

These comments are designed to try to clear up some of the confusion about case studies as a **research design**. Your text does a good job of treating the case study, but by virtue of what de Vaus **does not** discuss, he leaves some false impressions about the case study. The literature is even worse because the term is used to describe both a research design and a method of data collection and because different disciplines use the term to mean different things. I highly recommend that you read the Yin reading cited at the web site.

The Nature of the Case Study

From a design perspective, the distinguishing feature of the case study is that the sample is selected based on the **outcome or dependent variable(s)**. We decide which cases to study because we already know “how it turned out.” This is the exact opposite of the experimental designs where we select subjects because they meet criteria that we have defined *a priori*, randomly assign them to treatment and control groups, and then “do something” to the treatment group. E.g., in the true experiment we start with a group of test subjects who were “nearly identical” in terms of the **predictor or independent** variable(s) and then intervene with the treatment group to create a **difference between subjects** in terms of the outcome variable.. In the case study, we select subjects who **are already different in terms of the outcome variable** and then try to determine why they are different. We know the end state. Our task as researchers is to determine how the end state came about.

Smoking and Lung Cancer

Almost all of the **original** research that identified the relationship between smoking and lung cancer was case study research. I refer here to the research that originally **identified** the link, not later research that clarified and helped confirm the link. Think about this carefully and you will see that you can eliminate all of the other designs. The fundamental question was “Why do people get lung cancer?”

Not an experiment. We did not select a group of subjects, randomly assign them to treatment and control groups, and then have the treatment group smoke 3 packs of cigarettes per day for 20 years to see if they would get lung cancer. We could not have done this anyway, even if it were ethical and possible, because we did not know that smoking was a risk factor for lung cancer.

Not a quasi-experiment. This is harder to see, but remember that quasi- or natural experiments require the occurrence of some discrete event that creates a difference between two groups of people who are otherwise “nearly identical,” just as in the true experiment. There are two problems here. Smoking is a life-style choice that people practice for a long period of time. It is not a discrete event. Also, we did not know to select treatment and control groups on the basis of smoking.

Not longitudinal. This may be rather difficult to see, but remember that longitudinal studies require repeated measurement of the same variables **with the same subjects** over time. It certainly was not a prospective longitudinal study. We did not select a large, representative sample of people at say age 10 or 12, follow them for 20 years asking every year “Do you smoke now? If so, how many packs per day?”, and then see if the smokers ended up with lung cancer. Retrospective longitudinal

studies still have the requirement of repeated measurement over time. We simply had no way to make these repeated retrospective measurements of all the independent variables that might cause lung cancer with a group of smokers and non-smokers. Once again, we would not necessarily have known to ask about smoking, although certainly smoking was “suspected” as a cause of lung cancer.

Not cross-sectional. I suppose we could have used cross-sectional studies in the original research about the causes of lung cancer, but it would have been very difficult and time-consuming. You would have to have a large, representative sample of the population and then try to ask about every potential cause of lung cancer and find the relationship between smoking and lung cancer. Remember, I’m talking here about the **original identification of smoking as a factor in lung cancer**. This means that you would not even have known to take independent samples of smokers and non-smokers, or would have had no justification for this basis for sampling. Furthermore, remember that establishing cause and effect requires not only that you show a positive relationship between two factors like smoking and lung cancer, but also that you show that you can eliminate alternative explanations. This would have been very hard to do with a cross-sectional design because lung cancer is relatively rare. By taking a statistically representative sample of the population, you would have ended up with a lot of people who do not have lung cancer and only a few who do. It’s just not efficient. You might also have seen a statistical relationship between say eating pizza and getting lung cancer. This is the classic weakness of the cross-sectional design. It just is not a very powerful design for showing cause and effect.

So what did the researchers do? They took a large sample of people with severe lung cancer -- e.g., they selected the test or study subjects based on the **dependent variable or outcome state**, people lying in hospital beds croaking of lung cancer. Then they **eliminated** those subjects who had other known risk factors for cancer like coal miners or people who had suffered some form of radiation exposure or whatever. I do not know what the known risk factors were. The objective was to end up with a group of subjects who had lung cancer and for whom the existing explanations of its cause (like radiation) were eliminated. Then, the researchers tried to find out what this sample of subjects had in common among suspected risk factors. Smoking turned out to be the commonality. I do not know what the other suspected risk factors were. Does this “prove” that smoking causes lung cancer? No. There could still be some unidentified commonality among these people that we failed to identify. However, the internal validity is pretty high **because** of the case study design. We picked people who had the condition. We eliminated people who could have acquired lung cancer from known risk factors -- e.g., we eliminated as many alternative explanations for lung cancer as we possibly could. Then we examined the remaining subjects in terms of every suspected risk factor, including smoking.

