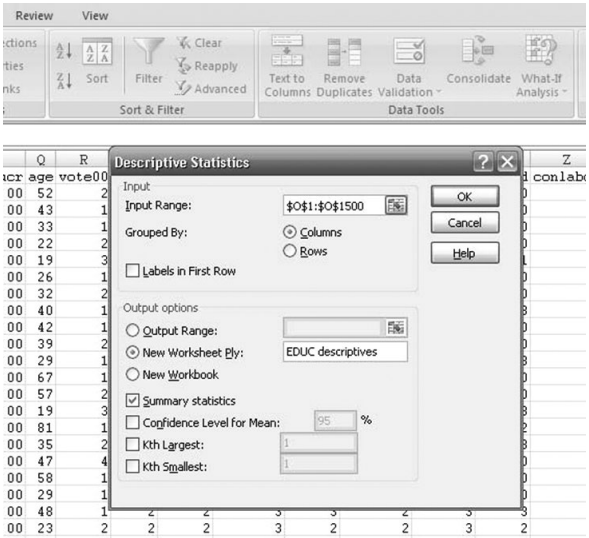

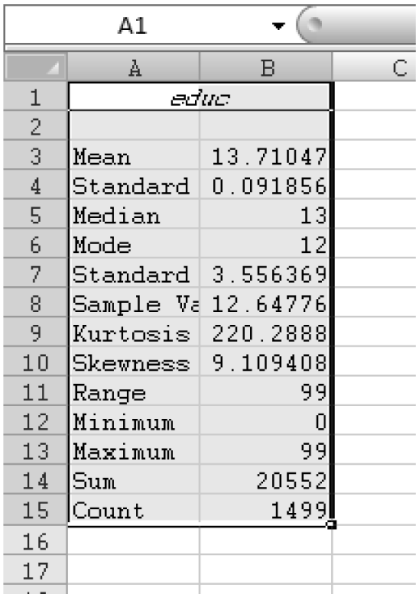

	A	B	C
1		*educ*	
2			
3	Mean	13.71047	
4	Standard	0.091856	
5	Median	13	
6	Mode	12	
7	Standard	3.556369	
8	Sample Va	12.64776	
9	Kurtosis	220.2888	
10	Skewness	9.109408	
11	Range	99	
12	Minimum	0	
13	Maximum	99	
14	Sum	20552	
15	Count	1499	
16			
17			

Graphing

There are two ways to create graphical representation of your data in Excel: using the Pivot Table or the Histogram Analysis Tool. First, let's try using the Pivot Table option. Follow the steps to create a frequency distribution Pivot Table for the variable ABRAPE. It is very easy to create a bar chart from here. Click on any cell in the Pivot Table, then click **Insert tab—column—2D column**. You can format your chart using a **Chart Styles** section under the **Design Tab.**

Now let's practice creating a histogram using the Histogram Analysis Tool. In order to create a histogram, we have to specify "bin values" that represent the entire range of values of your variable. By specifying bins we are creating slots to gauge how many times a specific value appears in our data. Our histogram is supposed to convey the age distribution of our respondents. First, create a frequency distribution of the respondent's age variable. You will see that our dataset the youngest respondent is 18 and the oldest is 89. Let's specify eight bins—28, 38, 48, 58, 68, 78, and 89—with each bin being the upper limit for a particular age group. Insert a new column in you data, title it AGEBINS, and enter the specified numbers.

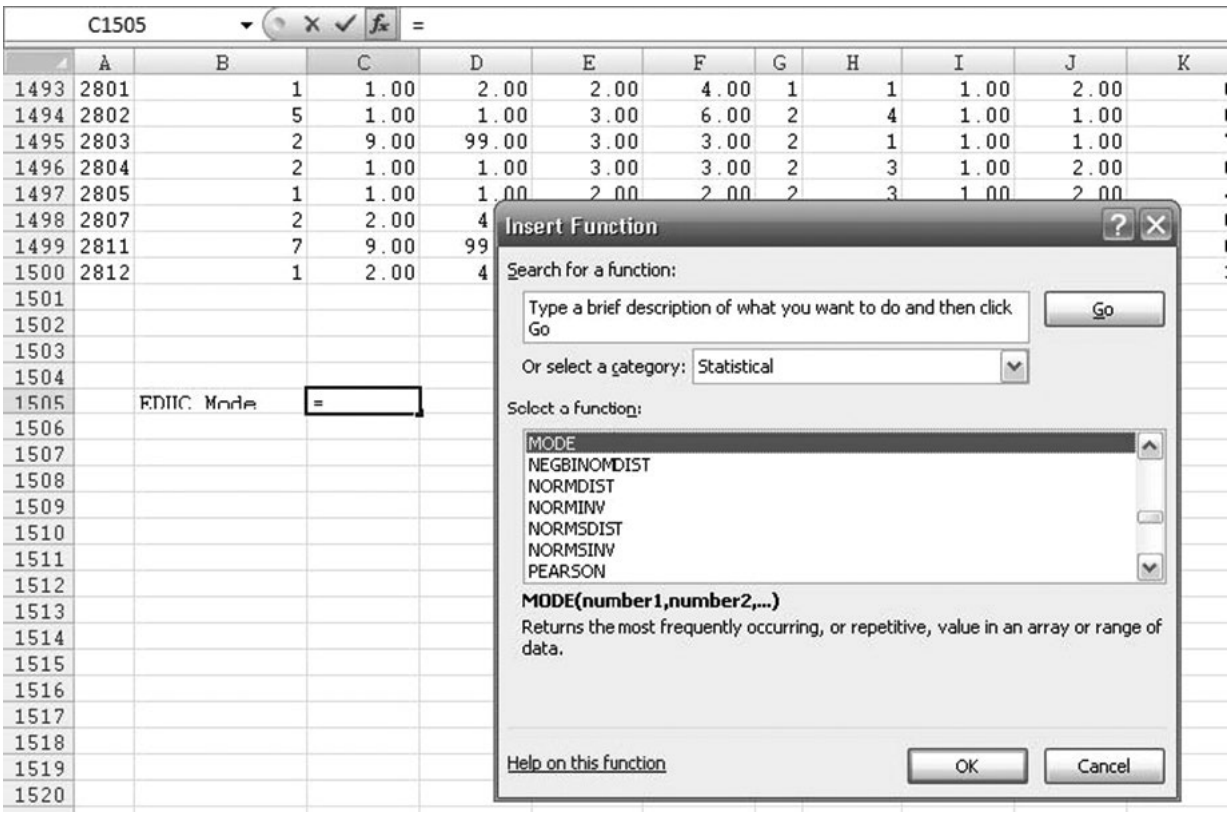

Click **Data—Data Analysis—Histogram—OK**.

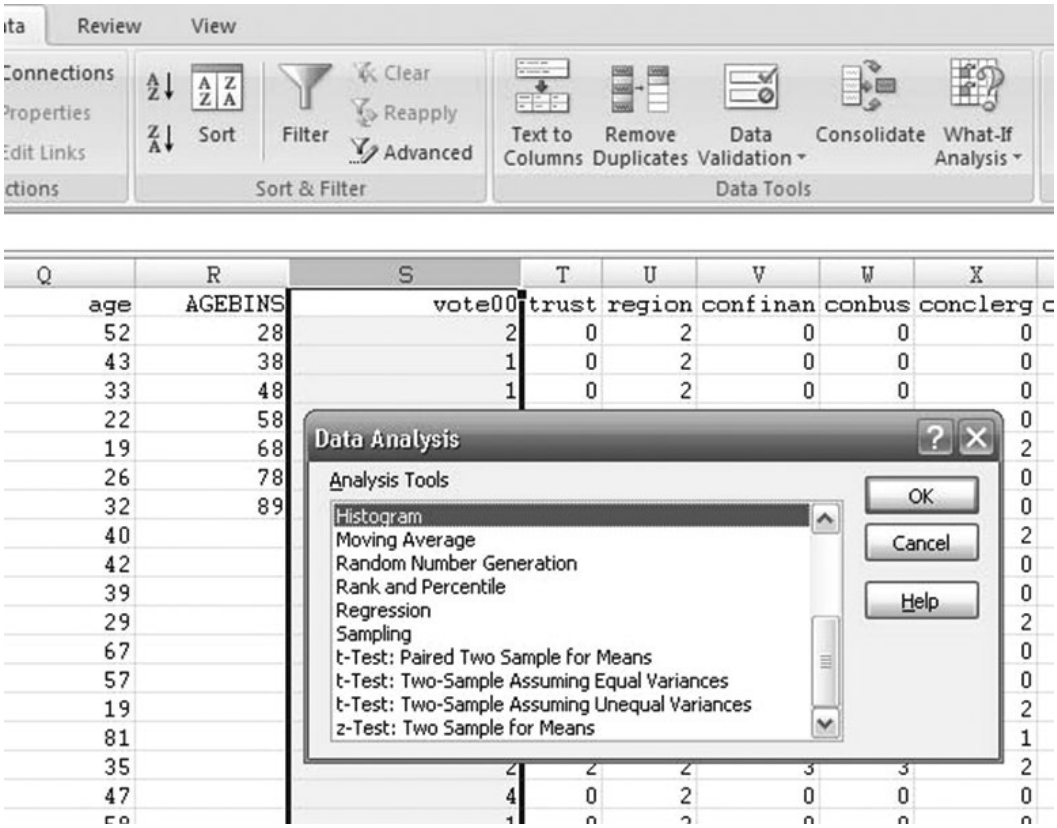

Q	R	S	T	U	V	W	X
age	AGEBINS	vote00	trust	region	confinan	conbus	conclerg
52	28	2	0	2	0	0	0
43	38	1	0	2	0	0	0
33	48	1	0	2	0	0	0
22	58						0
19	68						2
26	78						0
32	89						0
40							2
42							0
39							0
29							2
67							0
57							0
19							2
81							1
35							2
47		4	0	2	0	0	0

In the dialog window specify the *Input Range* (the range of the variable you are interested in) and the *Bin Range* (cell range where the bins values are specified). Make sure that the *Labels* box is checked. Choose the location of your output histogram (*New Worksheet Ply* is default). Check *Chart Output*. Click **OK**.

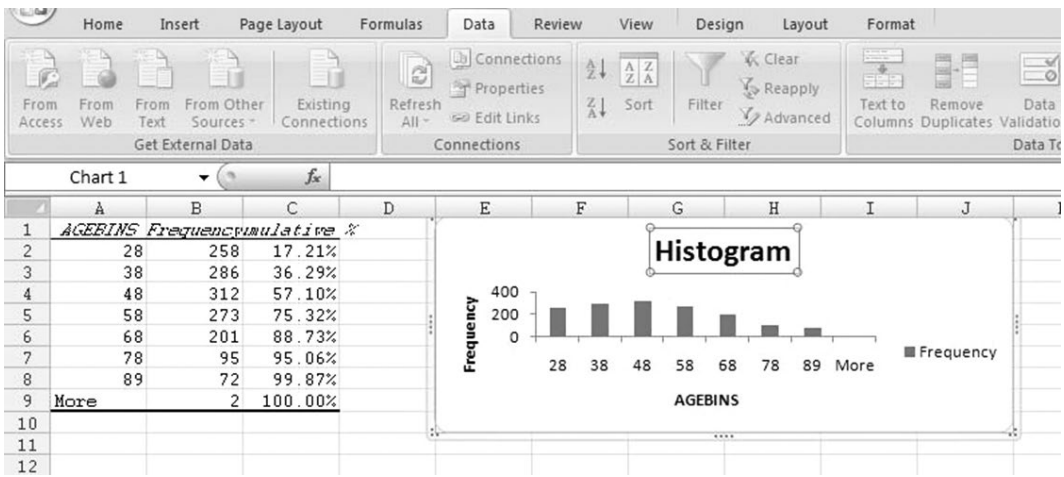

	A	B	C
1	AGEBINS	Frequency	Cumulative %
2	28	258	17.21%
3	38	286	36.29%
4	48	312	57.10%
5	58	273	75.32%
6	68	201	88.73%
7	78	95	95.06%
8	89	72	99.87%
9	More	2	100.00%
10			
11			
12			

Your histogram will open in a new worksheet. It does not look perfect, but you can modify it to make it look better. First, make the histogram taller by clicking on it and dragging its lower border down. Second, change the axis titles and the title of the histogram. Third, you can remove the space between the bars: right-click on one of the bars and select **Format Data Series**, then drag the *Gap Width* cursor all the way to the left. The possibilities of formatting a chart in Excel are endless.

2Recoding Variables

Recoding EDUC (highest level of education completed) variable. To recode EDUC into a new categorical variable you have to create a reference table that lists your new categories. Find a range of empty cells below your data and create a reference table where you specify:

Those who have 0-11 years of education are put in category 1

Those who have exactly 12 years of education put in category 2

Those who have 13 to 15 years of education are put in category 3

Those who have 16 or more years of education are put into category 4

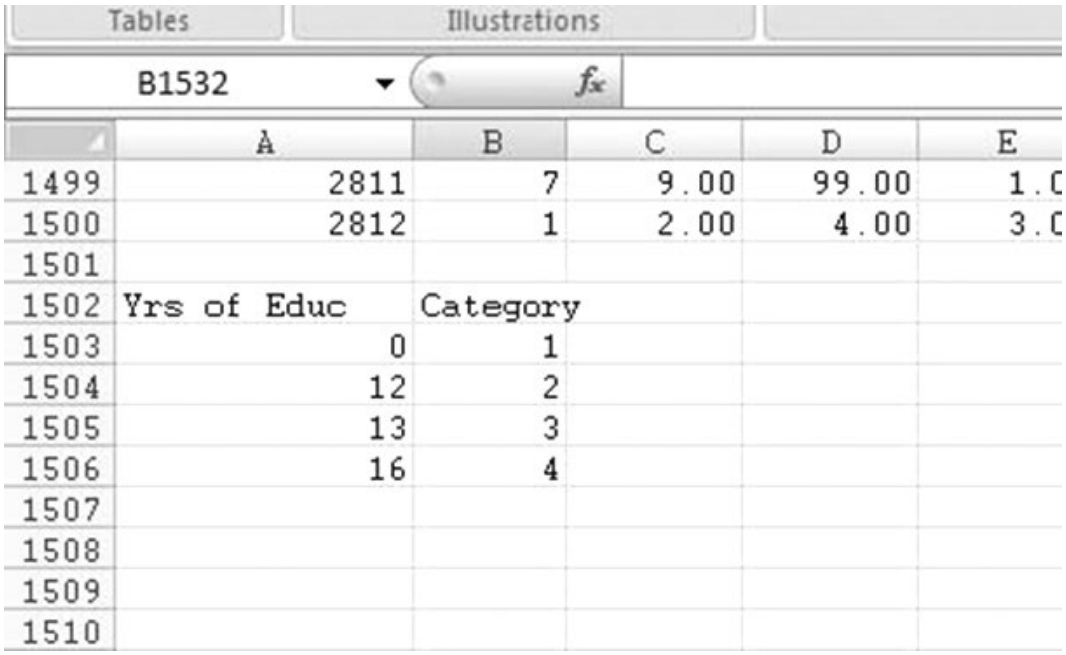

Next insert a row next to the EDUC variable, this is where you will compute your new variable—ED4CAT. Highlight the cell right under the variable label. Click: **Formulas—Insert Functions**. Select *Lookup & Reference category* and choose **VLOOKUP** function in the menu below. Click **OK**.

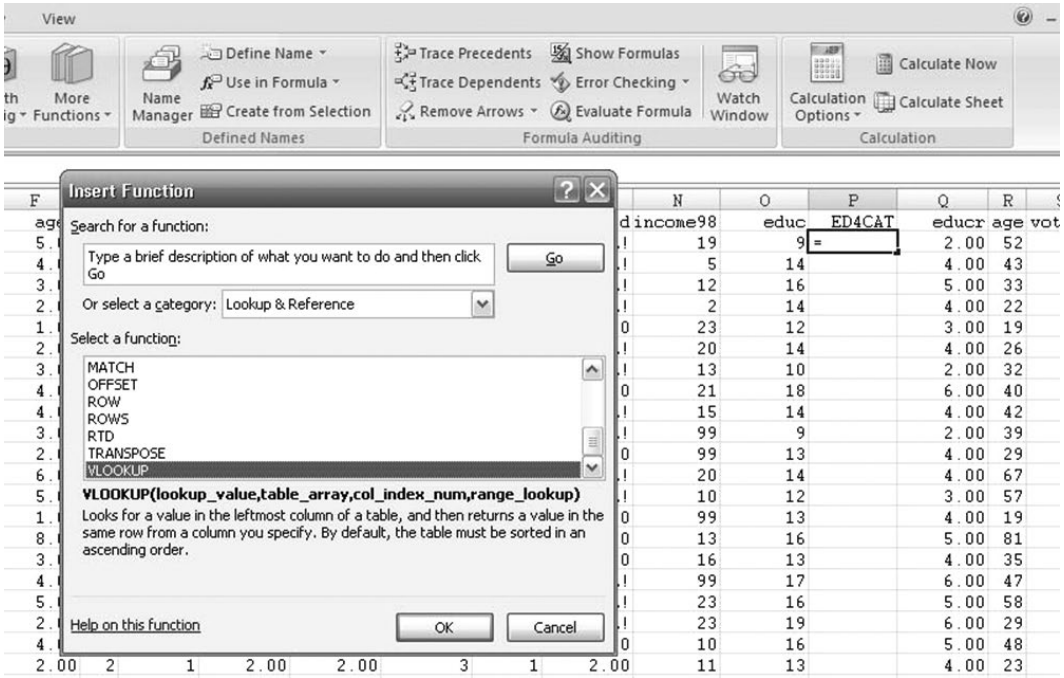

In the dialog window first specify the *Lookup_value*. The lookup value is the value of your original variable you want to look up and replace with some value of a new variable. *Table_array* is the location of the reference table we created to lookup values for the new variable. You can either highlight the location of the reference table you're your mouse or type the cell range manually. *Col_index_num* refers to the position of the reference table column we want to look up the new variable values from. In our case the column that contains all the potential values of the new variable is column number 2 (column B). Click **OK**.

