

TEACHERS' TOPICS

The Role of Matching in Epidemiologic Studies

Kevin W. Garey, PharmD

College of Pharmacy, University of Houston

Submitted April 20, 2004; accepted July 23, 2004; published August 17, 2004.

The elective graduate course in pharmacoepidemiology described here primarily focuses on epidemiologic methods and their applications in pharmacy. Several different epidemiology study designs are available for pharmacoepidemiology studies, each with their own advantages, disadvantages, and potential for bias. A commonly employed technique in epidemiology is matching. Matching is generally a principle that is not well understood by students; thus, the lecture is given during the second half of the semester, after material on study designs, bias, and confounding has been presented. Strategies to employ matching and common misconceptions are discussed.

Keywords: epidemiology, matching, case-control study

INTRODUCTION

Pharmacoepidemiology, the study of the use and effects of drugs in large numbers of people, is an applied science that draws from both clinical pharmacology, the study of the effects of drugs on humans, and epidemiology, the study of the distribution and determinants of disease in populations. The graduate course in pharmacoepidemiology primarily focuses on epidemiologic methods and their applications in pharmacy. This lecture on matching is scheduled during the latter half of the semester after extensive discussions on study designs, bias, and confounding. This lecture was chosen as it is generally a principle that is not well understood by students and usually sparks thoughtful and interesting discussion. Key principles important for this lecture are reviewed in this manuscript along with an in-depth discussion on the role of matching in epidemiologic studies.

INSTRUCTIONAL METHODS

This lecture is part of a 1-semester, 3-credit graduate course in pharmacoepidemiology, the study of the use and effects of drugs in large numbers of people. This is an elective course for students enrolled in the Master of Science degree program in Pharmacy Administration at the University of Houston College of Pharmacy. These students are generally pharmacists who will have little exposure to epidemiology beyond what is taught during their pharmacy degree. The class primarily uses 2 textbooks, *Epidemiology for Public Health Practice*¹ and *Epidemiology, An Introduction*.² For advanced lectures

Corresponding Author: Kevin W. Garey, PharmD. Address: University of Houston College of Pharmacy, 1441 Moursund Street, Houston, TX 77030. Tel: 713-795-8386. Fax: 713-795-8383. E-mail: kgarey@uh.edu.

(including this lecture on matching), selected chapters from *Modern Epidemiology*³ are chosen. In general, 10–15 students are enrolled in the class. At the beginning of the semester, the class is divided into groups of 2 students each. These groups are then randomly assigned to teach certain classes during the semester. Each week, the presenting group will independently formulate a study plan, prepare lecture notes and handouts, and make overheads or slides. The presenting group is then required to meet with the instructor 2–3 days before the lecture to receive feedback and critique on the topic and study materials. In addition, 2 other groups are assigned additional activities for the class. One group is required to review the previous lecture topic and another group is required to present a real-life application to the current lecture. Finally, each student enrolled in the course is required to write down the key points they learned from the assigned readings before they arrive for class. This critique is then discussed at the end of class along with any homework problems assigned for the week. Performance on a midterm and final examination, along with completion of extensive coursework, are the criteria used for assessing students' performance in the course. Students' course evaluations have generally been positive, with above average scores on the amount of work required and the perceived usefulness of the course to their career.

Content

Pharmacoepidemiology is an applied science that draws from clinical pharmacology, the study of the effects of drugs on humans, and epidemiology, the study of the distribution and determinants of disease in populations. From clinical pharmacology, pharmacoepidemi-

ology draws its focus on inquiry and formulating hypotheses. The methods by which these hypotheses are tested are derived from epidemiology. The graduate course in pharmacoepidemiology primarily focuses on epidemiologic methods and their applications in pharmacy. This lecture on matching is scheduled during the latter half of the semester after extensive discussions on study designs, bias, and confounding. This lecture was chosen as it is generally a principle that is not well understood by students and usually sparks thoughtful and interesting discussion. Key principles important for this lecture are reviewed in the next few paragraphs.