Of course, the people who manufacture cigarettes got pretty upset. The classic statement was “Show me the person who died from smoking.” We still have not been able to do this. If you think about this from a research design perspective, the people who made this statement were basically arguing that there is only one research design, the true experiment, that establishes cause and effect. I would hope that you know by now that this is not true, that internal validity is always a matter of degree, and that all designs, including true experiments, have some threats to internal validity. This is why we can never absolutely eliminate all alternative explanations. It’s just not possible, no matter what design you use. **Well designed** case studies actually tend to be fairly strong in terms of internal validity.

So how do we “know” that smoking causes lung cancer, something even cigarette manufacturers now admit. The basic answer is yet more case study research. After showing the strong commonality of smoking among the subjects with lung cancer, the researchers were still left with some people who did not smoke, but suffered from lung cancer. Over the years, they continued to trace common-

alities among these people. They identified yet more commonalities -- like working in textile factories (inhaling those little fiber particles). They then removed these subjects from the sample, which produced an even stronger relationship between smoking and lung cancer. Then maybe it was exposure to asbestos. Again, with the remaining subjects showed an even stronger relationship between smoking and lung cancer. Of course, we still have people who get lung cancer and never smoked, and people who smoke and do not get lung cancer. Put in research design terms, we cannot establish a direct causal relationship between smoking and lung cancer. In fact, we now know that the **direct** cause of lung cancer is the introduction of carcinogens into the lung tissue. Smoking just happens to be a very efficient, effective way of introducing the carcinogens. Working in a mine or around asbestos are good, too. E.g., we have been able to identify the specific mechanisms that **explain** the relationship between smoking and lung cancer, the presence of carcinogens in cigarette smoke. So, today we are (except for a few holdouts) convinced that smoking does indeed give you lung cancer, that smoking kills.

Fallacies about Case Studies

The sample is always small. Some case studies include thousands of subjects in the sample. Some case studies use a sample as small as one. There are no rules about the size of the sample.

The sample is always judgmental or purposive rather than statistically representative. You can use any kind of legitimate sampling procedure for a case study. The appropriate procedure depends on the research question, the time and money you have for the study, etc., all the factors that we have discussed with all the other designs. What is important is that you select your subjects based on the outcome state. If, for example, you want to study the effects of pre-marital counseling on long term marital stability, you would want to select a sample of people who have been married for a long time, however you define that. As is true for the experiment, the importance of the statistically representative sample is less in case studies than in longitudinal or cross-sectional designs because the case study depends more on theoretical generalization and less on statistical generalization than do those designs, and because case studies allow you to use screening or criteria matching to select study subjects.

If I pick a single place to conduct my research I automatically have a case study. Everyone picks a place to conduct research, even physicists. The question is, do you think the place you picked will somehow influence the results? For example, if you are a physicist, do you think it will matter if you complete an experiment in Phoenix or Gainesville? I doubt it if you are a physicist who studies subatomic particles. I can't think of a single reason why subatomic particles would somehow "act differently" in Gainesville than in Phoenix. On the other hand, if you're an astrophysicist who uses telescopes to collect data about stars, it would matter a lot. Telescope viewing in the humid subtropics just isn't prime time. If you report 10 globular clusters in the Milky Way and a colleague in Phoenix reports 20, I'd be very willing to bet he's right. All that water in the atmosphere in our part of the world gets in the way of seeing stars through a telescope. It will definitely affect something as simple as a star count. So don't set up your telescope in Gainesville. This is really the question of whether the accessible population is a reasonable representation of the theoretical population and it is an important question. Let's say you want to understand people's attitude in the U.S. about the relationship between educational achievement and incarceration. "People in the U.S.," is the theoretical population. Gainesville is probably a bad choice. It's a highly educated community and probably is not at all "like" the population of the U.S. as a whole in terms of attitudes about the importance of education in general. Starke's probably a bad choice, too, because of the prison there. It's presence may color people's attitudes about incarceration and incarcerated people in general. Go somewhere else. Ocala, or Tallahassee, or Orlando are probably fine choices. You'd want to check out some things and see if they are more or less like the US as a whole in terms of educational level at least, but there's no *a priori* reason to think they are bad choices.