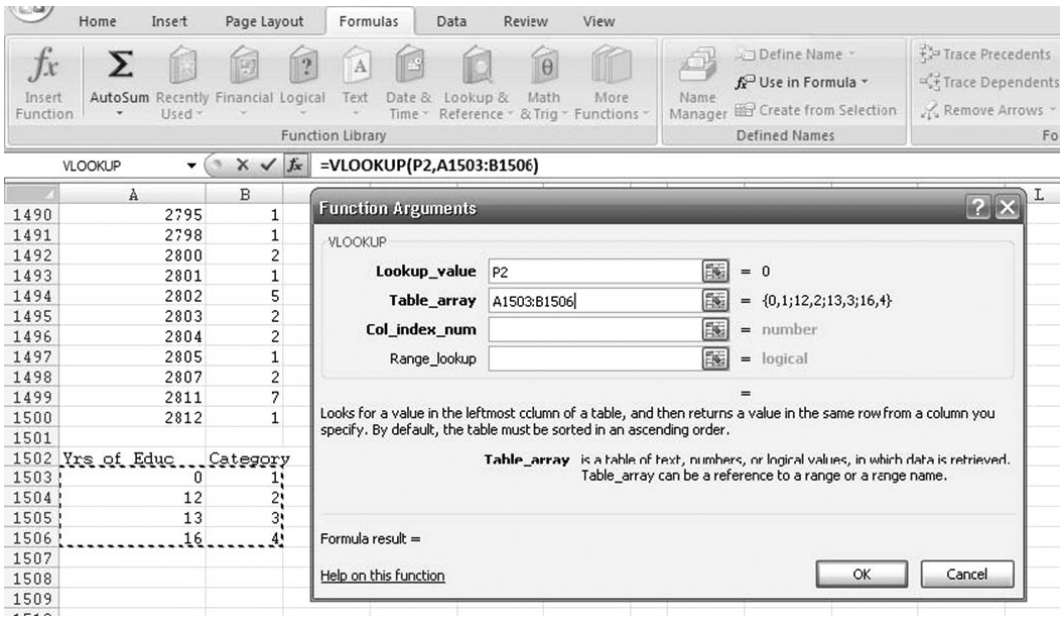

This gives you your first value of the new variable ED4CAT. To avoid repeating the above procedure for each case in the dataset, simply copy the formula from the cell just completed to other cells downward by clicking on the right bottom corner and dragging the mouse cursor down.

Computing a New Variable

It's quite easy to compute a new variable in Excel. All it takes is recollection of a few basic algebra rules. For example, to determine the age of each respondent's eldest child we need to compute the difference between the respondents' age and the age of the respondent when his or her first child was born from the respondent's. Let's start by inserting a new column anywhere in the worksheet and typing the name of our new variable (CHLDAGE) in the first row. In the empty cell right below type a formula for calculating the age of the respondent's eldest child. This formula will contain the "addresses" for cells that specify (1) the respondent's age (AGE) (Q2 in our example); and (2) respondent's age when his or her first child was born (AGEKDBRN) (R2). Our formula will be: **= Q2 – R2**. Press **Enter**, and there you have it the first respondent's eldest child is 18 years old. Copy the formula for the rest of the respondents.

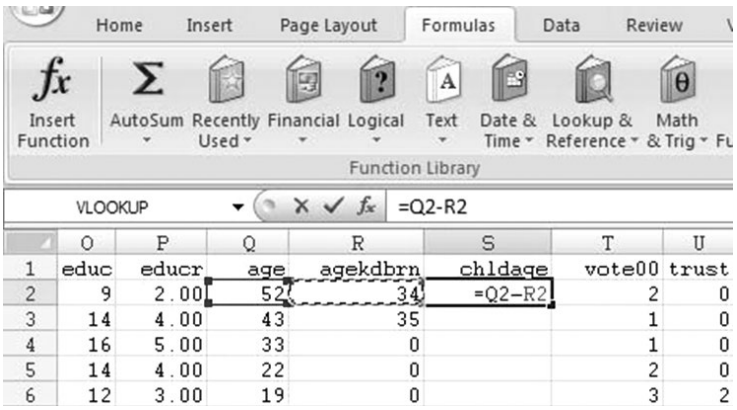

	O	P	Q	R	S	T	U
1	educ	educr	age	agekdbrn	chldage	vote00	trust
2	9	2.00	52	34	=Q2-R2	2	0
3	14	4.00	43	35		1	0
4	16	5.00	33	0		1	0
5	14	4.00	22	0		2	0
6	12	3.00	19	0		3	2

2 Bivariate and Multivariate Statistics

Crosstabulation

Exploring the relationship between two or more variables in Excel is done with the help of pivot tables. You will not find a Crosstab option like in SPSS. Let's say we are interested in the relationship between individual's level of education and whether or not he or she participates in elections. We will create a pivot table for two variables: EDUCR3 is a trichotomized measure of education and variable VOTE00, which assesses whether the respondent participated in 2000 Presidential elections.

Click **Insert** X **Pivot Table**. Select data range and the location for the Crosstab table (new worksheet). Click **OK**. On the right side of the new worksheet, in the *Pivot Table Field List* area first select **Field Section and Area Section Side-By-Side** option for more convenient display, then select (check the boxes) the two variables of interest VOTE00 and EDUCR3. Place VOTE00, the dependent variable, in the *Row Labels* box, and EDUCR3 in the *Column Labels*. Drag VOTE00 into the *Values* box, and change the **Value Field Settings** to ***Count*** instead of ***Sum***.

Your crosstab will look a little confusing due to presence of numbers instead of actual variable values. If you have the codebook for your data on hand, you can quickly change this manually. For example, we know that a value of 1 is assigned to VOTE00 variable if the respondent voted in the 2000 Presidential elections. We can change the cell accordingly.

Excel does not offer an option to simultaneously display both the count and percentage within the same crosstab. You can observe percentages (of rows or columns) by clicking **Value Field Settings—% of column—OK**.

	A	B	C	D	E
1	Drop Page Fields Here				
2					
3	Count of vote00	educr3			
4	vote00	Grade School	High School	Some College	Grand Total
5	Voted	34.27%	63.68%	69.99%	63.22%
6	Did not vote	51.64%	29.35%	22.42%	28.44%
7	Ineligible	12.21%	5.72%	7.13%	7.48%
8	Refused to answer	0.47%	0.00%	0.11%	0.13%
9	Don't know/remember	1.41%	1.24%	0.34%	0.73%
10	Grand Total	100.00%	100.00%	100.00%	100.00%
11					
12					
13					
14					
15					
16	Count of vote00	educr3			
17	vote00	Grade School	High School	Some College	Grand Total
18	Voted	73	256	618	947
19	Did not vote	110	118	198	426
20	Ineligible	26	23	63	112
21	Refused to answer	1		1	2
22	Don't know/remember	3	5	3	11
23	Grand Total	213	402	883	1498
24					

Since there were so few people that either refused to answer the question about voting or cannot recollect whether they voted, you have an option of dropping these two rows from the table (this will, of course, change the grand total counts, so you might not want to do this). Click the little arrow next to the VOTE00 cell in the table and uncheck these two rows. Click **OK**.

Three-Variable Crosstabs

It is very easy to convert the crosstab you have created into a three-variable crosstab.

The two-variable cross-tab we have created demonstrated that individuals with higher levels of education (those who have completed high school and/or some college) were more actively involved in voting in the 2000 Presidential election than individuals who have completed grade school. But perhaps education is not really a moving force behind individuals' propensity to be politically active, and the relationship we observed is due to extraneous factor. To explore the relationship

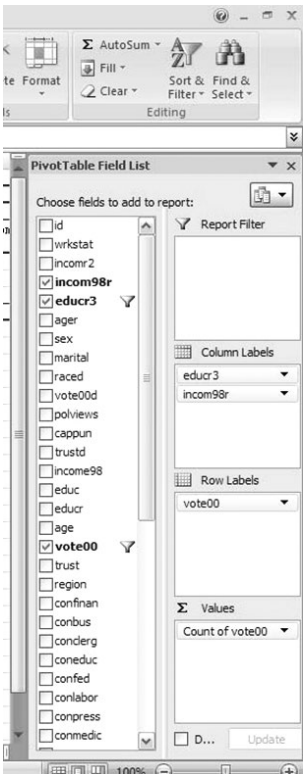

between voting and another variable while controlling for education all we have to do is add a third variable to the pivot table (to the column labels).

Comparing Means

Comparing means in Excel once again requires the use of Pivot tables. Let's examine whether individuals differ in terms of years of education by race. Click:

Insert—Pivot Table—[make sure the cell range for your data is specified in the *Table/Range* box]—**OK**.

In the *Pivot Table Field List* area check and drag the variables you are interested in (EDUC and RACE) into *Columns* and *Row Labels* boxes respectively. With your dependent variable (EDUC) also being in the *Values* box. Click **Value Field Settings—Average—OK**. This will produce a pivot table that shows mean years of education for each group within the race variable. Let's type the value labels of the race variable instead of numbers.

To accurately assess the difference between the mean years of education of Whites, Blacks, and those individuals who comprise the "Other" racial category, we need to calculate standard deviations by racial groups. This is also done in the pivot table. Let's copy the pivot table that we just created somewhere on the worksheet you have open. Now, all you have to do to obtain standard deviations by group is click **Value Field Settings** and choose **StdDev** under the *Summarize by* tab. In the figure below we have copied and pasted the Standard Deviations column adjacent to the Means column.

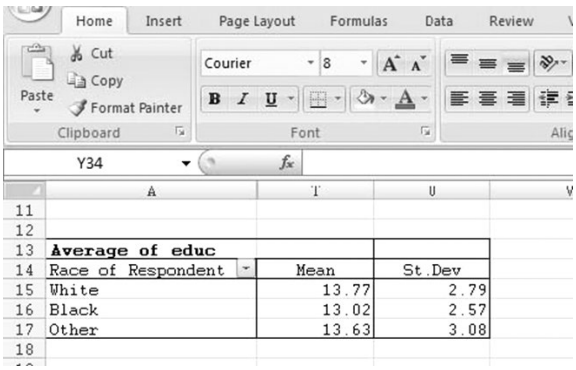

2Printing

To print portions of Excel worksheets click the **Office** button—**Print**. Before printing the **Page Setup** button under the **Page Layout** tab allows you to manipulate the orientation of the page (portrait or landscape), the margins of the page, rows and columns you want to print, as well as printing the output with or without gridlines.

Appendix F

Quantitative Data Analysis

LEARNING OBJECTIVES

1. Identify the types of graphs and statistics that are appropriate for analysis of variables at each level of measurement.
2. List the guidelines for constructing frequency distributions.
3. Discuss the advantages and disadvantages of using each of the three measures of central tendency.
4. Understand the difference between the variance and the standard deviation.
5. Define the concept of skewness and explain how it can influence measures of central tendency.
6. Explain how to calculate percentages in a cross-tabulation table and how to interpret the results.
7. Discuss the three reasons for conducting an elaboration analysis.
8. Write a statement based on inferential statistics that reports the confidence that can be placed in a statistical statement of a population parameter.
9. Define the statistics obtained in a multiple regression analysis and explain their purpose.

"Oh no, not data analysis and statistics!" We now hit the chapter that you may have been fearing all along, the chapter on data analysis and the use of statistics. This chapter describes what you need to do after your data have been collected. You now need to analyze what you have found, interpret it, and decide how to present your data so that you can most clearly make the points you wish to make.

What you probably dread about this chapter is something that you either sense or know from a previous course: Studying data analysis and statistics will lead you into that feared world of mathematics. We would like to state at the beginning, however, that you have relatively little to fear. The kind of mathematics required to perform the data analysis tasks in this chapter is minimal. If you can add, subtract, multiply, and divide and are willing to put some effort into carefully reading the chapter, you will do well in the statistical analysis of your data. In fact, it is our position that the analysis of your data will require more in the way of careful and logical thought than in mathematical skill. One helpful way to think of statistics is that

it consists of a set of tools that you will use to examine your data to help you answer the questions that motivated your research in the first place. Right now, the toolbox that holds your statistical tools is fairly empty (or completely empty). In the course of this chapter, we will add some fundamental tools to that toolbox. We would also like to note at the beginning that the kinds of statistics you will use on criminological data are very much the same as those used by economists, psychologists, political scientists, sociologists, and other social scientists. In other words, statistical tools are statistical tools, and all that changes is the nature of the problem to which those tools are applied.

This chapter will introduce several common statistics in social research and highlight the factors that must be considered in using and interpreting statistics. Think of it as a review of fundamental social statistics, if you have already studied them, or as an introductory overview, if you have not.

Two preliminary sections lay the foundation for studying statistics. In the first, we will discuss the role of statistics in the research process, returning to themes and techniques you already know. In the second preliminary section, we will outline the process of acquiring data for statistical analysis. In the rest of the chapter, we will explain how to describe the distribution of single variables and the relationships among variables. Along the way, we will address ethical issues related to data analysis. This chapter will be successful if it encourages you to see statistics responsibly and evaluate them critically and gives you the confidence necessary to seek opportunities for extending your statistical knowledge.

It should be noted that, in this chapter, we focus primarily on the use of statistics for descriptive purposes. Those of you looking for a more advanced discussion of statistical methods used in criminal justice and criminology should seek other textbooks (e.g., Bachman and Paternoster, 2017). Although many colleges and universities offer social statistics in a separate course, we don't want you to think of this chapter as something that deals with a different topic than the rest of the book. Data analysis is an integral component of research methods, and it's important that any proposal for quantitative research include a plan for the data analysis that will follow data collection.

Frequency distributions Numerical display showing the number of cases, and usually the percentage of cases (the relative frequencies), corresponding to each value or group of values of a variable.

Cross-tabulation (cross-tab) A bivariate (two-variable) distribution showing the distribution of one variable for each category of another variable.

Descriptive statistics Statistics used to describe the distribution of and relationship among variables.

Inferential statistics Mathematical tools for estimating how likely it is that a statistical result based on data from a random sample is representative of the population from which the sample is assumed to have been selected.

2 Introducing Statistics

Statistics play a key role in achieving valid research results in terms of measurement, causal validity, and generalizability. Some statistics are useful primarily to describe the results of measuring single variables and to construct and evaluate multi-item scales. These statistics include **frequency distributions**, graphs, measures of central tendency and variation, and reliability tests. Other statistics are useful primarily in achieving causal validity, by helping us describe the association among variables and control for, or otherwise take into account, other variables.

Cross-tabulation is one technique for measuring association and controlling other variables and is introduced in this chapter. All these statistics are called **descriptive statistics** because they are used to describe the distribution of and relationship among variables.

You learned in Chapter 5 that it is possible to estimate the degree of confidence that can be placed in generalizations for a sample and for the population from which the sample was selected. The statistics used in making these estimates are called **inferential statistics**, and they include confidence intervals, to which you were exposed in Chapter 5. In this chapter we will refer only briefly to inferential statistics, but we will emphasize later in the chapter their importance for testing hypotheses involving sample data.

Criminological theory and the results of prior research should guide our statistical plan or analytical strategy, as they guide the choice of other research methods. In other words, we want to use the statistical strategy that will best answer our research question. There are so many particular statistics and so many ways for them to be used in data analysis that even the best statistician can become lost in a sea of numbers if she is not using prior research and theorizing to develop a coherent analysis plan. It is also important for an analyst to choose statistics that are appropriate to the level of measurement of the variables to be analyzed. As you learned in Chapter 4, numbers used to represent the values of variables may not actually signify different quantities, meaning that many statistical techniques will be inapplicable. Some statistics, for example, will be appropriate only when the variable you are examining is measured at the nominal level. Other kinds of statistics will require interval-level measurement. To use the right statistic, then, you must be very familiar with the measurement properties of your variables (and you thought that stuff would go away!).

Case Study

The Causes of Delinquency

In this chapter, we will use research on the causes of delinquency for our examples. More specifically, our data will be a subset of a much larger study of a sample of approximately 1,200 high school students selected from the metropolitan and suburban high schools of a city in South Carolina. These students, all of whom were in the 10th grade, completed a questionnaire that asked about such things as how they spent their spare time; how they got along with their parents, teachers, and friends; their attitudes about delinquency; whether their friends committed delinquent acts; and their own involvement in delinquency. The original research study was designed to test specific hypotheses about the factors that influence delinquency. It was predicted that delinquent behavior would be affected by such things as the level of supervision provided by parents, the students' own moral beliefs about delinquency, their involvement in conventional activities such as studying and watching TV, their fear of getting caught, their friends' involvement in crime, and whether these friends provided verbal support for delinquent acts. All these hypotheses were derived from extant criminological theory, theories we have referred to throughout this book. One specific hypothesis, derived from deterrence theory, predicts that youths who believe they are likely to get caught by the police for committing delinquent acts are less likely to commit delinquency than others. This hypothesis is shown in Exhibit F.1. The variables from this study that we will use in our chapter examples are displayed in Exhibit F.2.

