Journal of Quantitative Criminology, Vol. 17, No. 4, December 2001 (© 2001)

Alternative Strategies for Identifying the Link Between Unemployment and Crime

Steven D. Levitt[1]

National-level time series data are a crude tool for distinguishing between two alternative behavioral explanations for a link between unemployment and crime. Consequently, inferences drawn from aggregate time series estimates are likely to be misleading. A more fruitful approach to learning about the link between unemployment and crime would be to utilize a menagerie of different methodological approaches such as cross-section and panel data analysis of less geographically aggregated areas, natural experiments, international data, individual-level data, and ethnography.

KEY WORDS: unemployment; crime; time series; methodology; natural experiments.

1. INTRODUCTION

The papers included in this symposium represent applications of state-of-the-art modern time series techniques applied to the question of the link between unemployment and crime using modern time series techniques. All of these papers are cogently argued and clearly written. Collectively, they make an invaluable contribution to the literature on the analysis of time series data in criminology. Unlike the other papers in this symposium, I have little to say about the particulars of time series estimation. Rather, I aim my discussion at a few broader points related to criminological research and the advancement of knowledge. Much of what I write is obvious and probably well understood by both the participants in the debate and the readers. Nonetheless, in my opinion these points bear repeating in the present context because they have not received adequate attention in the other papers in this volume.

The main point that I emphasize in this article is that national-level time series data are an extremely crude tool for answering criminological

[1]Department of Economics, University of Chicago, 1126 East 59th Street, Chicago, Illinois 60637. E-mail: slevitt@midway.uchicago.edu.

questions. There are certain circumstances in which such data are the ideal means of studying an issue, the particular question at hand—distinguishing between two alternative behavioral explanations for a link between unemployment and crime—does not satisfy these criteria. Consequently, inferences drawn from aggregate time series estimates are likely to be misleading.

By focusing so narrowly on this one particular methodological approach, the papers in this issue unintentionally make statistical methods, rather than understanding crime, the focus of the research. Of course, appropriate use of statistical techniques is critical to enhancing our understanding of the world. My contention, however, is that as technically advanced and elegant as the papers in this symposium are, if the goal is to learn about the link between unemployment and crime, the journal pages could have been better spent by devoting this issue to a menagerie of methodological approaches including cross-section and panel data analysis of less geographically aggregated areas, natural experiments, international data, individual-level data, and ethnography.

The remainder of this article is structured as follows. Section 2 considers the factors that influence the usefulness of aggregate time-series data for understanding a problem. Section 3 outlines a range of other approaches to analyzing the link between unemployment and crime. Although each of these alternatives has its own weaknesses, I argue that a portfolio of approaches is likely to ultimately shed more light on the issues. Section 4 concludes.

2. WHAT CAN WE LEARN FROM NATIONAL TIME SERIES DATA?

There are certain circumstances in which national time series data are the natural framework with which to analyze a problem. Such data are an ideal tool for describing long-run patterns and dynamics in macro variables. For instance, aggregate time series data certainly would seem to be a logical choice for studying the pattern of economic growth, inflation, or the national debt over the 20th century, as well as the relationship among these variables. The use of national data is particularly appropriate for studying these phenomena since all of them are inherently national in character, i.e., there is relatively little local variation, so there is little gain from analyzing more geographically disaggregated data.[2]

On the other hand, if there really is local variation in the variables that are being analyzed, then relying on national time series data fails to take

[2]In fact, when variation is strictly national, one may erroneously find that using smaller units of observation yields more precise estimates, but this is merely an artifact of failing to correct properly for correlated errors across the units of observation.

advantage of the full information provided by the data. National data average across all of these local fluctuations, removing potentially useful variation. Crime and unemployment clearly are variables that exhibit a tremendous degree of variability at the local level. The enormous variance across time and place in criminal activity is well documented (Lynch, 1995; Wilson, 1983). Similarly, unemployment is quite localized. For instance, in the 1990 Census, unemployment rates by neighborhood in the city of Chicago vary between 3 and 45%. State unemployment rates in April 2000 range from 2.2 (Iowa) to 6.6 (Alaska). Simply summarizing the Chicago unemployment rate in 1990 as 13%, or the national unemployment rate in April 2000 as 3.9, ignores all of this variation.