Epidemiologic Study Designs

Several different epidemiology study designs are available for pharmacoepidemiology studies, each with advantages and disadvantages. Study designs are classified as observational or experimental studies. Experimental studies are categorized as community or clinical trials. A clinical trial is designed to compare benefits of an intervention (for example, a drug) with a standard treatment or no treatment. Community trials are similar to clinical trials with the exception that a defined unit such as a state, county, or city is randomized to receive an intervention as opposed to randomization of single subjects in clinical trials. The limitations imposed by ethical considerations and financial costs often restrict the number of experimental studies that can be performed. For this reason, observational studies are much more commonly performed in epidemiology. Descriptive observational studies include case reports and case series which describe novel effects of drugs and ecologic studies which are group comparisons of aggregated data at one point in time. Cross-sectional studies or surveys are also considered descriptive observational studies that compare subject characteristics and outcomes at one point in time. Analytic observational studies include case-control and cohort studies. These 2 study designs are the most commonly performed study type in epidemiology. Thus, many lectures during this course (including the topic on matching) focus on these 2 study designs.

Case-control studies compare people who have a certain outcome (cases) to those who do not (controls) with respect to characteristics or exposure of interest. Cohort studies compare 2 or more groups differing in exposure status and follow these groups over time to determine incidence of a certain outcome. To properly discuss the role of matching in case-control and cohort studies, certain strengths and limitations of these study designs need to be stressed. Study designs are best

understood by visualizing a source population from which persons eligible for the study may be selected. For example, source populations may be diabetics living in Houston, Tex, premenopausal women at risk for breast cancer, or Hispanic males with latent tuberculosis. It is from these source populations that we select our study populations. Cohort studies are easy to visualize as 2 or more groups of people (called study cohorts) that are selected from this source population that are free of the disease or outcome currently being studied. These groups generally differ by one or more exposure variables. These cohorts are then assessed for the development of the disease or outcome being studied. Incidence of the disease or outcome in those exposed and unexposed is compared with assess the contribution of the exposure to the disease or outcome.

Case-control studies are best understood by first defining the source population from which the cases arose. From there, a hypothetical cohort study can be visualized in which the study population is assessed for certain exposure variables and followed for the development of the disease under study. In this hypothetical cohort study, the investigators main task would be to track the number of persons (or person-time) in the exposed and unexposed cohorts to compare the incidence of disease in each cohort. In a case-control study, cases are identified and their exposure status is assessed in the same manner as in a cohort study. However, as opposed to cohort studies, rates from which these diseases arose among the exposed and unexposed cannot be measured. Instead, a control group representative of the source population from which the cases arose is chosen. The purpose of this control group is then to determine the relative numbers of exposed or unexposed persons that developed the disease being studied (as opposed to the cohort study in which the absolute effect can be measured). A central issue with matching (and the main take home message from this section) is that cohort studies sample the source population by exposure status before the disease develops while case-control studies sample the source population by disease status.

Bias in Study Designs

Certain biases are inherent in epidemiologic studies and are important to recognize and control as much as possible. Biases can be categorized as random error or systematic error. The error that remains after systematic error is eliminated is called random error and is nothing more than variability in data that we cannot readily explain. Systematic error is usually classified as selection bias, information bias, or confounding. Selection

bias is a systematic error in a study that stems from the procedures used to select study subjects or influence study participation. Due to the methods to find non-diseased patients described above, selection bias can be problematic in case-control studies. Matching, described below, can also lead to selection bias in many circumstances. Information bias in a study can arise because the information collected from a study subject is erroneous (also called misclassification bias) and can occur in case-control or cohort studies. Confounding is a central issue for pharmacoepidemiology study designs and is defined as a confusion or mixing of effects. Specifically, a confounder is defined as a bias due to the association of a third variable with both the exposure and the disease independently and the failure to disassociate the third variable from the association under study. A confounder may not be an intermediary between the exposure and effect. A classic example of confounding is the observation that obese patients tend to have lower rates of lung cancer than non-obese patients. The confounding variable in this case is smoking. Smokers tend to be lower weight on average than non-smokers and smoking has been causally linked to lung cancer. Thus, smoking is independently associated with obesity as well as lung cancer and would not be considered an intermediary link. Confounding can be controlled either during the design or analysis stage of the study. The analysis techniques commonly used to control for confounding (covered later in the course) include stratification and regression analysis. Three methods to prevent confounding during the design or execution of the study include randomization, restriction, and matching. Randomization is only applicable to experimental studies. Restriction involves prohibiting any person who displays the confounding variable from entering the study. The third method, matching, is the topic of this discussion. This lecture is derived primarily from Chapter 10 of *Modern Epidemiology*.³