Case studies rely on “qualitative” or nominal data. You can collect any kind of data you want with a case study. Some research that you probably think of as very “quantitative” actually depends on case studies **and** on nominal data. For example, a lot of medical research depends on case studies that rely heavily on symptomology, and much of the latter is nominal data, things like the doctor’s observations, the patient’s observations or reported symptoms (lower back pain, frequent headaches, etc.).

Case studies do not use a control. Some do and some do not. In general, we know that including a control group enhances internal validity and is a strong mechanism for establishing causality. Many case study designs do include a control group. The Rodriguez, Perez & Garcia (2005) article, for example, describes research in which 160 Colombian women between 12 and 49 years of age with eating disorders were selected for study. The researchers wanted to know whether traumatic experiences and violent acts play a role in reducing the effectiveness of treatment for the eating disorder. Response to treatment for eating disorder is the outcome state. They divided the 160 women into two groups -- 70 with “poor response” and 90 with “better response.” Then they examined the relationship between the outcome variable, the response to treatment, and the exposure to traumatic experiences and violent acts at some time in the women’s lives (nominal data, you will notice, based on interviews and life histories of the subjects). Exposure was important. E.g., women who have experienced this kind of trauma and violence in their lives (not necessarily during the treatment for the eating disorder), are much more likely to drop out of treatment or relapse into the eating disorder after treatment. In other cases, the control is selected much like we select control groups in quasi-experiments. You select the “treatment” subjects based on the outcome state, and then you match them with other “control” subjects who do not have the outcome state on criteria that you want to **eliminate** as potential causes of the outcome state. The Szerlip, Desalvo and Szerlip (2005) article, for example, reports research about how to assess the risk that older patients will be HIV-positive (so that doctors know when to recommend HIV testing for older patients). The “treatment” group subjects were 53 patients over 55 who were HIV positive. They were matched on age and gender to 106 HIV-negative “control” group subjects. Then they examined the case histories of everyone in both groups to see how they differed. They took both nominal (medical and life history) and interval (things like albumin levels) data. They found that a history of alcohol abuse, a history of sexually transmitted disease, or an albumin to globulin ration of less than 1, were characteristics that distinguished between the two groups.

The case study depends on the interview or the ethnography as the primary method of data collection. Researchers use just about every method of data collection ever devised to collect data in case study designs. Samples of albumin and globulin are examples of interval data.

The case study requires that you examine a large number of variables, that you learn everything you can about the subject. This is really just a version of the “method dependent” argument and it is not true. Some case studies do examine large numbers of variables. Others examine just a few. It **is** true that the case study design is particularly useful when you suspect that a large number of factors interact to produce a given outcome state or result or outcome variable. It’s just plain hard to study a lot of variables in many designs. The experiment is the very limiting in this regard, for example, as we saw earlier in this course. Even four variables presents a big challenge for a researcher who wants to use an experimental design. Think about it. With four variables, you have four treatments plus a control, and you need all the possible combinations of the four treatments. So let’s count up the number of groups you have to establish: (1) the control, (2) factor 1 by itself, (3) factor 2 by itself, (4) factor 3 by itself, (5) factor 4 by itself, (6) factor 1 + factor 2, (7) factor 1 + factor 3, (8) factor 1 + factor 4, (9) factor 2 + factor 3, (10) factor 2 + factor 4, (11) factor 3 + factor 4, (12) factor 1 + factor 2 + factor 3, (13) factor 1 + factor 2 + factor 4, (14) factor 1 + factor 3 + factor 4, and (15) factor 2 + factor 3 + factor 4. Wow! I don’t know about you, but I don’t want to even try that. The *post hoc* sample selection of the case study lets you examine all these, plus