Exhibit F.1 Hypothesis for Perceived Fear of Being Caught and Delinquency

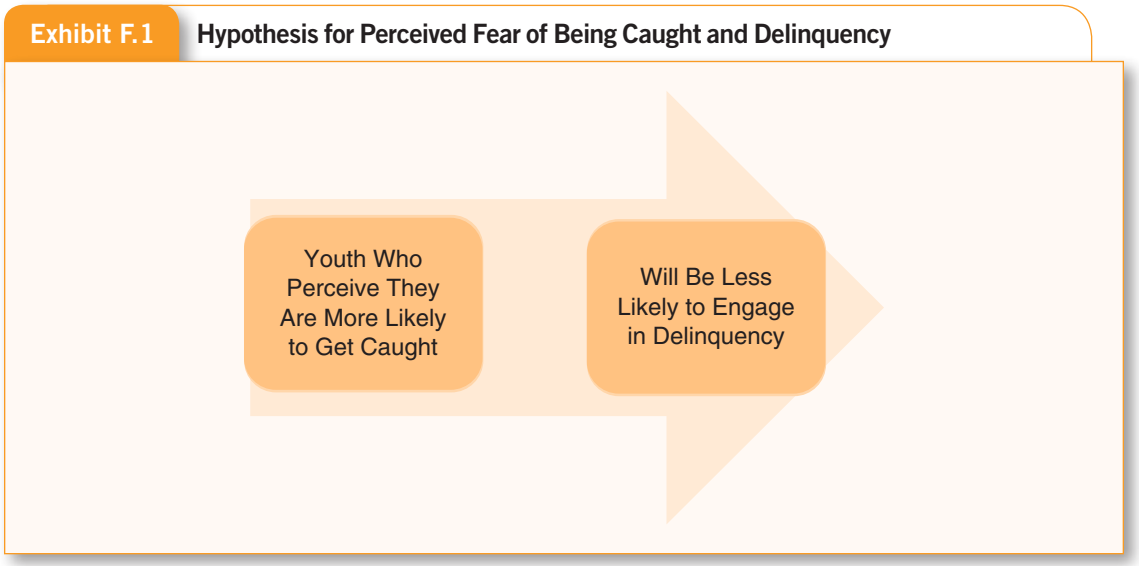

Exhibit F.2 List of Variables for Class Examples of Causes of Delinquency

Variable	*SPSS Variable Name*	*Description*
Gender	V1	Sex of respondent.
Age	V2	Age of respondent.
TV	V21	Number of hours per week the respondent watches TV.
Study	V22	Number of hours per week the respondent spends studying.
Supervision	V63	Do parents know where respondent is when he or she is away from home?
Friends think theft wrong	V77	How wrong do respondent's best friends think it is to commit petty theft?
Friends think drinking wrong	V79	How wrong do respondent's best friends think it is to drink liquor under age?
Punishment for drinking	V109	If respondent was caught drinking liquor under age and taken to court, how much of a problem would it be?
Cost of vandalism	V119	How much would respondent's chances of having good friends be hurt if he or she was arrested for petty theft?
Parental supervision	PARSUPER	Added scale from items that ask respondent if parents know where he or she is and whom he or she is with when away from home. A high score indicates high parental supervision.
Friend's opinion	FROPINON	Added scale that asks respondent if his or her best friends thought that committing various delinquent acts was all right. A high score means more support by friends for committing delinquent acts.
Friend's behavior	FRBEHAVE	Added scale that asks respondent how many of his or her best friends commit delinquent acts.
Certainty of punishment	CERTAIN	Added scale that measures how likely respondent thinks it is that he or she will be caught by police if he or she were to commit delinquent acts. A high score indicates youth perceive a greater probability of being caught.
Morality	MORAL	Added scale that measures how morally wrong respondent thinks it is to commit diverse delinquent acts. A high score means respondent has strong moral inhibitions.
Delinquency	DELINQ1	An additive scale that counts the number of times respondent admits to committing a number of different delinquent acts in the past year. The higher the score, the more delinquent acts she or he committed.

2 Preparing Data for Analysis

If you have conducted your own survey or experiment, your quantitative data must be prepared in a format suitable for computer entry. You learned in Chapter 8 that questionnaires and interview schedules can be precoded to facilitate data entry by representing each response with a unique number. This method allows direct entry of the precoded responses into a computer file, after responses are checked to ensure that only one valid answer code has been circled (extra written answers can be assigned their own numerical codes). Most survey research organizations now use a database management program to control data entry. The program prompts the data entry clerk for each response, checks the response

to ensure that it is a valid response for that variable, and then saves the response in the data file. Not all studies have used precoded data entry, however, and individual researchers must enter the data themselves. This is an arduous and time-consuming task, but not for us if we use secondary data. After all, we get the data only after they have been coded and computerized.

Of course, numbers stored in a computer file are not yet numbers that can be analyzed with statistics. After the data are entered, they must be checked carefully for errors, a process called **data cleaning**. If a data entry program has been used and programmed to flag invalid values, the cleaning process is much easier. If data are read in from a text file, a computer program must be written that defines which variables are coded in which columns, attaches meaningful labels to the codes, and distinguishes values representing missing data. The procedures for doing so vary with each specific statistical package. We used the Windows version of the Statistical Package for the Social Sciences (SPSS) for the analysis in this chapter; you will find examples of SPSS commands required to define and analyze data on the Student Study Site for this text, **edge.sagepub.com/bachmanprccj6e**.

Data cleaning The process of checking data for errors after the data have been entered in a computer file.

2 Displaying Univariate Distributions

The first step in data analysis is usually to display the variation in each variable of interest in what are called *univariate frequency distributions*. For many descriptive purposes, the analysis may go no further. Frequency distributions and graphs of frequency distributions are the two most popular approaches for displaying variation; both allow the analyst to display the distribution of cases across the value categories of a variable. Graphs have the advantage over numerically displayed frequency distributions because they provide a picture that is easier to comprehend. Frequency distributions are preferable when exact numbers of cases with particular values must be reported, and when many distributions must be displayed in a compact form.

No matter which type of display is used, the primary concern of the data analyst is to accurately display the distribution's shape—that is, to show how cases are distributed across the values of the variable. Three features of the shape of a distribution are important: **central tendency**, **variability**, and **skewness** (lack of symmetry). All three of these features can be represented in a graph or in a frequency distribution.

These features of a distribution's shape can be interpreted in several different ways, and they are not all appropriate for describing every variable. In fact, all three features of a distribution can be distorted if graphs, frequency distributions, or summary statistics are used inappropriately.

A variable's level of measurement is the most important determinant of the appropriateness of particular statistics. For example, we cannot talk about the skewness (lack of symmetry) of a qualitative variable (measured at the nominal level). If the values of a variable cannot be ordered from lowest to highest, if the ordering of the values is arbitrary, we cannot say whether the distribution is symmetric, because we could just reorder the values to make the distribution more (or less) symmetric. Some measures of central tendency and variability are also inappropriate for qualitative variables.

The distinction between variables measured at the ordinal level and those measured at the interval or ratio level should also be considered when selecting statistics to use, but social researchers differ on just how much importance they attach to this distinction. Many social researchers think of ordinal variables as imperfectly measured interval-level variables and believe that in most circumstances statistics developed for interval-level variables also provide useful summaries for ordinal variables. Other social researchers believe that variation in ordinal variables will often be distorted by statistics that assume an interval

Central tendency A feature of a variable's distribution, referring to the value or values around which cases tend to center.

Variability A feature of a variable's distribution; refers to the extent to which cases are spread out through the distribution or clustered in just one location.

Skewness A feature of a variable's distribution, referring to the extent to which cases are clustered more at one or the other end of the distribution rather than around the middle.

level of measurement. We will touch on some of the details of these issues in the following sections on particular statistical techniques.

We will now examine graphs and frequency distributions that illustrate these three features of shape. Summary statistics used to measure specific aspects of central tendency and variability will be presented in a separate section. There is a summary statistic for the measurement of skewness, but it is used only rarely in published research reports and will not be presented here.

Graphs

It is true that a picture often is worth a thousand words. Graphs can be easy to read, and they very nicely highlight a distribution's shape. They are particularly useful for exploring data, because they show the full range of variation and identify data anomalies that might be in need of further study. And good, professional-looking graphs can now be produced relatively easily with software available for personal computers. There are many types of graphs, but the most common and most useful are bar charts and histograms. Each has two axes, the vertical axis (*y*-axis) and the horizontal axis (*x*-axis), and labels to identify the variables and the values with tick marks showing where each indicated value falls along the axis. The vertical *y*-axis of a graph is usually in frequency or percentage units, whereas the horizontal *x*-axis displays the values of the variable being graphed. There are different kinds of graphs you can use to descriptively display your data, depending upon the level of measurement of the variable.

A **bar chart** contains solid bars separated by spaces. It is a good tool for displaying the distribution of variables measured at the nominal level and other discrete categorical variables, because there is, in effect, a gap between each of the categories. In our study of delinquency, one of the questions asked of respondents was whether their parents knew where the respondents were when the respondents were away from home. We graphed the responses to this question in a bar chart, which is shown in Exhibit F.3. In this bar chart we report both the frequency count for each value and the **percentage** of the total that each value represents. The chart indicates that very few of the respondents (only 16, or 1.3%) reported that their parents "never" knew where the respondents were when the respondents were not at home. Almost one half (562, or 44.3%) of the youths reported that their parents "usually" knew where the respondents were. What you can also see, by noticing the height of the bars above "usually" and "always," is that most youths report that their parents provide very adequate supervision. You can also see that the most frequent response was "usually" and the least frequent was "never." Because the response "usually" is the most frequent value, it is called the **mode** or modal response. With ordinal data like these, the mode is the most appropriate measure of central tendency (more about this later).

Bar chart A graphic for qualitative variables in which the variable's distribution is displayed with solid bars separated by spaces.

Percentage Relative frequencies, computed by dividing the frequency of cases in a particular category by the total number of cases, and multiplying by 100.

Mode The most frequent value in a distribution, also termed the probability average.

Histogram A graphic for quantitative variables in which the variable's distribution is displayed with adjacent bars.

Notice that the cases tend to cluster in the two values of "usually" and "always"; in fact, about 80% of all cases are found in those two categories. There is not much variability in this distribution, then.

A **histogram** is like a bar chart, but it has bars that are adjacent, or right next to each other, with no gaps. This is done to indicate that data displayed in a histogram, unlike the data in a bar chart, are quantitative variables that vary along a continuum (see the discussion of levels of measurement for variables in Chapter 4). Exhibit F.4 shows a histogram from the delinquency dataset we are using. The variable being graphed is the number of hours per week the respondent reported to be studying. Notice that the cases cluster at the low end of the values. In other words, there are a lot of youths who spend between 0 and 15 hours per week studying. After that, there are only a few cases at each different value, with "spikes" occurring at 25, 30, 38, and 40 hours studied. This distribution is clearly not symmetric. In a symmetric distribution there is a lump of cases or a spike with an equal number of cases to the left and right of that spike. In the distribution shown in Exhibit F.4, most of the cases are at the left end of the distribution (i.e., at low values), and the distribution trails off on the right side. The ends of a histogram

Exhibit F.3 Bar Chart Showing Youths' Reponses on Parents Knowing Where They Are

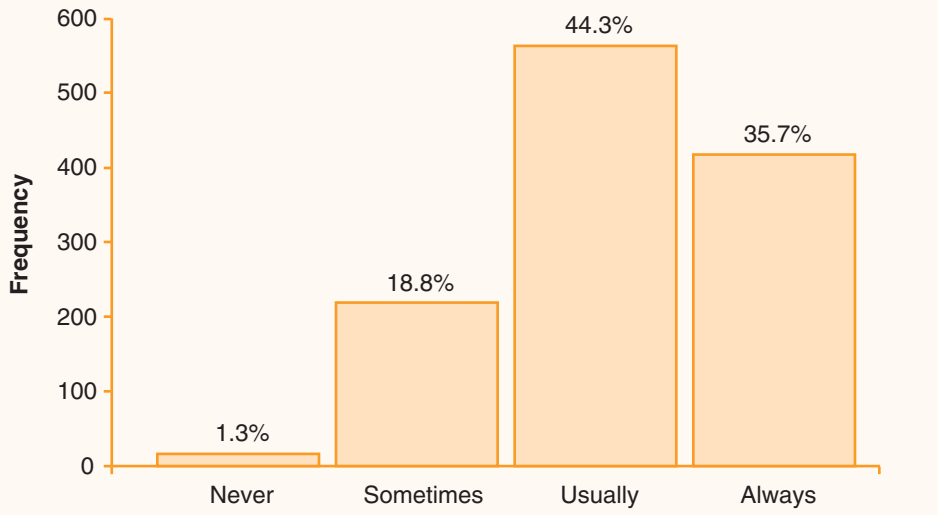

like this are often called the tail of a distribution. In a symmetric distribution, the left and right tails are approximately the same length. As you can clearly see in Exhibit F.4, however, the right tail is much longer than the left tail. When the tails of the distribution are uneven, the distribution is said to be asymmetrical or skewed. A skew is either positive or negative. When the cases cluster to the left and the right tail of the distribution is longer than the left, as in Exhibit F.4, our variable distribution is **positively skewed**. When the cases cluster to the right side and the left tail of the distribution is long, our variable distribution is **negatively skewed**.

Positively skewed Describes a distribution in which the cases cluster to the left and the right tail of the distribution is longer than the left.

Negatively skewed A distribution in which cases cluster to the right side, and the left tail of the distribution is longer than the right.

If graphs are misused, they can distort, rather than display, the shape of a distribution. Compare, for example, the two graphs in Exhibit F.5. The first graph shows that high school seniors reported relatively stable rates of lifetime use of cocaine between 1980 and 1985. The second graph, using exactly the same numbers, appeared in a 1986 *Newsweek* article on the coke plague (Orcutt and Turner, 1993). To look at this graph, you would think that the rate of cocaine usage among high school seniors increased dramatically during this period. But, in fact, the difference between the two graphs is due simply to changes in how the graphs are drawn. In the "plague" graph (B), the percentage scale on the vertical axis begins at 15 rather than 0, making what was about a one-percentage-point increase look very big indeed. In addition, omission from the plague graph of the more rapid increase in reported usage between 1975 and 1980 makes it look as if the tiny increase in 1985 were a new, and thus more newsworthy, crisis.

Adherence to several guidelines (Tufte, 1983) will help you spot these problems and avoid them in your own work:

- The difference between bars will be exaggerated if you cut off the bottom of the vertical axis and display less than the full height of the bars. Instead, begin the graph of a quantitative variable at 0 on both axes. It may at times be reasonable to violate this guideline, as when an age distribution is presented for a sample of adults, but in this case be sure to mark the break clearly on the axis.
- Bars of unequal width, including pictures instead of bars, can make particular values look as if they carry more weight than their frequency warrants. Always use bars of equal width.

Exhibit F.4 Histogram

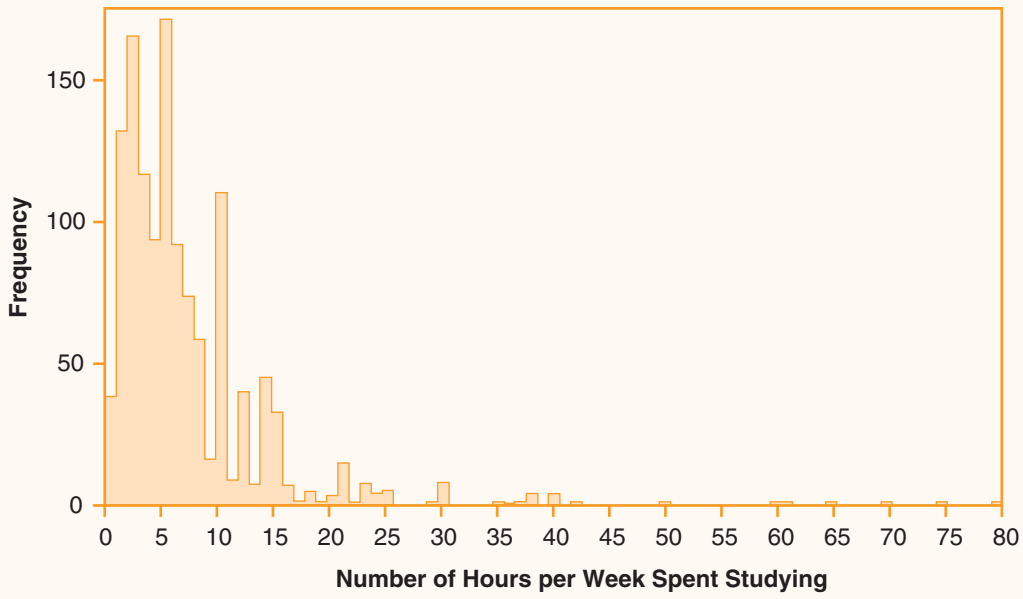

- Either shortening or lengthening the vertical axis will obscure or accentuate the differences in the number of cases between values. The two axes usually should be of approximately equal length.
- Avoid chart junk that can confuse the reader and obscure the distribution's shape (a lot of verbiage, numerous marks, lines, lots of cross-hatching, etc.).