DISCUSSION

Definition and forms of matching

Matching refers to the selection of unexposed subjects in a cohort study or to controls in a case-control study who are identical to the cases in certain characteristics. Matching is most frequently used for case-control studies but can be employed in cohort studies as well. In a cohort study, individual matching involves selecting one or more persons from the unexposed cohort who are identical in one or more characteristics to a person from the exposed cohort. In a case-control study, each case is matched to one or more controls who are identical to the

case in one or more characteristics. Regardless of study design, matching on a factor will generally require its control in the analysis. Matching should always be performed with confounding variables as matching by a nonconfounder can actually introduce bias into the study.

In general, many more controls are available than cases in epidemiologic studies. Cases can be matched to 1, 2, or more controls. A common question is “how many controls should I match for each case?” or more specifically, “what is the relative precision of matching r controls compared to $r-1$ controls.” Relative statistical efficiency decreases with each additional control so that matching 3 controls is only 12% more efficient than matching 2 controls, and matching 5 controls is only 4% more efficient than matching 4 controls. Thus, in case-control studies, it is generally considered to be of little benefit to match more than 4 controls for every case.^{4,5}

Matching may be performed at the subject level (called *individual* or *one-to-one matching*) or at a group level (called *frequency matching*). Frequency matching refers to matching on the anticipated distribution of the confounder and can only be used if the a priori distribution of the confounder is known. Unlike subject level matching, frequency matching does not have to wait until the identification of a case; it can operate independently of the cases.⁴

Advantages and Disadvantages of Matching

Matching is generally considered when there is substantial difference in the prevalence of confounders between study groups. However, the improvement in efficiency is generally small unless the variable is a strong confounder. There are other reasons to match in observational studies (Table 1).⁶ Matching can allow for the control of unmeasured confounders that are difficult to measure or obtain. For example, matching by neighborhood controls may allow for equal distributions of certain elements of socioeconomic and environmental status. Matching of identical twins is an extreme exam-

Table 1. Advantages and Disadvantages of Matching in Epidemiologic Studies

Advantages	Disadvantages
Control of unmeasured confounders	Potential additional costs and time
Sufficient controls in sub-analyses	Exclusion of cases without suitable controls
Time comparability	Complex data analysis
Direct control of confounders	Confounder effect cannot be estimated
	Overmatching

Table 2. Hypothetical Source Population Used for an Exercise in Matching, N= 2,000,000

	Total Population		Males		Females	
	Exposed	Unexposed	Exposed	Unexposed	Exposed	Unexposed
# Cases of rash	4,600	140	4,500	50	100	90
N	1,000,000	1,000,000	900,000	100,000	100,000	900,000
Incidence rate			0.005	0.0005	0.001	0.0001
Risk ratio	$\frac{(4,600/1,000,000)}{(140/1,000,000)} = 33$		$\frac{(4,500/900,000)}{(50/100,000)} = 10$		$\frac{(100/100,000)}{(90/900,000)} = 10$	

Adapted from Rothman KJ and Greenland S³ with permission

Note: Hypothetical source population of 2,000,000 persons either exposed or not exposed to a certain drug and who experienced a unique hypersensitivity rash. Considerable confounding exists within this population as the crude risk ratio is significantly higher than the sex-adjusted risk ratios.

ple of this matching. Matching can also help assure that there are sufficient controls to perform certain sub-analyses, especially if the confounder under investigation is rare. In data collected over time, matching by date can also provide time comparability between cases and controls. Finally, matching allows the investigator to directly control the confounder and allows for an automatic adjustment of the confounder.⁴

Unfortunately, matching has many disadvantages as well.^{4,6} Matching can add costs to the study if extra effort is required to identify and recruit the matched controls or if the matching process slows down the data collection and patient recruitment process. Matching may also require excluding cases if no suitable match is identified. Data analysis is more complex if matching has occurred and usually limits the analyses that can be performed. The effect of the counfounder cannot be estimated if a matching technique was employed in the design phase of the study. Finally, overmatching, a term referring to matching performed incorrectly or unnecessarily, may result in loss of study efficiency or biased results.

When is the appropriate time to consider using a matched study design. In general, there are 3 considerations when deciding whether to match variables in a study.⁷ The matched variable should be related to both exposure and disease (ie, a confounder), the confounder should be fairly significant so that matching will increase the statistical precision of any estimation, and finally, it should not be overly difficult to apply the matching technique.