many more, factors without establishing all those groups *a priori*. So it is true that you **can** examine many more factors in a case study design than an experiment. Imagine finding all those groups “naturally” occurring for a quasi-experiment. This is highly unlikely. About the only alternative besides the case study for examining the relationships between a large number of variables is the cross-sectional design and, as we know, that design has a lot of problems with spurious relationships and real problems with establishing causality because of the lack of both a control and a temporal component. The best cross-sectional design would require taking independent samples for each group. You’re stuck with trying to take 15 independent samples, or you’re stuck with the *post hoc* approach to establishing groups based on statistical relationships, with all the weaknesses we know that creates. So, case studies are particularly useful when you want to look at the relationships between a lot of variables, **but** you do not have to examine large numbers of variables to have a case study. The Grossman et al. (2005) article provides an example of a very narrow range of variables in a case study. They examined the role of just five factors in reducing the risk of youth suicide and unintentional firearm injuries, very specific things like whether the firearm was stored in an unloaded condition. The Purcell and Nevins (2005) article, on the other hand, examines a very large number of factors that are much more general in nature.

You cannot use statistical analyses of the data in case studies. you can use any kind of analysis with case study data. As always, the degree to which statistical analyses will be useful depends on the kind of data you collected. Statistical analyses in general are a way of helping us decide whether relationships are “real” or not. They are an aid to our interpretation of the data. They do not “go with” any particular kind of design.

Case studies are especially common in the social sciences. Scientists from many disciplines rely heavily on the case study design. Almost all of the historical sciences, from cosmology to paleontology to human history, use this design a lot. By their nature, these disciplines study single events that are not directly observable by the researcher. In essence, these scientists are always trying to understand “what caused the outcome state we observe?” The Pfefferkorn (2004) article is an excellent example of the use of the case study in paleobiology. However, many other disciplines also use the case study. Cabrera-Gomez et al. (2005) describe the use of the case study in medical research, in this case dealing with the hereditary component in disease. By the way, it’s one of those studies with a single case. Papapetropoulos (2005) uses a case study design to examine the causes of dementia in Parkinson’s disease. Even physicists use case studies. The Cockell (2004) article is a good example of the use of the case study in astrophysics.

Case study designs are less rigorous than other designs and are only used in “soft” science. By now, I would hope you understand that there is no such thing as “soft” science. But the perception that somehow case study research is less rigorous than research using other designs always amazes me. To give you just a few of the thousands of articles published in the last year describing what I believe virtually anyone would call very “hard,” rigorous science, take a look at the following. Fairen and Dohm (2004) use a case study to try to figure out how the lowlands of Mars formed. Lamb and Thompson’s (2005) case study focuses on the cause of the decline of the elm in northwestern Europe about 6000 years ago. Gradzinski et al. (2004) use a case study of geologic formations in Poland to try to figure out the ecological relationship between microbes and foraminifera in the early Jurassic. Candy, Black and Sellwood (2004) examine how climate change in the Late Quaternary affect a river system in what is now Spain. MacMillan, Gans and Alvarado (2004) use a case study to figure out the geologic history of southern Costa Rica. These are just a few studies that caught my eye because they are things that interest me. You can literally find thousands of case studies published every year covering every science discipline.

Terminology and Typology

The problem of distinguishing the **case study design** from other designs is simply made worse by the differences in terminology that we have already seen among disciplines. Some disciplines use the term “case study” to refer to any research that is place or event specific or even to studies that use a particular data collection method. For example, anthropologists are prone to calling any study that uses ethnographic data collection methods a “case study.” The whole terminology issue has been further clouded over the past 20 years by adding an epistemological component to the mix. Many post-modernists, for example, who use a different epistemological approach to the development of theory, use the term case study almost synonymously with the epistemology. However, these differences in terminology are in no way limited to the case study. For example, physicists call everything they do (pretty much) an “experiment.” Yet, from a design perspective, much of what they do is not at all an experiment. There is no control. There is no random assignment to treatment and control groups. From a research design perspective, these are the fundamental requirements of an experiment. The terminology is not important in and of itself. What is important is understanding the structure of the research design used for a particular study and then being able to determine the strengths and weaknesses of the design in terms of internal and external validity.

DeVaus makes an attempt at a typology of case study designs. It's an interesting exercise. Certainly, it would be useful to have a typology like we do for other design groups. I'm not sure he's got it “right,” but it's certainly worth considering in your response to Assignment 6.