Frequency Distributions

A frequency distribution displays the number, the percentage (the relative frequencies), or both for cases corresponding to each of a variable's values or a group of values. The components of the frequency distribution should be clearly labeled, with a title, a stub (labels for the values of the variable), a caption (identifying whether the distribution includes frequencies, percentages, or both), and perhaps the number of missing cases. If percentages are presented rather than frequencies (sometimes both are included), the total number of cases in the distribution (the **Base N**) should be indicated (see Exhibit F.6). Remember that a percentage is simply a relative frequency. A percentage shows the frequency of a given value relative to the total number of cases times 100.

Base N The total number of cases in a distribution.

Ungrouped Data

Constructing and reading frequency distributions for variables with few values is not difficult. In Exhibit F.6, we created the frequency distribution from the variable "Punishment for Drinking" found in the delinquency dataset (see Exhibit F.2). For this variable, the study asked the youths to respond to the following question: "How much of a problem would it be if you went to court for drinking liquor under age?" The frequency distribution in Exhibit F.6 shows the frequency for each value and its corresponding percentage.

Exhibit F.5 Two Graphs of Cocaine Usage

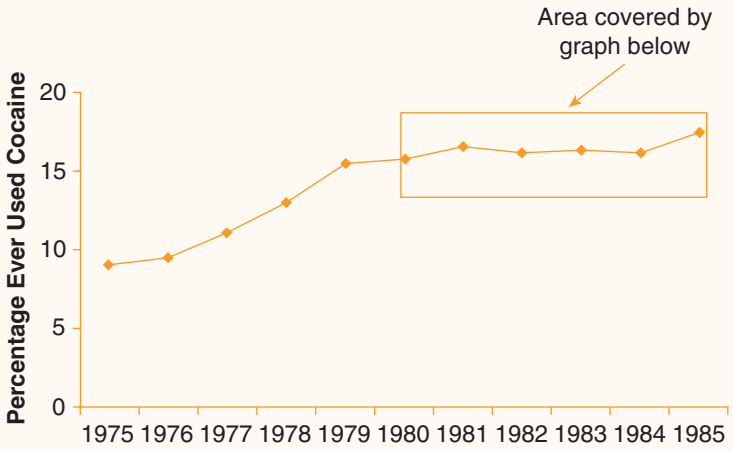

A. University of Michigan Institute for Social Research, Time Series for Lifetime Prevalence of Cocaine Use

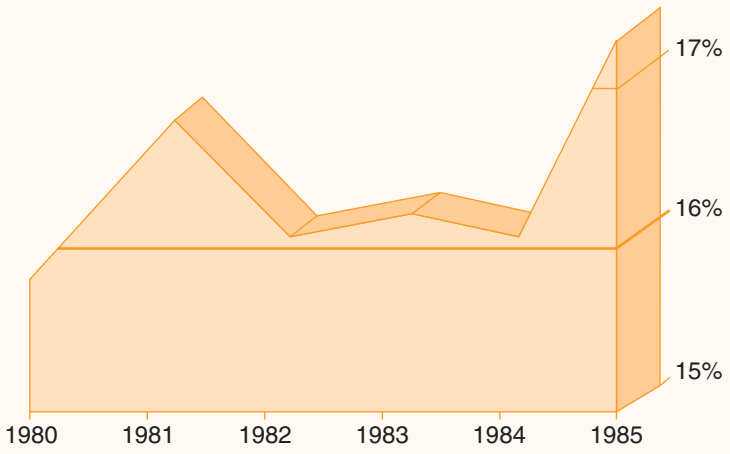

B. Final Stages of Construction

Source: James D. Orcutt and J. Blake Turner. "Shocking Numbers and Graphic Accounts." *Social Problems*, 40(2): 190–206. Copyright © 1993, The Society for the Study of Social Problems. Reprinted with permission from Oxford University Press.

Exhibit F.6 Frequency Distribution

How much of a problem would it be if you went to court for drinking liquor under age?

Value	*Frequency (f)*	*Percentage (%)*
No problem at all	14	1.1
Hardly any problem	53	4.2
A little problem	196	15.4
A big problem	421	33.1
A very big problem	588	46.2
Total	1,272	100.0

As another example of calculating the frequencies and percentages, suppose we had a sample of 25 youths and asked them their gender. From this group of 25 youths, 13 were male and 12 were female. The frequency of males (symbolized here by *f*) would be 13 and the frequency of females would be 12. The percentage of males would be 52%, calculated by *f*/ the total number of cases × 100 (13/25 × 100 = 52%). The percentage of females would be 12/25 × 100 = 48%.

In the frequency distribution shown in Exhibit F.6, you can see that only a very small number (14 out of 1,272) of youths thought that they would experience "no problem" if they were caught and taken to court for drinking liquor under age. You can see that most—in fact, 1,009—of these youths, or 79.3% of them, thought that they would have either "a big problem" or "a very big problem" with this. If you compare Exhibit F.6 to Exhibit F.3, you can see that a frequency distribution (see Exhibit F.6) can provide much of the same information as a graph about the number and percentage of cases in a variable's categories. Often, however, it is easier to see the shape of a distribution when it is graphed. When the goal of a presentation is to convey a general sense of a variable's distribution, particularly when the presentation is to an audience not trained in statistics, the advantages of a graph outweigh those of a frequency distribution.

Exhibit F.6 is a frequency distribution of an ordinal-level variable; it has a very small number of discrete categories. In Exhibit F.7, we provide an illustration of a frequency distribution with a continuous quantitative variable. This variable is one we have already looked at and graphed from the delinquency data, the number of hours per week the respondent spent studying. Notice that this variable, like many continuous variables in criminological research, has a large number of values. Although this is a reasonable frequency distribution to construct—you can, for example, still see that the cases tend to cluster in the low end of the distribution and are strung way out at the upper end—it is a little difficult to get a good sense of the distribution of the cases. The problem is that there are too many values to easily comprehend. It would be nice if we could simplify distributions like these that have a large number of different values. Well, we can. We can construct what is called a **grouped frequency distribution**.

Grouped frequency distribution A frequency distribution in which the data are organized into categories, either because there are more values than can be easily displayed or because the distribution of the variable will be clearer or more meaningful.

Grouped Data

Many frequency distributions, such as those in Exhibit F.7, and many graphs require grouping of some values after the data are collected. There are two reasons for grouping:

1. There are more than 15–20 values to begin with, a number too large to be displayed in an easily readable table.
2. The distribution of the variable will be clearer or more meaningful if some of the values are combined.

Inspection of Exhibit F.7 should clarify these reasons. In this distribution it is very difficult to discern any shape, much less the central tendency. What we would like to now do to make the features of the data more visible is change the values into intervals of values, or a range of values. For example, rather than having five separate values of 0, 1, 2, 3, and 4 hours studied per week, we can have a range of values or an interval for the first value, such as 0–4 hours studied. Then we can get a count or frequency of the number of cases (and percentage of the total) that fall within that interval.

Once we decide to group values, or categories, we have to be sure that in doing so we do not distort the distribution. Adhering to the following guidelines for combining values in a frequency distribution will prevent many problems:

- Categories should be logically defensible and preserve the distribution's shape.
- Categories should be mutually exclusive and exhaustive, so every case is classifiable in one and only one category.
- The first interval must contain the lowest value, and the last interval must contain the highest value in the distribution.

Exhibit F.7 **Frequency Distribution With Continuous Quantitative Data: Hours Studied per Week**

Value	*Frequency (f)*	*Percentage (%)*
0	38	3.0
1	132	10.4
2	165	13.0
3	116	9.1
4	94	7.4
5	171	13.4
6	92	7.2
7	73	5.7
8	58	4.6
9	16	1.3
10	110	8.6
11	9	0.7
12	40	3.1
13	7	0.6
14	45	3.5
15	32	2.5
16	7	0.6
17	5	0.4
18	4	0.3
19	1	0.1
20	15	1.2
21	8	0.6
22	1	0.1
23	1	0.1
24	4	0.3
25	5	0.4
29	1	0.1
30	8	0.6

(Continued)

(Continued)

Value	Frequency (f)	Percentage (%)
35	1	0.1
37	1	0.1
40	4	0.3
42	1	0.1
50	1	0.1
60	1	0.1
61	1	0.1
65	1	0.1
70	1	0.1
75	1	0.1
80	1	0.1
Total	1,272	100.0

- Each interval width, the number of values that fall within each interval, should be the same size.
- There should be between 7 and 13 intervals. This is a tough rule to follow. The key is not to have so few intervals that your data are clumped or clustered into only a few intervals (you will lose too much information about your distribution) and not to have so many intervals that the data are not much clearer than an ungrouped frequency distribution.

Let us use the data in Exhibit F.7 on the number of hours studied by these youths to create a grouped frequency distribution. We will follow a number of explicit steps:

Step 1. Determine the number of intervals you think you want. This decision is arbitrary, but try to keep the number of intervals you have in the 7–13 range. For our example, let us say we initially decided we wanted to have 10 intervals. (Note, if you do your frequency distribution and it looks too clustered or there are too many intervals, redo your distribution with a different number of intervals.) Don't worry; there are no hard and fast rules for the correct number of intervals, and constructing a grouped frequency distribution is as much art as science. Just remember that the frequency distribution you make is supposed to convey information about the shape and central tendency of your data.

Step 2. Decide on the width of the interval (symbolized by w_i). The interval width is the number of different values that fall into your interval. For example, an interval width of 5 has five different values that fall into it, say, the values 0, 1, 2, 3, and 4 hours studied. There is a simple formula to approximate what your interval width should be given the number of intervals you decided on in the first step: Determine the range of the data, where the range is simply the highest score in the distribution minus the lowest score. In our data, with the number of hours studied, the range is 80 because the high score is 80 and the low score is 0, so range = 80 – 0 = 80. Then determine the width of the interval by dividing the range by the number of intervals you want from Step 1. We wanted 10 intervals, so our interval width would be $w_i = 80/10 = 8$. We should therefore have an interval width of 8. If you use this simple formula for determining your interval width and you end up with a decimal, say 8.2 or 8.6, then simply round up or down to an integer.

Step 3. Make your first interval so that the lowest value falls into it. Our lowest value is 0 (for studied 0 hours per week), so our first interval begins with the value 0. Now, if the beginning of our first interval is 0 and we want an interval width of 8, is the last value of our interval 7 (with a first interval of 0–7 hours), or is the last value of our interval 8 (with a first interval of 0–8 hours)? One easy way to make a grouped frequency distribution is to do the following: Take the beginning value of your first interval (in our case, it is 0), and add the interval width to that value (8). This new value is the first value of your next interval. What we know, then, is that the first value of our first interval is 0, and the first value of our second interval is 8 (0–?, 8–?). This must mean that the last value to be included in our first interval is one less than 8, or 7. Our first interval, therefore, includes the range of values 0–7. If you count the number of different values in this interval, you will find that it includes eight different values (0, 1, 2, 3, 4, 5, 6, 7). This is our interval width of 8.

Step 4. After your first interval is determined, the next intervals are easy. They must be the same width and not overlap (mutually exclusive). You must make enough intervals to include the last value in your variable distribution. The highest value in our data is 80 hours per week, so we construct the grouped frequency distribution as follows:

0–7

8–15

16–23

24–31

32–39

40–47

48–55

56–63

64–71

72–79

80–87

Notice that in order to include the highest value in our data (80 hours) we had to make 11 intervals instead of the 10 we originally decided upon in Step 1. No problem. Remember, the number of intervals is arbitrary and this is as much art as science.

Step 5. Count the number or frequency of cases that appear in each interval and their percentage of the total. The completed grouped frequency distribution is shown in Exhibit F.8. Notice that this grouped frequency distribution conveys the important features of the distribution of these data. Most of the data cluster at the low end of the number of hours studied. In fact, more than two thirds of these youths studied less than 8 hours per week. Notice also that the frequency of cases thins out at each successive interval. In other words, there is a long right tail to this distribution, indicating a positive skew because fewer youths studied a high number of hours. Notice also that the distribution was created in such a way that the interval widths are all the same, and each case falls into one and only one interval (i.e., the intervals are exhaustive and mutually exclusive). We would have run into trouble if we had two intervals like 0–7 and 7–14, because we would not know where to place those youths who spent 7 hours a week studying. Should we put them in the first or second interval? If the intervals are mutually exclusive, as they are here, you will not run into these problems.

Exhibit F.8 Example of a Grouped Frequency Distribution From Hours Studied

Value	*Frequency (f)*	*Percentage (%)*
0–7	881	69.26
8–15	317	24.92
16–23	42	3.30
24–31	18	1.42
32–39	2	0.16
40–47	5	0.39
48–55	1	0.08
56–63	2	0.16
64–71	2	0.16
72–79	1	0.08
80–87	1	0.08
Total	1,272	100.00

Note: Total may not equal 100.0% due to rounding error.

2 Summarizing Univariate Distributions

Summary statistics, sometimes called descriptive statistics, focus attention on particular aspects of a distribution and facilitate comparison among distributions. For example, suppose you wanted to report the rate of violent crimes for each city in the United States with over 100,000 in population. You could report each city's violent crime rate, but it is unlikely that two cities would have the same rate, and you would have to report approximately 200 rates, one for each city. This would be a frequency distribution that many, if not most, people would find difficult to comprehend. One way to interpret your data for your audience would be to provide a summary measure that indicates what the average violent crime rate is in large U.S. cities. That is the purpose of the set of summary statistics called measures of *central tendency*. You would also want to provide another summary measure that shows the variability or heterogeneity in your data—in other words, a measure that shows how different the scores are from each other or from the central tendency. That is the purpose of the set of summary statistics called measures of *variation* or *dispersion*. We will discuss each type of measurement in turn.

Measures of Central Tendency

Central tendency is usually summarized with one of three statistics: the mode, the median, or the mean. For any particular application, one of these statistics may be preferable, but each has a role to play in data analysis. To choose an appropriate measure of central tendency, the analyst must consider a variable's level of measurement, the skewness of a quantitative variable's distribution, and the purpose for which the statistic is used. In addition, the analyst's personal experiences and preferences inevitably will play a role.

Mode

The mode is the most frequent value in a distribution. For example, refer to the data in Exhibit F.8, which shows the grouped frequency distribution for the number of hours studied. The value with the greatest frequency in

those data is the interval 0–7 hours; this is the mode of that distribution. Notice that the mode is the most frequently occurring value; it is not the frequency of that value. In other words, the mode in Exhibit F.8 is 0–7 hours; the mode is not 881, which is the frequency of the modal category. To show how the mode can also be thought of as the value with the highest probability, refer to Exhibit F.9. Suppose you had this grouped frequency distribution but knew nothing else about each of the 1,272 youths in the study. If you were to pick a case at random from the distribution of 1,272 youths and were asked how many hours the youth studied per week, what would your best guess be? Well, since 881 of the 1,272 youths fall into the first interval of 0–7 hours studied, the probability that a randomly selected youth studied from 0 to 7 hours would be .696 (881/1,272). This is higher than the probability of any other interval. It is the interval with the highest probability because it is the interval with the greatest frequency or mode of the distribution. When a variable distribution has one case or interval that occurs more often than the others, it is called a **unimodal distribution**. The ordinal variable of "parents knowing kids' whereabouts" in Exhibit F.3 is also unimodal. The category with the highest percentage is "usually."

Unimodal distribution A distribution of a variable in which there is only one value that is the most frequent.

Sometimes a distribution has more than one mode because there are two values that have the highest frequency. This distribution would be called **bimodal**. Some distributions are trimodal in that there are three distinctively high frequency values. When there is no frequency much higher than another, it is even possible to have a distribution without a mode. In saying that there is no mode, though, you are communicating something very important about the data: that no case is more common than the others. Another potential problem with the mode is that it might happen to fall far from the main clustering of cases in a distribution. It would be misleading in this case, then, to say simply that the variable's central tendency was the same as the modal value.

Bimodal distribution A distribution that has two nonadjacent categories with about the same number of cases, and these categories have more cases than any other categories.

Nevertheless, there are occasions when the mode is very appropriate. Most important, the mode is the only measure of central tendency that can be used to characterize the central tendency of variables measured at the nominal level. In Exhibit F.9 we have the frequency distribution of the conviction offense for 1,000 offenders convicted in a criminal court. The central tendency of the distribution is property offense, because more of the 1,000 offenders were convicted of a property crime than any other crime. For the variable "type of offense convicted of," the most common value is property crime. The mode also is often referred to in descriptions of the shape of a distribution. The terms *unimodal* and *bimodal* appear frequently, as do descriptive statements such as "The typical (most probable) respondent was in her 30s." Of course, when the issue is determining the most probable value, the mode is the appropriate statistic.

Exhibit F.9 Frequency Distribution of Offense for 1,000 Convicted Offenders

Type of Offense	*Frequency (f)*
Violent	125
Drug	210
Property	480
Public order	100
Other	85
Total	1,000

Median

The **median** is the score in the middle of a rank-ordered distribution. It is, then, the score or point that divides the distribution in half (the 50th percentile). The median is inappropriate for variables measured at the nominal level because their values cannot be put in ranked order (remember, there is no "order" to nominal-level data), and so there is no meaningful middle position. To determine the median, we simply need to do the following. First, rank-order the values from lowest to highest. Because the median is

Median The position average, or the point that divides a distribution in half (the 50th percentile).

a positional measure, we then have to find the position of the median in the rank order of scores by using the following simple formula:

$$\frac{N+1}{2}$$

where *N* is equal to the total number of cases.