Purpose and Effect of Matching

A common misconception concerning matching is that its purpose is to enhance the validity of observational studies. The role of matching is not to increase validity, but rather to increase the efficiency of a study. However, depending on an investigator's definition of efficiency (sample size or economic costs), matching can be shown to actually decrease study efficiency in some circumstances. We will begin this discussion using an example of matching in a cohort and case-control study.

Consider a study that wishes to investigate whether a certain drug used to treat cancer causes a unique hypersensitivity rash. Study investigators highly suspect that male patients are much more likely to display this hypersensitivity reaction, and this drug is much more commonly prescribed to male patients than female patients. A common question posed to investigators would be "should we match on gender in this study?" In this hypothetical example, we will construct a source population in which we "know" the answer to this question. We will then construct cohort and case-control studies using matching techniques to see how matching would influence our results from the true answer. In this example, consider that the source population is comprised of 2 million people equally divided into 1 million males and 1 million females. Half of the population (1,000,000) has been exposed to the drug and 4,740 persons have been diagnosed with the hypersensitivity rash. True to the belief of the investigator, the exposure (drug) and disease (hypersensitivity rash) are both much more common in male patients than in female patients. Male patients accounted for 900,000 of the persons exposed to the drug and 4,550 of the hypersensitivity cases were in male patients. A depiction of this case is shown in Table 2. There is considerable confounding in this example as 90% of the persons given the drug are male and male patients are much more likely than female patients to get the rash. This is highlighted by the fact that the crude risk ratio (33) was much different than the gender-stratified risk ratio (10). It is important to pay attention to the "true" answer to this question, which is that persons who take the drug are 10 times more likely to develop this rash than persons not taking the drug. Male patients, for an unknown other reason, develop this rash through another cause and are also much more likely to receive this drug, thus confounding the results.

For our first study derived from this hypothetical source population, we will perform a matched cohort study of 100,000 patients selected from this source population. We will match each person exposed to the drug,

Table 3. Expected Results from a Matched Cohort Study (Matched on Gender) of 100,000 Patients

	Males		Females	
	Exposed	Unexposed	Exposed	Unexposed
# Cases of rash	450	45	10	1
N	90,000	90,000	10,000	10,000
Expected risk ratio	$\frac{(450/90,000)}{(45/90,000)}$	= 10	$\frac{10/10,000}{1/10,000}$	= 10
Crude risk ratio	$\frac{(450+10)/(90,000+10,000)}{(45+1)/(90,000+10,000)}$		= 10	

Adapted from Rothman KJ and Greenland S³ with permission

Note: Expected results from a matched cohort study (matched on gender) of 100,000 patients exposed to a certain drug expected to cause a unique hypersensitivity rash and 100,000 patients not exposed to the certain drug. These patients were sampled without error from the population described in Table 2. In this case, matching by gender successfully removed the confounding.

Table 4. Expected Results From a Matched Case-control Study (Matched on Gender) of 4,740 Patients

	Total		Males		Females	
	Exposed	Unexposed	Exposed	Unexposed	Exposed	Unexposed
Cases	4,600	140	4,500	50	100	90
Controls	4,114	626	4,095	455	19	171
Risk ratio	$\frac{(4,600 \times 626)}{(4114 \times 140)}$	= 5	$\frac{(4,500 \times 455)}{(50 \times 4,095)}$	= 10	$\frac{(100 \times 171)}{(19 \times 90)}$	= 10

Adapted from Rothman KJ and Greenland S³ with permission

Note: Expected results from a matched case-control study (matched on gender) of 4,740 patients with a hypersensitivity rash (cases) and 4,740 patients who did not experience a rash (controls). These patients were sampled without error from the population described in Table 2. In this case, although the gender-specific risk ratios give the correct value of 10, the crude risk ratio is much less than the true risk ratio.