Analysis of Case Study Data

DeVaus discusses non-statistical analysis techniques under case studies. There is a real problem here for two reasons. First, this implies that you cannot or should not use statistical analyses with case study data. Second, it also implies that non-statistical analytic techniques are not useful with data from other designs. These are both incorrect. My own personal belief is that every analysis should use both statistical and non-statistical analytic techniques to examine data. They complement each other. So, whatever you may take from this course, please do not take the message that case study design = nominal data = non-statistical analysis. That certainly is not the message I want to leave with you.

Balance in Design

In your final project you must examine a body of research from a theoretical perspective and determine the overall balance in design. Every design group has strengths and weaknesses in terms of internal and external validity. Therefore, any body of research as a whole that depends too heavily on a single design approach tends to develop an overall set of weaknesses that can lead us to poor or erroneous conclusions or poor generalizability.

For example, agricultural researchers rely very, very heavily on the use of the true experiment, and even on a single design within that group, the randomized complete block design. The true experiment, as we know, has two major weaknesses from the point of view of external validity or generalizability: explanatory narrowness and artificiality. Given the over-reliance on this design group, we should not be surprised that the research produced in the 20 to 25 years after World War II led to erroneous conclusions. I refer to the unwanted and undesirable impacts of many agricultural practices on the environment. While it may be true that some in agriculture were somehow “anti-environment,” this can hardly explain why entire disciplines failed so dramatically to identify potential “side effects” of the technology that was developed and why, when those side effects were identified, there was so much resistance to accepting them as real and true impacts from agricultural practice and technology. The researchers relied on a design that greatly restricts the number

of variables that can be examined, a key weakness of the design that results in explanatory narrowness. Further, the true experiment extracts the relationships under study from the environment in which they occur in order to be able to manipulate the independent or predictor variables and measure change in one or a very few outcome or dependent variables. It greatly “simplifies” what happens in a natural setting in order to better understand a very limited number of relationships that the researcher identifies. This creates the weakness of artificiality. These two weaknesses created a large body of research that simply failed to examine how specific practices affect non-experimental variables. We knew a lot about how DDT suppresses pests. We knew nothing about how DDT affects non-target organisms like birds, or how this pesticide interacts with the environment to reach the non-target organisms. Further, the researchers who raised these issues, biologists for the most part, use different research designs. In fact, they rarely use experiments. Agricultural researchers, who tended to only understand one design, the experiment, rejected the findings of biologists at least in part, in my opinion, because they simply did not understand either the weaknesses of the experimental design, or the strengths of other design groups.

This is not the only time this has occurred. In fact, the story could be repeated many times. This is why it is important for you to be able to evaluate not just individual pieces of research, but the overall body of research in a specific area. The individual experiments that agricultural researchers performed were fine. They were not flawed. They did not lead to false results. The problem was the over-reliance on one single research design. The same thing can happen when a body of knowledge rests too heavily on any single research design.

We have seen that some disciplines face real difficulties in using some of the design groups. Let’s take geology as an example, and specifically that part of geology that deals with tectonic activity or earthquakes. Scientists in this area have had to rely very heavily on the case study. After all, earthquakes are pretty rare and you certainly can’t create one to conduct an experiment. There’s no real way to set up a quasi-experiment either. Longitudinal designs are pretty much out. A cross-sectional design doesn’t work very well because there are not that many earthquakes and certainly not that many big ones. For many years, geologists pretty much had to use a case study design. They had to observe what happened during or, more commonly, go look for signs of what had happened after an earthquake, and then ask “What caused this earthquake of this magnitude to happen here at this time?” What’s more, geologists have even had to rely on, dare I say it, participant observation, for data; that is, the stories that people tell them about what happened. There was not a single instrument in place when the biggest earthquake in U.S. history occurred, for example, and geologists have used a lot of eye witness reports from regular old folks like you and me to piece together what happened and how strong that earthquake was.

Since we know over-reliance on a single design does not produce a robust body of knowledge, what have geologists done to expand the designs they can use? For one thing, they have put a lot of instrumentation in place in areas where earthquakes are more frequent to collect secondary data. They monitor all sorts of things. This boils down to a longitudinal design. They resample the same places over time and get the same data. It’s not “earthquake” data, but it’s data about places where earthquakes occur. This has helped them understand what kinds of phenomena are associated with earthquakes. Someday, they hope, this will help them actually predict when and where earthquakes will occur. As the body of data about earthquakes has increased, they’ve also been able to use the cross-sectional design to compare data from different places, to sample earthquakes. This has allowed them to look for commonalities among groups of earthquakes. As a result, they’ve been able to develop a typology (remember that basic level of theory development) of earthquakes. They can divide them into groups and hypothesize, based on theory, that different kinds of earthquakes have different causal mechanisms.