In Exhibit F.10, we first list a sample of 17 U.S. cities and their rate of violent crime. We are going to calculate the median from two samples taken from this list, one sample of nine cities and another sample of eight cities.

Exhibit F.10 **Hypothetical Rate of Violent Crime for Selected U.S. Cities**

	City	*Number of Violent Crimes per 100,000*
1.	Atlanta	3,571
2.	Boston	1,916
3.	Cleveland	1,530
4.	Dallas	1,589
5.	Los Angeles	2,059
6.	New Orleans	1,887
7.	New York	1,861
8.	Philadelphia	1,322
9.	San Francisco	1,461
1.	Atlanta	3,571
2.	Boston	1,916
3.	Cleveland	1,530
4.	Dallas	1,589
5.	Los Angeles	2,059
6.	New Orleans	1,887
7.	New York	1,861
8.	Philadelphia	1,322

The first sample of nine cities is shown in Exhibit F.10a.

In this sample of nine cities, we first must find the median position, which is determined by (9 + 1)/2 = 10/2 = 5. The median violent crime rate, then, is in the fifth position in this rank order. Starting either at the top of the scores and counting down to the fifth position or at the bottom and counting up, we find that in the fifth position is the score 1,861 violent crimes per 100,000, which is the median violent crime rate for these nine U.S. cities. Now, let us find the median in the second list, which has only eight cities that are rank ordered in Exhibit F.10b.

Now our median position is: (8+1)/2 = 9/2 = 4.5. Because we now have to find the value of the median between the fourth and fifth positions, we have to find the average of the values that fall in these two positions. The score at the fourth position is 1,861, and the score at the fifth is 1,887. The value of the median can now be found by adding these two scores and dividing by 2. The median rate of violent crime for this sample of eight cities, then, is equal to (1,861 + 1,887)/2 = 1,874 violent crimes per 100,000 population. This tells us that 50% of the cities have violent crime rates lower than 1,874 and 50% of the cities have violent crime rates higher than 1,874.

Because the median is the score at the 50th percentile, we can also identify it in a frequency distribution by finding the value corresponding to a cumulative percentage of 50. We show you how to do this in Exhibit F.11. These data are a repeat of the data in Exhibit F.7, and show the number of hours studied for the youths in the delinquency dataset.

To find the 50th percentile, we simply added a new column to these data, labeled "cumulative percentage." Cumulative percentages are found by taking the percentage of the interval percentage plus all others below it. So the first value (3.0%) would be entered as the first cumulative percentage, because there are no other intervals below the first. This cumulative percentage simply means that 3% of the youths studied for 0 hours per week. Then we add the percentage in the next value (10.4%) to this to arrive at a cumulative percentage of 13.4%. This means that 13.4% of

Exhibit F.10a Sample of Nine Cities From Exhibit F.10

Rank	Crime Rate
1	1,322
2	1,461
3	1,530
4	1,589
5	1,861
6	1,887
7	1,916
8	2,059
9	3,571

Exhibit F.10b Sample of Eight Cities From Exhibit F.10

Rank	Crime Rate
1	1,322
2	1,530
3	1,589
4	1,861
5	1,887
6	1,916
7	2,059
8	3,571

the youths studied for 1 hour per week or less. This becomes the second entry in the cumulative percentage column. We continue adding each adjacent percentage value until we reach 50%. There is a cumulative percentage of 56.3% at the value of 5 hours per week. The median number of hours studied per week, then, is 5 hours. Of the respondents, 50% studied less than 5 hours per week, and 50% studied more than 5 hours per week.

Mean

The **mean** is simply the arithmetic average of all scores in a distribution. It is computed by adding up the value of all the cases and dividing by the total number of cases, thereby taking into account the value of each case in the distribution:

Mean The arithmetic, or weighted, average, computed by adding up the value of all the cases and dividing by the total number of cases.

$$\text{Mean} = \text{Sum of value of all cases / number of cases}$$

The symbol for the mean is $\bar{X}$ (pronounced "X-bar"). In algebraic notation, the equation is

$$\bar{X} = \frac{\Sigma_1^N x_i}{N}$$

where x_i is a symbol for each ith score and i's go from 1 to N; N is the total number of cases. What the algebraic equation says to do is to sum all scores, starting at the first score and continuing until the last, or Nth, score; then divide this sum by the total number of cases (N).

We will calculate the mean rate of violent crime for the nine U.S. cities listed in Exhibit F.10a:

$$\bar{X} = \frac{(1{,}322 + 1{,}461 + 1{,}530 + 1{,}589 + 1{,}861 + 1{,}887 + 1{,}916 + 2{,}059 + 3{,}571)}{9} = 1{,}910.7$$

Exhibit F.11 **Frequency Distribution With Continuous Quantitative Data: Hours Studied per Week**

Value	*Frequency (f)*	*Percentage (%)*	*Cumulative Percentage*
0	38	3.0	3.0
1	132	10.4	13.4
2	165	13.0	26.4
3	116	9.1	35.5
4	94	7.4	42.9
5	141	13.4	56.3 (includes 50th percentile)
6	92	7.2	
7	73	5.7	
8	58	4.6	
9	16	1.3	
10	110	8.6	
11	9	0.7	
12	40	3.1	
13	7	0.6	
14	45	3.5	
15	32	2.5	
16	7	0.6	
17	5	0.4	
18	4	0.3	
19	1	0.1	
20	15	1.2	
21	8	0.6	
22	1	0.1	
23	1	0.1	
24	4	0.3	
25	5	0.4	
29	1	0.1	
30	8	0.6	

Value	Frequency (f)	Percentage (%)	Cumulative Percentage
35	1	0.1	
37	1	0.1	
40	4	0.3	
42	1	0.1	
50	1	0.1	
60	1	0.1	
61	1	0.1	
65	1	0.1	
70	1	0.1	
75	1	0.1	
80	1	0.1	
Total	1,272	100.0	

The mean rate of violent crime for these nine U.S. cities, then, is 1,910.7 violent crimes per 100,000 population. When calculating the mean, we do not have to first rank-order the scores. The mean takes every score into account, so it does not matter if we add 3,571 first, in the middle, or last.

Computing the mean requires adding up the values of the cases, so it makes sense to compute a mean only if the values of the cases can be treated as actual quantities—that is, if they reflect an interval or ratio level of measurement, or if they are ordinal and we assume that ordinal measures can be treated as intervals. It would make no sense, however, to calculate the mean for the variable racial or ethnic status. Imagine a group of four people in which there were two Caucasians, one African American, and one Hispanic. To calculate the mean you would need to solve the equation (Caucasian + Caucasian + African American + Hispanic)/4 = ? Even if you decide that Caucasian = 1, African American = 2, and Hispanic = 3 for data entry purposes, it still does not make sense to add these numbers, because they do not represent real numerical quantities. In other words, just because you code Caucasian as "1" and African American as "2," that does not mean that African Americans possess twice the race or ethnicity that Caucasians possess. To see how numerically silly this is, note that we could just as easily have coded African Americans as "1" and Caucasians as "2." Now, with one arbitrary flip of our coding scheme, Caucasians have twice as much race or ethnicity as African Americans. Thus, both the median and the mean are *not* appropriate measures of central tendency for variables measured at the nominal level.

Median or Mean?

Both the median and the mean are used to summarize the central tendency of quantitative variables, but their suitability for a particular application must be carefully assessed.

The key issues to be considered in this assessment are the variable's level of measurement, the shape of its distribution, and the purpose of the statistical summary. Consideration of these issues will sometimes result in a decision to use both the median and the mean and will sometimes result in neither measure being seen as

preferable. But in many other situations, the choice between the mean and median will be clear-cut as soon as the researcher takes the time to consider these three issues.

Level of measurement is a key concern, because to calculate the mean, we must add up the values of all the cases, a procedure that assumes the variable is measured at the interval or ratio level. So even though we know that coding Agree as 2 and Disagree as 3 does not really mean that Disagree is one unit more of disagreement than Agree, the mean assumes this evaluation to be true. Calculation of the median requires only that we order the values of cases, so we do not have to make this assumption. Technically speaking, then, the mean is simply an inappropriate statistic for variables measured at the ordinal level (and you already know that it is completely meaningless for nominal variables). In practice, however, many social researchers use the mean to describe the central tendency of variables measured at the ordinal level, for the reasons outlined earlier.

The shape of a variable's distribution should also be taken into account when deciding whether to use the median or the mean. When a distribution is perfectly symmetric (i.e., when the distribution is bell shaped), the distribution of values below the median is a mirror image of the distribution of values above the median, and the mean and median will be the same. But the values of the mean and median are affected differently by skewness, or the presence of cases with extreme values on one side of the distribution but not the other side. The median takes into account only the number of cases above and below the median point, not the value of these cases, so it is not affected in any way by extreme values. The mean is based on adding the value of all the cases, so it will be pulled in the direction of exceptionally high (or low) values. When the value of the mean is larger than the median, we know that the distribution is skewed in a positive direction, with proportionately more cases with lower than higher values. When the mean is smaller than the median, the distribution is skewed in a negative direction.

The differential impact of skewness and/or outliers on the median and the mean can be illustrated with a simple thought exercise. Let's assume your class has 20 people and we ask you each to tell us your family of origin's family

Research in the News

Questions About the Article

MEDIAN LIFETIME EARNINGS

If you are feeling a bit overwhelmed and wondering whether going to college was worth it, a story from the *Washington Post* will lift your spirits. It highlights a study that utilized census data to investigate the lifetime earnings of people by their level of education. They study also examined the difference in lifetime earnings across many different college majors. If you are taking this class, you are probably not getting your major to make millions of dollars, but to help people and improve society in some way, right? The article presents a bar graph of the "median" lifetime earnings by college major. While engineering and computer science majors are at that top of the pack in terms of earnings, criminal justice and criminology majors are above many majors.

1. Why do you think the research presented median earnings rather than mean earnings over the lifetime?
2. What other statistics would you like to know from this article?

Source: Guo, Jeff. 2014. "Want Proof College is Worth It?" *Washington Post,* September 29. http://www.washingtonpost.com/news/storyline/wp/2014/09/29/want-proof-college-is-worth-it-look-at-this-list-of-the-highest-paying-majors/

income for the past year. We determine that the mean income for the families for your class members is $72,000. We also find that the median income is $54,000, which tells us that 50% of the families make less than $54,000 and 50% of families make more. Now imagine one of Bill Gate's kids enrolls in the class. Bill Gates is estimated to make over $3.5 billion annually. Wow. That makes the mean income for the class $166,735,238. Clearly, this figure does not represent the "typical" family income any longer. Notice that despite Bill Gates's child entering the class, the median family income would still remain $54,000. As you can see, the median now becomes a much better measure to use when describing the "typical" family income!

Measures of Variation

You have learned that central tendency is only one aspect of the shape of a distribution. Although the measure of center is the most important aspect for many purposes, it is still just a piece of the total picture. A summary of distributions based only on their central tendency can be very incomplete, even misleading. For example, three towns might have the same mean and median crime rate but still be very different in their social character due to the shape of the crime distributions. We show three distributions of community crime rates for three different towns in Exhibit F.12. If you calculate the mean and median crime rate for each town, you will find that the mean and median crime rate is the same for all three. In terms of its crime rate, then, each community has the same central tendency.

As you can see, however, there is something very different about these towns. Town A is a very heterogeneous town; crime rates in its neighborhoods are neither very homogeneous nor clustered at either the low or high end. Rather, the crime rates in its neighborhoods are spread out from one another. Crime rates in these neighborhoods are, then, very diverse. Town B is characterized by neighborhoods with very homogeneous crime rates; there are no real high or low crime areas, because the rate in each neighborhood is not far from the overall mean of 62.4 crimes per 1,000. Town C is characterized by neighborhoods with either very low crime rates or very high crime rates. Crime rates in the first four neighborhoods are much lower than the mean (62.4 crimes per 1,000), whereas those in the last four neighborhoods are much higher than the mean. Although they share identical measures of central tendency, these three towns have neighborhood crime rates that are very different.

The way to capture these differences is with statistical measures of variation. Four popular measures of variation are the range, the interquartile range, the variance, and the standard deviation (which is the most popular measure of variability). To calculate each of these measures, the variable must be at the interval or ratio level. Statistical measures of variation are used infrequently with qualitative variables, so statistical measures will not be presented here.

Range

The **range** is a simple measure of variation, calculated as the highest value in a distribution minus the lowest value:

Range = Highest value – Lowest value

Range The true upper limit in a distribution minus the true lower limit (or the highest rounded value minus the lowest rounded value, plus one).

It often is important to report the range of a distribution, to identify the whole range of possible values that might be encountered. However, because the range can be drastically altered by just one exceptionally high or low value (called an **outlier**), it does not do an adequate job of summarizing the extent of variability in a distribution. For our three towns in Exhibit F.12, the range in crime rates for Town A is 89.9 (109.4 – 19.5), for Town B it is 6.9 (65.0 – 58.7), and for Town C it is 106.4 (115.3 – 8.9).

Exhibit F.12 Neighborhood Crime Rates in Three Different Towns

Town A	Town B	Town C
19.5	58.1	8.9
28.2	59.7	15.4
35.7	60.1	18.3
41.9	62.7	21.9
63.2	63.2	63.2
75.8	63.9	103.5
92.0	64.2	104.2
95.7	64.5	110.7
109.4	65.0	105.3

Outlier An exceptionally high or low value in a distribution.

Interquartile range The range in a distribution between the end of the first quartile and the beginning of the third quartile.

Quartiles The points in a distribution corresponding to the first 25% of the cases, the first 50% of the cases, and the top 25% of the cases.

Interquartile Range

A version of the range statistics, the **interquartile range**, avoids the problem created by unusually high or low scores in a distribution. It is the difference between the scores at the first and third quartiles. **Quartiles** are the points in a distribution corresponding to the first 25% of the cases (the first quartile), the first 50% of the cases (the second quartile), and the first 75% of the cases (the third quartile). You already know how to determine the second quartile, corresponding to the point in the distribution covering half of the cases; it is another name for the median. The first and third quartiles are determined in the same way, but by finding the points corresponding to 25% and 75% of the cases, respectively.

Variance

If the mean is a good measure of central tendency, then it would seem that a good measure of variability would be the distance each score is away from the mean. Unfortunately, we cannot simply take the average distance of each score from the mean. One property of the mean is that it exactly balances negative and positive distances from it, so if we were to sum the difference between each score in a distribution and the mean of that distribution, it would always sum to zero. What we can do, though, is to square the difference of each score from the mean so the distance retains its value. This is the notion behind the variance as a measure of variability.

Variance A statistic that measures the variability of a distribution as the average squared deviation of each score from the mean of all scores.

The **variance** is the average square deviation of each case from the mean, so it takes into account the amount by which each case differs from the mean. The equation to calculate the variance is:

$$s^2 = \frac{\Sigma(x - \bar{X})^2}{N - 1}$$

In words, this formula says to take each score and subtract the mean, then square this difference, then sum all these differences, and then divide this sum by N or the total number of scores. Calculations for the variance for the crime rate data from Town A in Exhibit F.12 are shown in the table that follows.

x	$(x - \bar{X})$	$(x - \bar{X})^2$
19.5	(19.5 ~ 62.4) = −42.9	1,840.41
28.2	(28.2 ~ 62.4) = −34.2	1,169.64
35.7	(35.7 ~ 62.4) = −26.7	712.89
41.9	(41.9 ~ 62.4) = −20.5	420.25
63.2	(63.2 ~ 62.4) = −0.8	0.64
75.8	(75.8 ~ 62.4) = 13.4	179.56
92.0	(92.0 ~ 62.4) = 29.6	876.16
95.7	(95.7 ~ 62.4) = 33.3	1,108.89
109.4	(109.4 ~ 62.4) = 47.0	2,209.00
	$\Sigma(x - \bar{X}) = 0$	$\Sigma (x - \bar{X})^2 = 8{,}517.44$

We can now determine that the variance is

$$S^2 = \frac{8{,}517.44}{8} = 1{,}064.68$$

The variance of these data, then, is 1,064.68. In "squared deviation units," the variance tells us the amount of variation the distribution has around its mean. We had to square the original deviation units before summing them, because $\Sigma(x - \bar{X})^2 = 0$. For most people, however, it is difficult to grasp "squared deviation units." For this reason, we typically take the square root of this value, called the standard deviation, to bring the variable back to its original units of measurement.

Standard Deviation

The **standard deviation** is simply the square root of the variance. It is the square root of the average squared deviation of each case from the mean:

Standard deviation The square root of the average squared deviation of each case from the mean.

$$s = \sqrt{\frac{\Sigma(x - \bar{X})^2}{N - 1}}$$

To find the standard deviation, then, simply calculate the variance and take the square root. For our example, the standard deviation is

$$s = \sqrt{1{,}064.68} = 32.62$$

This value tells us that, on average, the neighborhood crime rates in Town A vary 32.62 around their mean of 62.4.