with an unexposed person from the same gender. In the source population, there are 90% (900,000) exposed males and 10% (100,000) exposed females. Thus, assuming we were able to appropriately select from our source population, out of 100,000 persons randomly selected, we would expect to have 90,000 males and 10,000 females. We would then randomly match these 90,000 exposed males to 90,000 unexposed males and the 10,000 exposed females to 10,000 unexposed females. Using the expected incidence rates (assuming no sampling error) calculated in Table 2 (50 cases and 5 cases per 10,000 exposed and unexposed males, respectively, and 10 cases and 1 cases per 10,000 exposed and unexposed females, respectively), we would anticipate 450 cases of rash among the exposed males, 45 cases among the unexposed males, 10 cases among the exposed females, and 1 case among the unexposed females. Results from this hypothetical cohort study are shown in Table 3. Matching in this study effectively removed the confounding from the study results. The crude risk ratio as well as the gender stratified risk ratios now all show a 10-times increased risk of rash for persons exposed to the oncology drug. Thus, matching in this cohort study was effectively able to remove the confounding from the results.

We will next consider a matched case-control study derived from the same source population. In this case, we will assume that the investigators were able to identify all

4,740 cases of the hypersensitivity rash (4,550 male and 190 female). Assuming no misclassification bias, results of the study will show that 4,500 male cases were exposed to the drug and 50 were not exposed. Likewise, the results will show that 100 female cases were exposed to the drug and 90 were not. In total, 4,600 cases were exposed to the drug and 140 were not. From there, gender-matched controls (persons without the hypersensitivity rash) will be randomly selected and assessed for exposure. Recall that 90% of males and 10% of females were exposed to the drug. Thus, from the 4,550 male controls selected, we would anticipate that 90% of these (4,095) were exposed to the drug and 10% (455) were not exposed. Of the 190 control females, we would anticipate that 10% (19) of these women were not exposed to the drug and 90% (171) were exposed. Results from this study are shown in Table 4. The results from this study are very surprising. Although the gender-specific rate ratio correctly identifies a 10-times increased risk of developing a rash in patients given this drug, the crude rate ratio is much less than the true risk ratio. Thus, unlike cohort matching, case-control matching has not eliminated confounding in the crude estimate of the risk ratio. The matching in the case-control study has introduced a selection bias by selecting controls related to the exposure (gender). This bias is similar to confounding, but note that in this case matching has changed the direction of the bias.

The following example illustrates the following effects of matching: In a well-executed cohort study (without competing risks, loss to follow up, or other bias), matching effectively removes confounding from the analysis. In case-control studies, matching can introduce a selection bias if the matching variable is associated with the exposure in the source population. The analysis requires stratification by the matching factor even if the variable is not a risk factor for the disease. What causes these discrepant results? In a cohort study, matching is based on exposure status (without regard to disease status). This alters the distribution of the matching factor in the entire study population. In contrast, matching in a case-control study involves selection of a matched control for every identified case. Thus, only the distribution of the control variables is affected by matching. If the exposure is related to the matched variable, this will introduce a selection bias into the analysis.

Matching in Case-Control Studies

Matching in case-control studies appears to be especially prone to error. What are the properties of case-control studies that make these study designs especially prone to bias? Controls selected for a case-control study are supposed to reflect the distribution of exposure in the source population. If the matching variable influences the exposure variable, this will change the distribution of the control population away from the distribution in the source population. This will introduce a selection bias that is very similar to confounding. If the matching variable does not influence the exposure variable, no bias will be introduced. However, in many research studies, the influence of the matching variable on the exposure of interest is not well appreciated or is unknown. Worse yet, if the matching variable is not related to the disease it would not be a confounder in an unmatched case-control study. However, if the matched variable is associated with exposure, it will introduce a controllable selection bias into the study where none existed before.

The above discussion may make one start to wonder why an investigator would want to consider matching at all in a case-control study. The above discussion illustrates that matching does not eliminate confounding; however, it may enhance the statistical efficiency of a study by controlling confounding. For example, suppose a study is being conducted in which it is known that the age of the subjects will confound the exposure-disease relationship being studied and that the age distribution of the cases is much older than the source population. If the results of the analysis are divided into various age strata, many strata will likely contain many cases and only a

few controls. If cases were matched to controls by age, each age strata would have a constant ratio of cases to controls. This is especially important in studies in which resources only permit selection of a certain number of controls. Unfortunately, in strata with cases in which very few controls exist, identifying controls may entail a considerable amount of work. Also, if matching is performed on a variable that is only associated with exposure and not disease (ie, a nonconfounder), matching may introduce selection bias into the study. In this case, non-matching may be preferable.