Even more recently, they've actually ventured into experiments. They cannot create an earthquake, it's true. However, a lot of what we want to know about earthquakes is how much damage they will cause. Nobody really cares much, after all, if an earthquake occurs in the middle of nowhere. An earthquake in LA is a big problem. So geologists and engineers have started to study risk reduction. They want to know what kinds of structures are at risk in earthquakes and why. To do this, they actually build big machines that can mimic the ground motions that occur during earthquakes. Then they put different kinds of structures on the platform of this machine and shake them. Because of this we now know, for example, that the open faced garage under apartments are a very bad idea. They cause buildings to collapse. In other experiments they load the platform with different kinds of soil (or regolith), like sands and clays, and they can do things like change the moisture content of the soil. By doing this, they have learned that building on landfill in earthquake country is a really bad idea. They combine their machine experiments with models and use the experiments to see if their models "work" to predict what will happen during an earthquake under different conditions.

Overall, then, this group of scientists who study a pretty rare phenomenon and who don't actually get much first-hand experience have done a comendable job of overcoming some inherent limits on research design. They have been creative in figuring out ways to extend the repertoire of designs they can use. As a result, they have been able to enhance the overall internal and external reliability of the body of scientific knowledge that they have created. This is the take home lesson. There is NO single best design. Any body of knowledge that rests too completely on a single design approach is bound to have weaknesses in its internal and external validity. Your job as a researcher or practitioner is to understand the strengths and weaknesses of individual designs or individual studies, and to be able to evaluate the overall validity of the body of knowledge available to you.

References

- Cabrera-Gomez, J.A., Eschazabal-Santana, N., Real-Gonzalez, Y., Garcia, K.R. et al. (2005) *Multiple Sclerosis* 11(3), 364-366.
- Candy, I., Black, S. & Sellwood, B.W. (2004) Interpreting the response of a dryland river system to Late Quaternary climate change. *Quaternary Science Reviews* 23(23-24), 2513-2523.
- Cockell, C.S. (2004) Impact-shocked rocks -- insights into archean and extraterrestrial microbial habitats (and sites for prebiotic chemistry?) *Advances in Space Research* 33(8), 1231-1235.
- Fairen, A.G. & Dohm, J.M. (2004) Age and origin of the lowlands of Mars. *Icarus* 168(2), 277-284.
- Gradzinski, M., Tyszka, J., Uchman, A. & Jach, R. (2004) *Palaeogeography, Palaeoclimatology, Palaeoecology* 213(1-2), 133-151.
- Grossman, D.C., Mueller, B.A., Riedy, C., Dowd, M.D. et al. (2005) Gun storage practices and risk of youth suicide and unintentional firearm injuries. *Journal of the American Medical Association* 293(6), 707-714.
- Lamb, H. & Thompson, A. (2005) Unusual mid-Holocene abundance of *Ulmus* in western Ireland -- human impact in the absence of a pathogen? *The Holocene* 15(3), 447-452.
- MacMillan, I., Gans, P.B. & Alvarado, G. (2004) *Tectonophysics* 392(1-4), 325-348.
- Papapetropoulos, S., Gonzalez, J., Lieberman, A., Villar, J.M. & Mash, D.C. (2005) Dementia in Parkinson's disease: A post-mortem study in a population of brain donors. *International Journal of Geriatric Psychiatry* 20(5), 418-422.

Pfefferkorn, H.W. (2004) The complexity of mass extinction. *Proceedings of the National Academy of Sciences* 101 (35), 12779-12780.

Purcell, M. & Nevins, J. (2005) Pushing the boundary: State restructuring, state theory, and the case of U.S.-Mexico border enforcement in the 1990s. *Political Geography* 24(2), 211-235.

Rodriguez, M., Perez, V., & Garcia, Y. (2005) Impact of traumatic experiences and violent acts upon response to treatment of a sample of Colombian women with eating disorders. *International Journal of Eating Disorders* 37(4), 299-306.

Szerlip, M.A., Desalvo, K.B. & Szerlip, H.M. (2005) Predictors of HIV-infection in older adults. *Journal of Aging and Health* 17(3), 293-304.