The standard deviation has mathematical properties that make it the preferred measure of variability in many cases. In particular, the calculation of confidence intervals around sample statistics, which you learned about in Chapter 5, relies on an interesting property of normal curves. Areas under the normal curve correspond

to particular distances from the mean, expressed in standard deviation units. If a variable is normally distributed, 68% of the cases will lie between plus and minus 1 standard deviation from the distribution's mean, and 95% of the cases will lie between 1.96 standard deviations above and below the mean. Cases that fall beyond plus or minus 1.96 standard deviations from the mean are termed outliers. Because of this property, the standard deviation tells us quite a bit about a distribution, if the distribution is normal. This same property of the standard deviation enables us to infer how confident we can be that the mean (or some other statistic) of a population sampled randomly is within a certain range of the sample mean (see Chapter 5).

2 Cross-Tabulating Variables

Most data analyses focus on relationships among variables to test hypotheses or just to describe or explore relationships. For each of these purposes, we must examine the association among two or more variables. Cross-tabulation (cross-tab) is one of the simplest methods for doing so. A cross-tabulation displays the distribution of one variable for each category of another variable; it can also be called a *bivariate distribution*. Cross-tabs also provide a simple tool for statistically controlling one or more variables while examining the associations among others. In this section, you will learn how cross-tabs used in this way can help test for spurious relationships and evaluate causal models. Cross-tabulations are usually used when both variables are measured at either the nominal or the ordinal level—that is, when the values of both variables are categories.

We are going to provide a series of examples of cross-tabulations from our delinquency data. In our first example, the independent variable we are interested in is the youth's gender (V1, see Exhibit F.2), and the dependent variable is the youth's self-reported involvement in delinquent behavior (DELINQ1). To use the delinquency variable in a cross-tabulation, however, we first need to recode it into a categorical variable. We will make three approximately equal categories of self-reported delinquency: low, medium, and high. Using the SPSS recode command, we will create another variable called DELINQ2 using the following recode commands:

$$(0 - 2 = 1)$$

$$(3 - 13 = 2)$$

$$(14 - 118 = 3)$$

Anyone who reported from none to two delinquent acts is now coded as 1, or low delinquency; anyone reporting from three to 13 delinquent acts is now coded as 2, or medium delinquency; and anyone reporting 14 or more delinquent acts is now coded as 3, or high delinquency. If you were to do a frequency distribution of this new variable, DELINQ2, you would see that there are three approximately equal groups.

We are interested in the relationship between gender and delinquency because a great deal of delinquency theory would predict that males are more likely to be delinquent than females. The gender of the youth is the independent variable, and the level of self-reported delinquency is the dependent variable.

Exhibit F.13 shows the cross-tabulation of gender with DELINQ2. Some explanation of this table is in order. Notice that there are two values of gender (male and female) that comprise the values in the two rows of the table, and three values of delinquency (low, medium, and high) that comprise the values in the three columns of the table. Cross-tabulations are usually referred to by the number of rows and columns the table has. Our cross-tabulation in Exhibit F.13 is a 2 × 3 (pronounced "two-by-three") table because there are two rows and three columns. Notice also that there are values at the end of each row and at the end of each column. These totals are referred to as the *marginals* of the table. These **marginal distributions** provide the sum of the frequencies for each column and each row of the table. For example, there are 680 females in the data and

Marginal distributions The summary distributions in the margins of a cross-tabulation that correspond to the frequency distribution of the row variable and of the column variable.

Exhibit F.13 **Cross-Tabulation of Respondents' Gender by Delinquency**

		Self-Reported Delinquency			
		Low	*Medium*	*High*	*Total*
Gender	Female	275 40.4%	182 26.8%	223 32.8%	680 100%
	Male	175 29.6%	166 28.0%	251 42.4%	592 100%
	Total	450	348	474	1,272

592 males. These row marginals should sum to the total number of youths in the dataset: 1,272. There are 450 youths who are low in delinquency, 348 youths who are medium in delinquency, and 474 youths who are high on the delinquency variable. These column marginals should also sum to the total number of youths in the dataset: 1,272.

Now notice that there are 2 × 3 or 6 data entries in the table (let us ignore the percentages for now). These data entries are called the *cells* of the cross-tabulation and represent the joint distribution of the two variables: gender and delinquency. The table in Exhibit F.13 has six cells for the joint distribution of two levels of gender with three levels of delinquency. In other words, notice where the value for female converges with the value of low for delinquency. You see a frequency number of 275 in this cell. This frequency is how many times there is the joint occurrence of a female and low delinquency; it shows that 275 females were also low in delinquency. Moving to the cell to the right of this, we see that there are 182 females who were medium in delinquency, and moving to the right again we see that there are 223 females who were high in delinquency. The sum of these three numbers is equal to the total number of females, 680. The row for the males shows the joint distribution of males with each level of delinquency.

What we would like to know is whether there is a relationship or an association between gender and delinquency. In other words, are males more likely to be delinquent than females? Because raw frequencies can provide a deceptive picture, we determine whether there is any relationship between our independent and dependent variables by looking at the percentages. Keep in mind that the idea in looking at relationships is that we want to know if variation on the independent variable has any effect on the dependent variable. To determine this, what we always do in cross-tabulation tables is to calculate our percentages on each value of the independent variable. For example, notice that in Exhibit F.13, gender is our independent variable. We calculated our percentages so that for each value of gender the percentages sum to 100% at the end of each row. The percentages for both females and males, therefore, sum to 100% at the end of the row. Now we take a given category of the dependent variable and ask what percentage of each independent variable value falls into that category of the dependent variable. Another way to say this is that we calculate our percentages on the independent variable and compare them to percentages on the dependent variable. We compare the percentages for different levels of the independent variable on the same category or level of the dependent variable.

In Exhibit F.13, for example, notice that 40.4% of the female youths were low in delinquency, but only 29.6% of the males were low. This tells us that females are more likely to be low in delinquency than males. Now let us look at the *high* category. We can see that 32.8% of the females were high in delinquency and 42.4% of the males were high. Together, this tells us that females are more likely to be low in delinquency and males are more likely to be high in delinquency. There is, then, a relationship between gender and delinquency. Also notice that the independent variable was the row variable and the dependent variable was the column variable. It does not always have to be this way; the independent variable could just as easily have been the column variable. The important general rule to remember is to always calculate your percentages on the levels of the independent variable (e.g., use marginal totals for the independent variable as denominators), and compare percentages on a level of the dependent variable.

In Exhibit F.14, we report the same data as in Exhibit F.13, this time switching the rows and the columns. Now, the independent variable (gender) is the column variable, so we calculate our percentage going down each of the two columns. We then compare percentages across rows. For example, we still see that 40.4% of the females were low in delinquency, whereas only 29.6% of the males were. And 42.4% of the males were high in delinquency, but only 32.8% of the females were high in delinquency.

Describing Association

A cross-tabulation table reveals four aspects of the association between two variables:

- *Existence*. Do the percentage distributions vary at all among categories of the independent variable?
- *Strength*. How much do the percentage distributions vary among categories of the independent variable?
- *Direction*. For quantitative variables, do values on the dependent variable tend to increase or decrease with an increase in value of the independent variable?
- *Pattern*. For quantitative variables, are changes in the percentage distribution of the dependent variable fairly regular (simply increasing or decreasing), or do they vary (perhaps increasing, then decreasing, or perhaps gradually increasing, then rapidly increasing)?

Exhibit F.14 shows that an association exists between delinquency and gender, although we can say only that it is a modest association. The percentage difference at the low and high ends of the delinquency variables is approximately 10 percentage points.

We provide another example of a cross-tabulation in Exhibit F.15. This is a 3 × 3 table that shows the relationship between how morally wrong a youth thinks delinquency is (the independent variable) and his or her self-reported involvement in delinquency (the dependent variable). This table reveals a very strong relationship between moral beliefs and delinquency. We can see that 5.6% of youths with weak moral beliefs are low on delinquency; this increases to 33.8% for those with medium beliefs and to 62.8% for those with strong moral beliefs. At the high end, over two thirds (72.1%) of those youths with weak moral beliefs are high in delinquency, 29.4% of those with medium moral beliefs are high in delinquency, and only 16.9% of those youths with strong moral beliefs are high in delinquency. Clearly, then, having strong moral beliefs serves to effectively inhibit involvement in delinquent behavior. This is exactly what control theory would have us believe.

Exhibit F.14 **Cross-Tabulation of Respondents' Delinquency by Gender**

		Gender		
		Female	*Male*	*Total*
Self-Reported Delinquency	Low	275 40.4%	175 29.6%	450
	Medium	182 26.8%	166 28.0%	348
	High	223 32.8%	251 42.4%	474
	Total	680 100%	592 100%	1,272

Exhibit F.15 **Cross-Tabulation of Respondents' Morals by Delinquency**

		Self-Reported Delinquency			
		Low	*Medium*	*High*	*Total*
Morals	Weak	20 5.6%	79 22.3%	256 72.1%	355 100%
	Medium	170 33.8%	185 36.8%	148 29.4%	503 100%
	Strong	260 62.8%	84 20.3%	70 16.9%	414 100%
	Total	450	348	474	1,272

Exhibit F.15 shows an example of a negative relationship between an independent and a dependent variable. As the independent variable increases (i.e., as one goes from weak to strong moral beliefs), the likelihood of delinquency decreases (one becomes less likely to commit delinquency). The independent and dependent variables move in opposite directions, so this is a negative relationship. The pattern in this table is close to what is called monotonic. In a **monotonic relationship**, the value of cases consistently increases (or decreases) on one variable as the value of cases increases (or decreases) on the other variable. *Monotonic* is often defined a bit less strictly, with the idea that as the value of cases on one variable increases (or decreases), the value of cases on the other variable tends to increase (or decrease), and at least does not change direction. This describes the relationship between moral beliefs and delinquency. Delinquency is most likely when moral beliefs are low, less likely when moral beliefs are medium, and least likely when moral beliefs are strong.

Monotonic relationship A pattern of association in which the value of cases on one variable increases or decreases fairly regularly across the categories of another variable.

We present another cross-tabulation table for you in Exhibit F.16. This table shows the relationship between the variable "number of hours studied" and the variable "certainty of punishment" (see Exhibit F.2). Both variables were originally continuous variables that we recoded into three approximately equal groups for this example. We hypothesize that those youths who study more will have a greater perceived risk of punishment than those who study less, so hours studied is our independent variable and certainty is the dependent variable. Comparing levels of hours studied for those with high certainty, we see that there is not much variation. Of those who did not study very much (0–3 hours), 39.2% were high in perceived certainty. Of those who studied from 4 to 6 hours, 35.6% were high in perceived certainty, and 40.3% of those who studied more than 7 hours per week were high in perceived certainty. Much the same levels prevail at low levels of perceived certainty. Those who do not study very much are no more or less likely to perceive a low certainty of punishment than those who study a lot. Variation in the independent variable, then, is not related to variation in the dependent variable. It looks like there is no association between the number of hours a youth studies and the extent to which he or she thinks punishment for delinquent acts is certain.

You will find when you read research reports and journal articles that social scientists usually make decisions about the existence and strength of association on the basis of more statistics than just percentage differences in a cross-tabulation table. A **measure of association** is a type of descriptive statistic used to summarize the strength of an association. There are many measures of association, some of which are appropriate for variables measured at particular levels. One popular measure of association in cross-tabular analyses with variables measured at the ordinal level is **gamma**. As with many

Measure of association A type of descriptive statistic that summarizes the strength of an association.

Gamma A measure of association sometimes used in cross-tabular analyses.

Exhibit F.16 Cross-Tabulation of Respondents' Hours Studied and Perceived Certainty of Punishment

		Certainty of Punishment			
		Low	*Medium*	*High*	*Total*
Number of Hours Studied	0–3 Hours	126 27.9%	148 32.8%	177 39.2%	451 100%
	4–6 Hours	117 32.8%	113 31.7%	127 35.6%	357 100%
	7 + Hours	129 27.8%	148 31.9%	187 40.39%	464 100%
	Total	372	409	491	1,272

measures of association, the possible values of gamma vary from –1, meaning the variables are perfectly associated in a negative direction; to 0, meaning there is no association of the type that gamma measures; to +1, meaning there is a perfect positive association of the type that gamma measures.

Inferential statistics are used in deciding whether it is likely that an association exists in the larger population from which the sample was drawn. Even when the association between two variables is consistent with the researcher's hypothesis, it is possible that the association was just due to chance or to the vagaries of sampling on a random basis. (Of course, the problem is even worse if the sample is not random.) It is conventional in statistics to avoid concluding that an association exists in the population from which the sample was drawn unless the probability that the association was due to chance is less than 5%. In other words, a statistician normally will not conclude that an association exists between two variables unless he or she can be at least 95% confident that the association was not due to chance. This is the same type of logic that you learned about in Chapter 5, which introduced the concept of 95% confidence limits for the mean. Estimation of the probability that an association is not due to chance will be based on one of several inferential statistics, **chi-square** being the one used in most cross-tabular analyses. The probability is customarily reported in a summary form such as "$p < .05$," which can be translated as "the probability that the association was due to chance is less than 5 out of 100 [5%]."

Chi-square An inferential statistic used to test hypotheses about relationships between two or more variables in a cross-tabulation.

When an association passes muster in this way, when the analyst feels reasonably confident (at least 95% confident) that it was not due to chance, it is said that the association is statistically significant. **Statistical significance** means that an association is not likely to be due to chance, according to some criterion set by the analyst. Convention (and the desire to avoid concluding that an association exists in the population when it does not) dictates that the criterion be a probability less than 5%.

Statistical significance A relationship that is not likely to be due to chance, judged by a criterion set by the analyst (often that the probability is less than 5 out of 100, or $p < .05$).

But statistical significance is not everything. You may remember from Chapter 5 that sampling error decreases as sample size increases. For this same reason, an association is less likely to appear on the basis of chance in a larger sample than in a smaller sample. In a table with more than 1,000 cases, such as those involving the delinquency dataset, the odds of a chance association are often very low indeed. For example, with our table based on 1,272 cases, the probability that the association between gender and delinquency (see Exhibit F.14) was due to chance was less than 1 in 1,000 ($p < .001$)! The association in that table was fairly weak, as indicated by a gamma of .20. Even weak associations can be statistically significant with such a large sample, which means that the analyst must be careful not to assume that just because a statistically significant association exists, it is therefore important. In a large sample, an association may be statistically significant but still be too weak to be substantively significant or important. All this boils down to another reason for evaluating carefully both the existence and the strength of an association.

Controlling for a Third Variable

Cross-tabulation can also be used to study the relationship between two variables while controlling for other variables. We will focus our attention on controlling for a third variable in this section, but we will say a bit about controlling for more variables at the section's end. We will examine three different uses for three-variable cross-tabulation: identifying an intervening variable, testing a relationship for spuriousness, and specifying the conditions for a relationship. Each of these uses for three-variable cross-tabs helps determine the validity of our findings, either by evaluating criteria for causality (nonspuriousness and identification of a causal mechanism) or by increasing our understanding of the conditions required for a relationship to hold, an indication of the cross-population generalizability of the findings. All three uses are aspects of **elaboration analysis**—the process of introducing control variables into a bivariate relationship in order to better understand—to elaborate the relationship (Rosenberg, 1968). We will examine the gamma and chi-square statistics for each table in this analysis.

Elaboration analysis The process of introducing a third variable into an analysis in order to better understand—to elaborate—the bivariate (two-variable) relationship under consideration; additional control variables also can be introduced.

Intervening Variables

We have already discovered that females are less likely to be delinquent than males (see Exhibit F.14). Finding this relationship between gender and delinquency is just the beginning of our work, however. What we would now like to know and investigate is why this relationship exists. What is it about females that makes them less likely to commit delinquent acts than males? Let us first rule out strictly biological factors and explore some possible social reasons for this gender difference in delinquency. One possibility is that because they are more closely supervised than males, females have fewer opportunities to be delinquent. In other words, females are under more strict parental supervision, and it is because they are under more strict supervision that they are less likely than males to commit delinquency. This possible relationship is shown in Exhibit F.17. Notice that in this relationship the variable "parental supervision" intervenes between gender and delinquency. It explains why females are at lower risk for delinquency compared to males. To determine whether parental supervision intervenes in the relationship between gender and delinquency and whether it explains this relationship, we must examine the relationship between gender and delinquency while controlling for difference in parental supervision. If parental supervision intervenes in the gender-delinquency relationship, the effect of controlling for this third variable would be to eliminate, or at least substantially reduce, the original relationship between gender and delinquency.

To examine this possibility, we first recode the parental supervision variable (PARSUPER; see Exhibit F.2) into two approximately equal levels: weak supervision and strong supervision. We then look at two **subtables** of the gender-delinquency relationship: once under the condition of weak parental supervision and once under the condition of strong parental supervision (see Exhibit F.17). For ease of presentation, we will report only the cell percentages and not the frequencies. What we see is that once parental supervision is controlled, there is no real relationship between gender and delinquency. That is, if males and females have the same amount of supervision from their parents, they do not differ that much in their risk of being delinquent. For example, among females with weak parental supervision, 46.0% are high in delinquency, and among males with weak parental supervision, 49.6% are high in delinquency. There is less than four percentage points' difference between males and females in their risk of being high delinquents under these conditions. Among those with strong parental supervision, 19.8% of the females were high in delinquency and 23.6% of the males were high, less than four percentage points' difference.

Subtables Tables describing the relationship between two variables within the discrete categories of one or more other control variables.