By improving the statistical efficiency of a study, matching would appear to decrease the financial costs of a study by requiring fewer control subjects. However, even this fact has been debated. If several factors are being matched, it may be difficult to identify appropriate controls, and more time, effort, and money may be required to identify these matched controls. Thus, if efficiency is considered as cost efficiency and not size efficiency, matching may decrease the efficiency of study designs.

There are certain, limited situations in which matching is desirable or necessary. If exposure assessment can only be obtained at a great expense (for example, an expensive laboratory test), the costs of finding matched controls may be less than the exposure assessment. Another situation is coined the "sparse-data" problem in which it is well recognized that each strata will have a limited number of potential controls. Without matching, each strata may have only one case or one control and would not be useful for analysis. Also, if a large number of confounders are suspected and the investigator wishes to control for each variable, matching may be required to assure that each strata contains an adequate number of controls. However, as mentioned previously, the cost of finding these controls may outweigh the potential for collecting this information. In summary, matching in case-control studies is a useful means of improving study efficiency in terms of the amount of information required per subject studied in some situations. However, it has the potential to introduce or augment study bias and should be approached cautiously in case-control studies.

Matching in Cohort Studies

Unlike case-control studies, matching in cohort studies can effectively eliminate confounding from the study analysis. Despite this, matched cohort studies are quite uncommon for multiple reasons. Cohort studies generally require a much larger sample size than case-control studies, and identifying the matched, unexposed cohort would require such a large amount of time and effort that

it would be unreasonable to conduct. Also, in certain circumstances, matching does not eliminate the need to control for the matching factor in the analysis (for example when the exposure and matching factors affect disease risk, competing risk, or loss to follow up).

Analysis of Matched Data

Further lectures in this course teach how to appropriately analyze the complex data produced by using a matched design experiment. Chapter 15, "Introduction to Stratified Analysis," and Chapter 16, "Applications of Stratified Analysis Methods," in *Modern Epidemiology*³ provide an excellent overview on analyzing matched data. Other sources that detail the analysis of matched data include Chapter 6, "Case control studies," in *Epidemiology: Study Design and Data Analysis*⁴ and Chapter 9, "Case-control studies: II. Further design considerations and analysis" in *Methods in Observational Epidemiology*.⁷

SUMMARY

In summary, matching is a commonly applied tool to control for confounding in epidemiologic studies. Whether or not to perform matching in a study is not an easy question to answer. An investigator sometimes feels compelled to match on a certain factor (eg, age) simply because of peer pressure from the research community. Because of the potential for introducing bias into the study, if the matching variable is not a confounder, matching should not be employed. For true confounders, matching will likely increase study efficiency by requiring a smaller sample size. However, for case-control

studies, matching sometimes does more harm than good by introducing bias into the study. Cohort studies rarely employ matching due to the prohibitive costs of identifying this matched, unexposed cohort. Whether to match in a study depends upon several factors including the strength of the relationships, the confounding variables, and the desire to include the confounding variables in the analysis. In case-control studies, matching should be used when statistical efficiency is required due to a small or costly control sample. In most other cases, matching should be approached cautiously if at all in case-control studies.

ACKNOWLEDGEMENTS

I would like to acknowledge Melinda Neuhauser, PharmD for careful review of this manuscript.

REFERENCES

1. Friis RH, Sellers TA. *Epidemiology for Public Health Practice*. 2nd ed. Gaithersburg: Aspen Publications; 1999.
2. Rothman KJ. *Epidemiology. An Introduction*. 1st ed. New York: Oxford University Press; 2002.
3. Rothman KJ, Greenland S. *Modern Epidemiology*. 2nd ed. Philadelphia: Lippincott Williams & Wilkins; 1998.
4. Woodward M. *Epidemiology: study design and data analysis*. Boca Raton, FL: Chapman & Hall/CRC Press; 1999.
5. Ury HK. Efficiency of case-control studies with multiple controls per case: continuous or dichotomous data. *Biometrics*. 1975;31:643-649.
6. Wacholder S, Silverman DT, McLaughlin JK, Mandel JS. Selection of controls in case-control studies. III. Design options. *Am J Epidemiol*. 1992;135:1042-1050.
7. Kelsey JL. *Methods in Observational Epidemiology*. Vol 26. 2nd ed: New York: Oxford University Press; 1996.