This percentage analysis is borne out by the chi-square tests and measures of association. Under both the weak and strong levels of parental supervision, the relationship between gender and delinquency is not significant, and gamma is only .067 when supervision is weak and .136 when supervision is strong. In neither case is the obtained gamma very different from zero (indicating no relationship). Collectively, these results would lead us to the conclusion that parental supervision intervenes in the relationship between gender and delinquency. A very important reason females are less delinquent than males, therefore, is that females are under stricter supervision from their parents than are males, and strong parental supervision leads to a reduced risk of delinquency.

Exhibit F.17 Cross-Tabulation of Respondents' Gender by Delinquency Within Levels of Parental Supervision

Weak Parental Supervision

		Self-Reported Delinquency			
Gender		*Low*	*Medium*	*High*	*Total*
	Female	26.1%	27.9%	46.0%	337
	Male	23.2%	27.2%	49.6%	427
	Total				764

$\chi^2 = 1.220$ ($p > .05$), Gamma = .067

Strong Parental Supervision

		Self-Reported Delinquency			
Gender		*Low*	*Medium*	*High*	*Total*
	Female	54.4%	25.7%	19.8%	343
	Male	46.1%	30.3%	23.6%	165
	Total				508

$\chi^2 = 3.193$ ($p > .05$), Gamma = .136

Extraneous Variables

Another reason for introducing a third variable into a bivariate relationship is to see whether the original relationship is spurious due to the influence of an **extraneous variable**, which is a variable that causes both the independent and dependent variables. The only reason the independent and dependent variables are related, therefore, is that they both are the effects of a common cause (another independent variable).

Extraneous variable A variable that influences both the independent and dependent variables so as to create a spurious association between them that disappears when the extraneous variable is controlled.

Exhibit F.18 shows what a spurious relationship would look like. In this case, the relationship between *x* and *y* exists only because both are the effects of the common cause *z*. Controlling for *z*, therefore, will eliminate the *x*-*y* relationship. Ruling out possible extraneous variables will help considerably strengthen the conclusion that the relationship between the independent and dependent variables is causal, particularly if all the variables that seem to have the potential for creating a spurious relationship can be controlled.

Notice that if a variable is acting as an extraneous variable, then controlling for it will cause the original relationship between the independent and dependent variables to disappear or substantially diminish. This was also the empirical test for an intervening variable. Therefore, the difference between intervening and extraneous variables is a logical one and not an empirical one. In both instances, controlling for the third variable will cause the original relationship to diminish or disappear. There should, therefore, be sound theoretical grounds for suspecting that a variable is acting as an intervening variable, explaining the relationship between the independent and dependent variables.

As an example of a possible extraneous relationship, we will look at the association between a youth's perception of the certainty of punishment and self-reported involvement in delinquency. Deterrence theory should lead us to predict a negative relationship between perceived certainty and delinquency. Indeed, this is exactly what we

observe in our delinquency data. We will not show you the cross-tabulation table, but when we looked at the relationship between perceived certainty and delinquency, we found that 53.2% of youth who were low in certainty were high in delinquency; 39.1% of those who perceived medium certainty were high in delinquency; and only 23.6% of those who perceived a high certainty of punishment were high in delinquency. Youth who believed they would get caught if they engaged in delinquency, then, were less likely to be delinquent. The gamma value for this table was –.382, indicating a moderate negative relationship between perceived certainty and delinquency, exactly what deterrence theory would lead us to expect.

Exhibit F.18 Example of a Spurious Relationship

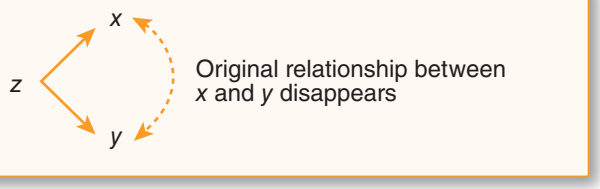

Someone may reasonably argue, however, that this discovered negative relationship may not be causal but instead may be spurious. It could be suggested that what is actually behind this relationship is the extraneous variable, moral beliefs. The argument is that those with strong moral inhibitions against committing delinquent acts think that punishment for morally wrongful actions is certain *and* refrain from delinquent acts. Thus, the observed negative relationship between perceived certainty and delinquency is really due to the positive effect of moral beliefs on perceived certainty and the negative effect of moral beliefs on delinquency (see Exhibit F.19). If moral beliefs are actually the causal factor at work, then controlling for them will eliminate or substantially reduce the original relationship between perceived certainty and delinquency.

Exhibit F.19 A Spurious Relationship Between x and y

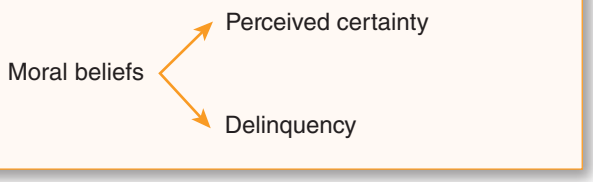

To look at this possibility, we examined the relationship between perceived certainty and delinquency under three levels of moral beliefs (weak, medium, and strong). The cross-tabulations are shown in Exhibit F.20. What we can see is that in each of the subtables there is a negative and significant association between the perceived certainty of punishment and delinquency. In two of the three subtables, however, the relationship is weaker than what was in the original table (there the gamma was –.382); we obtained gammas of –.271 and –.197. Under the condition of strong moral beliefs, however, the original relationship is unchanged. What we would conclude from this elaboration analysis is that the variable "moral beliefs" is not acting as a very strong extraneous variable. Although some of the relationship between perceived risk and delinquency is due to their joint relationship with moral beliefs, we cannot dismiss the possibility that the perceived certainty of punishment has a causal influence on delinquent behavior.

Specification

By adding a third variable to an evaluation of a bivariate relationship, the data analyst can also specify the conditions under which the bivariate relationship occurs. A **specification** occurs when the association between the independent and dependent variables varies across the categories of one or more other control variables—that is, when the original relationship is stronger under some condition or conditions of a third variable and weaker under others.

Specification A type of relationship involving three or more variables in which the association between the independent and dependent variables varies across the categories of one or more other control variables.

In criminology, social learning theory would predict that youths who are exposed to peers who provide verbal support for delinquency are at greater risk for their own delinquent conduct. We found support for this hypothesis in our delinquency dataset. We examined this relationship by recoding into two approximately equal groups the variable FROPINON (see Exhibit F.2). The first group had weak verbal support from peers, whereas the second group had strong verbal support. Among those youths who reported that their peers provided only weak verbal support for delinquency, 15% were highly delinquent. Among those with strong verbal support from peers, nearly 58% were

Exhibit F.20 Cross-Tabulation of Perceived Risk by Delinquency Within Levels of Moral Beliefs

Weak Moral Beliefs

		Self-Reported Delinquency			
		Low	*Medium*	*High*	*Total*
Perceived Certainty	Low	3.8%	14.7%	81.4%	156
	Medium	8.6%	27.3%	64.1%	128
	High	4.2%	29.6%	66.8%	71
	Total				355

$\chi^2 = 13.646$ ($p < .001$), Gamma = –.271

Medium Moral Beliefs

		Self-Reported Delinquency			
		Low	*Medium*	*High*	*Total*
Perceived Certainty	Low	22.0%	42.5%	35.4%	127
	Medium	33.9%	35.6%	30.5%	174
	High	41.1%	34.2%	24.8%	202
	Total				503

$\chi^2 = 13.646$ ($p < .001$), Gamma = –.197

Strong Moral Beliefs

		Self-Reported Delinquency			
		Low	*Medium*	*High*	*Total*
Perceived Certainty	Low	42.7%	28.1%	29.2%	89
	Medium	58.9%	17.7%	23.4%	107
	High	72.9%	18.3%	8.7%	218
	Total				414

$\chi^2 = 13.646$ ($p < .001$), Gamma = –.393

highly delinquent. The gamma value for this relationship was .711, a very strong positive relationship. Clearly, then, having friends give you verbal support for delinquent acts (e.g., "it's okay to steal") puts you at risk for delinquency.

It is entirely possible, however, that this relationship exists only when friends' verbal support is backed up by their own behavior. That is, verbal support from our peers might not affect our delinquency when they do not themselves commit delinquent acts or when they commit only a very few. In this case, their actions (inaction in this case) speak louder than their words, and their verbal support does not influence us. When they also commit delinquent acts, however, the verbal support of peers carries great weight.

We looked at this possibility to examine the relationship between friends' verbal support for delinquency and a youth's own delinquency within two levels of friends' behavior (FRBEHAVE; see Exhibit F.2). We recoded FRBEHAVE into two approximately equal groups. In the first group, fewer of one's friends are delinquent (few delinquent friends) than the other (many delinquent friends). This attempt to specify the relationship between friends' opinions and a youth's own delinquency is shown in Exhibit F.21. What we see is a little complex. When only a few of a youth's friends are committing delinquent acts, their verbal support still has a significant and positive effect on self-reported

Exhibit F.21 Cross-Tabulation of Friends' Verbal Support by Delinquency Within Levels of Friends' Delinquent Behavior

Few Delinquent Friends

		Self-Reported Delinquency			
Friends' Verbal Support		*Low*	*Medium*	*High*	*Total*
	Weak	67.3%	23.8%	8.9%	437
	Strong	44.3%	34.3%	21.4%	140
	Total				577

$\chi^2 = 27.374$ ($p > .001$), Gamma =.416

Many Delinquent Friends

		Self-Reported Delinquency			
Friends' Verbal Support		*Low*	*Medium*	*High*	*Total*
	Weak	30.5%	38.5%	31.0%	174
	Strong	7.5%	24.8%	67.4%	521
	Total				695

$\chi^2 = 87.508$ ($p > .001$), Gamma =.608

delinquency. The gamma value in this subtable is .416, which is moderately strong but less than the original gamma of .771. When many of a youth's friends are delinquent, however, the positive relationship between peers' verbal support and self-reported delinquency is much stronger, with a gamma of .608. The behavior of our peers, then, only weakly specifies the relationship between peer opinion and delinquency. Clearly, then, what our peers say about delinquency matters, even if they are not committing delinquent acts all the time themselves.

2 Regression and Correlation

Our goal in introducing you to cross-tabulation has been to help you think about the associations among variables and to give you a relatively easy tool for describing association. To read most statistical reports and to conduct more sophisticated analyses of social data, you will have to extend your statistical knowledge. Many statistical reports and articles published in the social sciences use statistical techniques called **regression analysis** and **correlation analysis** to describe the associations among two or more quantitative variables. The terms actually refer to different aspects of the same technique. Statistics based on regression and correlation are used frequently in social science and have many advantages over cross-tabulation—as well as some disadvantages.

We provide only a brief overview of this approach here. Take a look at Exhibit F.22. It's a plot, termed a *scatterplot,* of the bivariate relationship

Regression analysis A statistical technique for characterizing the pattern of a relationship between two quantitative variables in terms of a linear equation and for summarizing the strength of this relationship.

Correlation analysis A standardized statistical technique that summarizes the strength of a relationship between two quantitative variables in terms of its adherence to a linear pattern.

Exhibit F.22 **Example of a Positive Relationship. Scatterplot of Murder Rate (dependent variable) and Poverty Rate (independent variable) in U.S. States, 2010.**

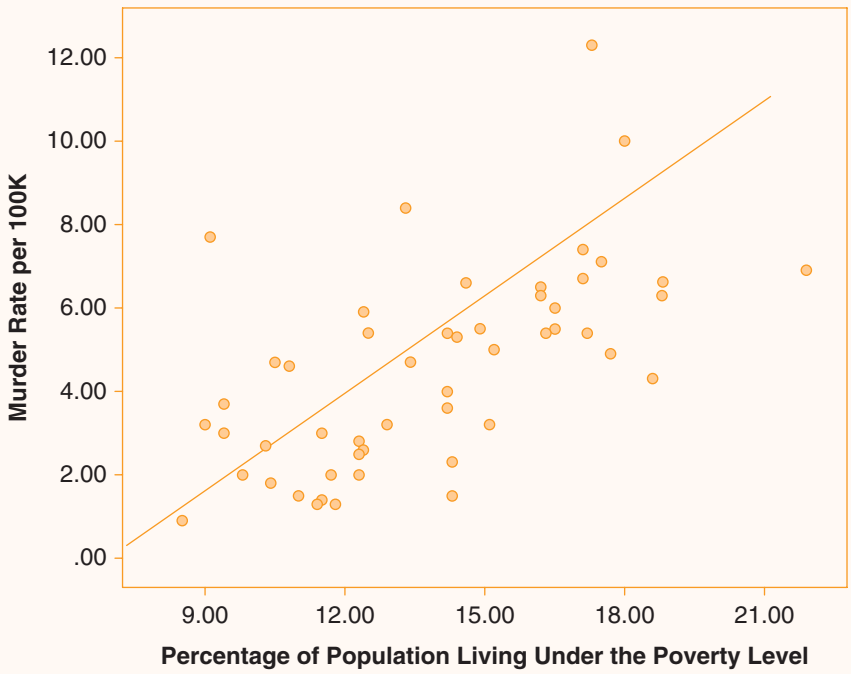

between two interval/ratio-level variables. The variables were obtained from a U.S. state-level dataset. The dependent variable, presented on the *y*-axis (vertical) is the murder rate per 100,000 population, and the independent variable, presented on the *x*-axis (horizontal), is the poverty rate (percentage of each state's population living under the poverty level).

You can see that the data points in the scatterplot tend to run from the lower left to the upper right of the chart, indicating a positive relationship. States with higher levels of poverty also tend to have higher rates of murder. This regression line is the "best fitting" straight line for this relationship—it is the line that lies closest to all the points in the chart, according to certain criteria. But you can easily see that quite a few points are pretty far from the regression line.

How well does the regression line fit the points? In other words, how close does the regression line come to the points? (Actually, it's the square of the vertical distance, on the *y*-axis, between the points and the regression line that is used as the criterion.) The **correlation coefficient**, also called *Pearson's r,* or just *r,* gives one answer to that question. The value of *r* for this relationship is .60, which indicates a moderately strong positive linear relationship (if it were a negative relationship, *r* would have a negative sign). The value of *r* is 0 when there is absolutely no linear relationship between the two variables, and it is 1 when all the points representing all the cases lie exactly on the regression line (which would mean that the regression line describes the relationship perfectly).

Correlation coefficient (r) A summary statistic that varies from 0 to 1 or –1, with 0 indicating the absence of a linear relationship between two quantitative variables and 1 or –1 indicating that the relationship is completely described by the line representing the regression of the dependent variable on the independent variable.

So the correlation coefficient does for two interval/ratio-level variables what gamma does for a cross-tabulation table: It is a summary statistic that tells us about the strength of the association between the two variables. Values of *r* close to 0 indicate that the relationship is weak; values of *r* close to ±1 indicate the relationship is strong—in between there is a lot of room for judgment. You will learn in a statistics course that r^2 is often used

Exhibit F.23 A Guide to Interpreting Strong to Weak Relationships

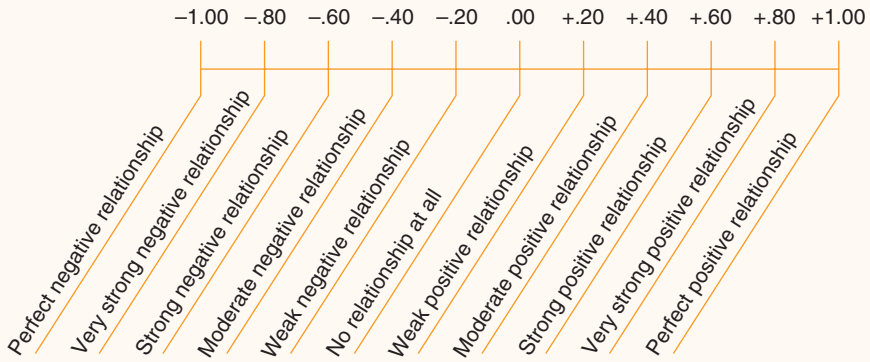

Source: Frankfort-Nachmais and Leon-Guerrero (2006:230). Reprinted with permission from SAGE Publications, Inc.

instead of *r*. Exhibit F.23 provides an overview of how to interpret the values of *r*. Although not all possible values of *r* are displayed in Exhibit F.23, it highlights how the use of adjectives can describe various values between 0 and 1.

An example of a negative relationship is shown in Exhibit F.24, where we provide a scatterplot of the robbery rate in states (dependent variable) on the *y*-axis and the percentage of each state's population that resides in rural

Exhibit F.24 Example of a Negative Relationship. Scatterplot of Robbery Rate (dependent variable) and Percentage Rural (independent variable) in U.S. States, 2010.

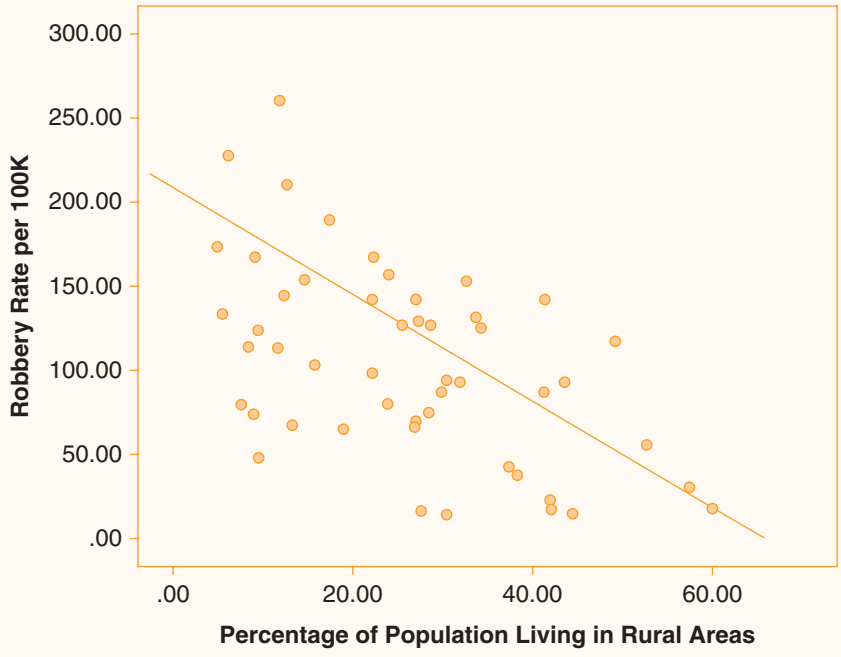

areas as the independent variable (*x*-axis). You can see here a clear negative relationship; a state that has a higher percentage of its population residing in rural areas will tend to have lower robbery rates. The correlation coefficient for this relationship is $r = -.53$, indicating a moderate negative relationship.

You can also use correlation coefficients and regression analysis to study simultaneously the association among three or more variables. Let's use the murder rate as the dependent variable to illustrate. In a *multiple regression analysis,* you could test to see whether several other variables in addition to poverty are associated simultaneously with the murder rate—that is, whether the variables have independent effects on murder after statistically controlling for each other.

Controlling for the geography in a state is also important for predicting murder rates, so we will be including percentage rural in our equation. We also know that robberies sometimes have lethal outcomes, so controlling for the robbery rate is also important. Let's examine what a multiple regression equation would look like predicting the murder rate using the poverty rate, the percentage rural, and the robbery rate as the three independent variables. Interpreting regression output is way beyond the scope of this text; we are simply going to examine the standardized regression coefficients, called *betas,* and their significance level for this illustration. Results are displayed in Exhibit F.25.

First, look at the numbers under the Beta Coefficient heading. Beta coefficients are standardized statistics that indicate how strong the linear relationship is between the dependent variable (murder rate, in this case) and each independent variable, while the other independent variables are controlled. Like the correlation coefficient (*r*), values of beta range from 0, when there is no linear association, to ±1.0, when the association falls exactly on a straight line. You can see in the beta column that rural population is not significantly related to the murder rate when the other variables are controlled. Both the percentage poor and the robbery rate, however, are still significant predictors of murder. R^2 (*r*-squared) is a model fit statistic and tells us, when multiplied by 100, the percentage of the dependent variable's variation that is explained by all the independent variables in the model. In this model, we learn from R^2 that the three independent variables together explain, or account for, 68% of the total variation in murder rates. Our goal is to explain as much variation as possible of the 100%, so explaining over two-thirds of the variation is not bad!

You will need to learn more about when correlation coefficients and regression analysis are appropriate (e.g., both variables have to be quantitative, and the relationship has to be linear [not curvilinear]), but that's for another time and place. To learn more about correlation coefficients and regression analysis, you should take an entire statistics course. For now, this short introduction will enable you to make sense of more of the statistical analyses you find in research articles. You can also learn more about these techniques with the tutorials on the text's study site.

Exhibit F.25 **Multiple Regression Predicting the Murder Rate in States Using Poverty, the Divorce Rate, and the Robbery Rate as Independent Variables**

Variable	*Beta Coefficient*	*Significance Level*
Percentage poor	.462	$p = .001$
Percentage rural	–.001	$p = .382$
Robbery rate	.630	$p = .001$
R^2	.68	
N	50	

Source: Dana Hunt

CAREERS AND RESEARCH

Dana Hunt, PhD, Principal Scientist

In the study site video for this chapter, Dana Hunt discusses two of the many lessons she has learned about measurement in a decades-long career in social research. Hunt received her BA in sociology from Hood College in Pennsylvania and then earned her PhD in sociology at the University of Pennsylvania. After teaching at Hood for several years, she took an applied research position at National Development and Research Institutes (NDRI) in New York City. NDRI's description on its website gives you an idea of what drew the attention of a talented young social scientist.

Founded in 1967, NDRI is a nonprofit research and educational organization dedicated to advancing scientific knowledge in the areas of drug and alcohol abuse, treatment, and recovery; HIV, AIDS, and HCV (hepatitis C virus); therapeutic communities; youth at risk; and related areas of public health, mental health, criminal justice, urban problems, prevention, and epidemiology.

Hunt moved from New York to the Boston area in 1990, where she is now a principal scientist at Abt Associates, Inc., in Cambridge. Abt's website description conveys the scope of the research projects the company directs.

Abt Associates applies scientific research, consulting, and technical assistance expertise on a wide range of issues in social, economic, and health policy; international development; clinical trials; and registries. One of the largest for-profit government and business research and consulting firms in the world, Abt Associates delivers practical, measurable, high-value-added results.

Two of Hunt's major research projects in recent years are the nationwide Arrestee Drug Abuse Monitoring Program for the Office of National Drug Control Policy and a study of prostitution and sex trafficking demand reduction for the National Institute of Justice.

2 Analyzing Data Ethically: How Not to Lie About Relationships

When the data analyst begins to examine relationships among variables in some real data, social science research becomes most exciting. The moment of truth, it would seem, has arrived. Either the hypotheses are supported or not. But, in fact, this is also a time to proceed with caution and to evaluate the analyses of others with even more caution. Once large datasets are entered into a computer, it becomes very easy to check out a great many relationships; when relationships are examined among three or more variables at a time, the possibilities become almost endless.

This range of possibilities presents a great hazard for data analysis. It becomes very temping to search around in the data until something interesting emerges. Rejected hypotheses are forgotten in favor of highlighting what's going on in the data. It is not wrong to examine data for unanticipated relationships; the problem is that inevitably some relationships between variables will appear just on the basis of chance association alone. If you search hard and long enough, it will be possible to come up with something that really means nothing.

A reasonable balance must be struck between deductive data analysis to test hypotheses and inductive analysis to explore patterns in a dataset. Hypotheses formulated in advance of data collection must be tested as they were originally stated; any further analyses of these hypotheses that involve a more exploratory strategy must be labeled as such in research reports. Serendipitous findings do not need to be ignored, but it must be

reported that they were serendipitous. Subsequent researchers can try to deductively test the ideas generated by our explorations.

We also have to be honest about the limitations of using survey data to test causal hypotheses. The usual practice for those who seek to test a causal hypothesis with nonexperimental survey data is to test for the relationship between the independent and dependent variables, controlling for other variables that might possibly create spurious relationships. This is what we did by examining the relationship between the perceived certainty of punishment and delinquency while controlling for moral beliefs. But finding that a hypothesized relationship is not altered by controlling for just one variable does not establish that the relationship is causal, nor does controlling for two, three, or many more variables. There always is a possibility that some other variable that we did not think to control, or that was not even measured in the survey, has produced a spurious relationship between the independent and dependent variables in our hypothesis (Lieberson, 1985). We must always think about the possibilities and be cautious in our causal conclusions.

12 Conclusion

This chapter has demonstrated how a researcher can describe phenomena in criminal justice and criminology, identify relationships among them, explore the reasons for these relationships, and test hypotheses about them. Statistics provide a remarkably useful tool for developing our understanding of the social world, a tool that we can use to test our ideas and generate new ones.

Unfortunately, to the uninitiated, the use of statistics can seem to end debate right there; you cannot argue with the numbers. But you now know better than that. The numbers will be worthless if the methods used to generate the data are not valid, and the numbers will be misleading if they are not used appropriately, taking into account the type of data to which they are applied. And even assuming valid methods and proper use of statistics, there is one more critical step, for the numbers do not speak for themselves. Ultimately, it is how we interpret and report the numbers that determines their usefulness. It is this topic we turn to in the next chapter.

Key Terms

Bar chart F-6
Base N F-8
Bimodal distribution F-15
Central tendency F-5
Chi-square F-28
Correlation analysis F-33
Correlation coefficient (r) F-34
Cross-tabulation (cross-tab) F-2
Data cleaning F-5
Descriptive statistics F-2
Elaboration analysis F-29
Extraneous variable F-30
Frequency distributions F-2
Gamma F-27
Grouped frequency distribution F-10
Histogram F-6
Inferential statistics F-2
Interquartile range F-22
Marginal distributions F-24
Mean F-17
Measure of association F-27
Median F-15
Mode F-6
Monotonic relationship F-27
Negatively skewed F-7
Outlier F-21
Percentage F-6
Positively skewed F-7
Quartiles F-22
Range F-21
Regression analysis F-33
Skewness F-5
Specification F-31
Standard deviation F-23
Statistical significance F-28
Subtables F-29
Unimodal distribution F-15
Variability F-5
Variance F-22

Highlights

- Data collection instruments should be precoded for direct entry, after verification, into a computer. All data should be cleaned during the data entry process.
- Use of secondary data can save considerable time and resources but may limit data analysis possibilities.
- Bar charts, histograms, and frequency polygons are useful for describing the shape of distributions. Care must be taken with graphic displays to avoid distorting a distribution's apparent shape.
- Frequency distributions display variation in a form that can be easily inspected and described. Values should be grouped in frequency distributions in a way that does not alter the shape of the distribution. Following several guidelines can reduce the risk of problems.
- Summary statistics are often used to describe the central tendency and variability of distributions. The appropriateness of using the mode, mean, and median for a description varies with a variable's level of measurement, the distribution's shape, and the purpose of the summary.
- The variance and standard deviation summarize variability around the mean. The interquartile range is usually preferable to the range to indicate the interval spanned by cases, due to the effect of outliers on the range. The degree of skewness of a distribution is usually described in words rather than with a summary statistic.
- Cross-tabulations should normally be divided into percentages within the categories of the independent variable. A cross-tabulation can be used to determine the existence, strength, direction, and pattern of an association.
- Elaboration analysis can be used in cross-tabular analysis to test for spurious and intervening relationships and to identify the conditions under which relationships occur.
- Inferential statistics are used with sample-based data to estimate the confidence that can be placed in a statistical estimate of a population parameter. Estimates of the probability that an association between variables may have occurred on the basis of chance are also based on inferential statistics.
- Regression analysis is a statistical method for characterizing the relationship between two or more quantitative variables with a linear equation and for summarizing the extent to which the linear equation represents that relationship. Correlation coefficients summarize the fit of the relationship to the regression line.

Exercises

1. Create frequency distributions from lists in the Federal Bureau of Investigation (FBI) Uniform Crime Reports on characteristics of arrestees in at least 100 cases (cites). You will have to decide on grouping schemes for the distribution of data for variables such as race, age, and crime committed, and how to deal with outliers in the frequency distribution.

   a. Decide what summary statistics to use for each variable of interest. How well were the features of each distribution represented by the summary statistics? Describe the shape of each distribution.

   b. Propose a hypothesis involving two of these variables, and develop a cross-tabulation to evaluate the support for this hypothesis.

   c. Describe each relationship in terms of the four aspects of an association, after making percentages within each table within the categories of the independent variable. Which hypotheses appear to have been supported?

2. Become a media critic. For the next week, scan a newspaper or some magazines for statistics related to crime or criminal victimization. How many can you find using frequency distributions, graphs, and the summary statistics introduced in this chapter? Are these statistics used appropriately and interpreted correctly? Would any other statistics have been preferable or useful in addition to those presented?

3. The table that follows shows a frequency distribution of "trust in people" as produced by SPSS with the General Social Survey data. As you can see, the table includes abbreviated labels for the variable and its response choices, as well as the raw frequencies and three percentage columns. The first percentage column (Percentage) shows the percentage in each category; the next percentage column (Valid Percentage) is based on the total number of respondents who gave valid answers (3,929 in this instance). It is the Valid Percentage column that normally should be used to construct a frequency distribution for presentation. The last percentage column is Cumulative Percentage, adding up the valid percentages from top to bottom.

   Redo the table for presentation, using the format of the frequency distributions presented in the text.

		Frequency	*Percentage*	*Valid Percentage*	*Cumulative Percentage*
Valid	CAN TRUST	1279	28.4		
	CANNOT TRUST	2458	54.5		
	DEPENDS	192	4.3		
	Total	3929	87.1		
		Frequency	*Percentage*		
Missing	NAP	575	12.7		
	NA	6	.1		
	Total	581	12.9		
Total		4510	100.0		

Developing a Research Proposal

Use the General Social Survey data to add a pilot study to your proposal. A pilot study is a preliminary effort to test out the procedures and concepts that you have proposed to research.

1. Review the GSSCRJ2K variable list, and identify some variables that have at least some connection to your research problem. If possible, identify one variable that might be treated as independent in your proposed research and one that might be treated as dependent.
2. Request frequencies for these variables.
3. Request a cross-tabulation of the dependent variable by the independent variable (if you were able to identify any). If necessary, recode the independent variable to three or fewer categories.
4. Write a brief description of your findings and comment on their implications for your proposed research. Did you learn any lessons from this exercise for your proposal?

Web Exercises

1. Search the web for a crime-related example of statistics. The Bureau of Justice Statistics is a good place to start: www.ojp.usdoj.gov/bjs/. Using the key terms from this chapter, describe the set of statistics you have identified. What phenomena does this set of statistics describe? What relationships, if any, do the statistics identify?
2. Do a web search for information on a criminological subject that interests you. How much of the information that you find relies on statistics as a tool for understanding the subject? How do statistics allow researchers to test their ideas about the subject and generate new ideas? Write your findings in a brief report, referring to the websites that you found.

Ethics Exercises

1. Review the frequency distributions and graphs in this chapter. Change one of these data displays so that you are "lying with statistics."
2. Consider the relationship between gender and delinquency that is presented in Exhibit F.13. What third variable do you think should be controlled in the analysis to better understand the basis for this relationship? How might criminal justice policies be affected by finding out that this relationship was due to differences in teacher expectations rather than to genetic differences in violence propensity?

SPSS or Excel Exercises

Data for Exercise	
Dataset	**Description**
Youth.sav	This dataset is from a random sample of students from schools in a southern state. While not representative of the United States, it covers a variety of important delinquent behaviors and peer influences.
Variables for Exercise	
Variable Name	**Description**
delinquency	An interval/level variable that measures self-reported delinquency.
D1	A binary variable based on the number of delinquent acts a respondent reported. A 0 indicates that the respondent reported 1 or fewer acts, while 1 indicates 2 or more.
Variables for Exercise	
Variable Name	**Description**
lowcertain_bin	Binary indicator for whether the respondent felt there was certainty that he or she would be punished for delinquent behaviors, where 1 = low certainty and 0 = high certainty.
certain	A scale indicating how likely the individual feels it is that he or she will be punished for delinquent behavior. High values indicate high certainty.

1. For this exercise let's take a look at whether a person's expectation of punishment after Delinquency is associated with the number of deviant behaviors a student engages in, as measured by the variable Delinquency.
   a. Run a frequency of the dependent variable, delinquency, and answer the following questions:
      i. What level of measurement is this item?
      ii. What forms of descriptive analysis are appropriate?
      iii. How would you best represent this data in a graph?
   b. Based on your responses to Part 1a, conduct all appropriate descriptive analyses. Be sure to describe what you can about the data's distribution and what measures of central tendency are most appropriate. If one or another measure may produce misleading results, be sure to caution the reader why.
2. D1 measures delinquency differently than the interval/ratio level variable called Delinquency. Is the variable D1 appropriate for use in an ordinary least squares (OLS) regression analysis? Why or why not? If you have been

taught them, consider how it will or will not meet different assumptions of OLS.

3. Repeat Part 1 for the independent variable, which is called lowcertain_bin. Again, describe the variable and how you would go about presenting it. Remember that you are required to conduct only the analyses that are appropriate.
4. On to the actual analysis! First, let's compare mean delinquency scores for lowcertain_bin and the delinquency variable. This can be done under analyze->means->compare means.
    a. What is the difference between the two group means?
    b. What do these results suggest substantively?
5. Second, let's estimate a linear regression model. This can be accessed by selecting analyze->regression->linear.
    a. What are your results? How do they compare with your results in Part 4?
    b. Do you notice any similarities between your regression coefficient and the results from Part 4? Think carefully about why this is the case—would this similarity apply to all independent variables in a regression model or just binary ones?
    c. Test your answer to Part 5b by running the regression model again, but this time use the continuous version of lowcertain_bin, named "certain." High values on this measure indicate high levels of certainty, which is the inverse of the original measure.
        i. How have your results changed?
        ii. Do these results lead to substantively similar conclusions?
6. Return to your answer for Part 2. How sound do you think these specific analyses are, given that they are all based on the analysis of means? If you think they may be biased, explain how they are biased and any ideas you might have for overcoming them.

STUDENT STUDY SITE

The companion Student Study Site for *Fundamentals of Research in Criminology and Criminal Justice* can be found at **https://study.sagepub.com/bachmanfrccj4e**.

Visit the Student Study Site to enhance your understanding of the chapter content and to discover additional resources that will take your learning one step further. You can enhance your understanding of the chapters by using the comprehensive study material, which includes SAGE journal and reference articles, e-flashcards, quizzes, multimedia links, and more.