A second weakness of national time series data is that they have little potential for identifying causal (as opposed to merely correlational) relationships between variables. Because of the severe limitation on the available degrees of freedom, it is generally not possible to include a wide range of covariates in the analysis.[3] To interpret the parameters on the variable of interest as causal, one needs to believe that all other factors influencing crime are taken into account. Examples of time series analysis leading to parameter estimates that are not reasonably interpreted as causal are widespread. Marvell and Moody (1997), using annual national aggregated data on homicides and the prison population over the period 1930–1994, report an elasticity of homicide with respect to the prison population of −1.3. Back-of-the-envelope calculations suggest that this coefficient is implausibly large; indeed it is almost 10 times higher than the estimates obtained from other approaches (Marvell and Moody, 1994; Spelman, 1994; Levitt, 1996).

Another instance in which aggregate time series estimates generate suspicious results is the impact of changes in the age distribution on crime. A priori, one would expect that aggregate time series analysis would be an ideal source of identification for demographic characteristics: these variables tend to move slowly and there is little local variation to exploit. The age distribution example is particularly useful because data on arrests by age are available to compute directly a back-of-the-envelope calculation of a reasonable magnitude for the coefficient.[4] The relationship between age and criminal involvement at the level of the individual is one of the most well-known and robust relationships in all of criminology (Goring, 1913; Wilson and Herrnstein, 1985; Blumstein *et al.*, 1986) and dates back to Quetelet

[3]Land *et al.* (1990) highlight the substantial degree of colinearity between potential determinants of crime in cross-sectional regressions. Those authors use a factor-analytical approach to reduce the dimensionality of the system.

[4]This is in contrast to the unemployment coefficient in a crime regression, for which we have no definitive benchmark with which to compare the time series estimate.

(1831). There is a sharp rise in criminal involvement with the onset of adolescence, followed by a steady decline with age. As noted by Marvell and Moody (1991), however, most of the aggregate time series studies that attempt to estimate a link between the population of young males and crime uncover only weak evidence for such a relationship.[5] The results reported in Tables II and III of Greenberg (2001) provide a more recent example. Although the variable Greenberg uses to capture changes in the age distribution (the percentage of males between 15 and 29 years of age) enters regressions of homicide and robbery rates with a positive coefficient that is often statistically significant, even the largest of these estimates are an order of magnitude smaller than one would expect based on reasonable methodologies based on the distribution of arrests by age (Levitt, 1999; Steffensmeier and Harer, 1987).[6]

Turning to the Greenberg (2001) estimates of the impact of unemployment on crime, the only other factors included in the specification explaining crime are the percentage of males aged 15–29 and the divorce rate. Yet other factors that influence crime have certainly changed over time. For instance, there has been a quadrupling of the prison population over the last 25 years that undoubtedly must have dampened crime, even if not as much as the Marvell and Moody (1997) estimates suggest. Over this same period, there have been dramatic changes in income inequality that may have exacerbated crime (Freeman, 1983; Chiricos, 1987; Land *et al.*, 1990).

A third reason why national-level data seem ill suited for the debate taking place in this symposium is that the focus is not the overall link between unemployment and crime but, rather, fairly subtle predictions about a variety of possible behavioral channels through which the crime–-unemployment nexus operates. As the symposium authors themselves recognize, differentiating between these behavioral dimensions through the use of various lags in the timing of unemployment is extremely crude.

3. ALTERNATIVE METHODOLOGICAL APPROACHES TO UNDERSTANDING THE UNEMPLOYMENT–CRIME LINK

If it were the case that there were no other data available to study this issue, then the exclusive focus on national time series data could be justified

[5]Cohen and Land (1987) provide a notable exception to this tendency.

[6]It is possible that the methods that rely on individual-level data are biased because of indirect effects of changes in the age distribution on crime. I am skeptical of this argument for two reasons. First, most plausible stories of indirect links between cohort size and crime imply that being part of a large cohort encourages crime, for instance, due to more competition in the legitimate labor market. Second, empirically, there seems to be little evidence to support a non-linear relationship between cohort size and per capita crime rates (O'Brien, 1989; Steffensemeier *et al.*, 1987).

as a necessary concession to data limitations. In reality, however, there are many alternative means of identifying the relationship between unemployment and crime. In this section, I discuss these approaches, some of which have already been put into practice in previous studies and others that I suggest as the possible focus of future research.

This section is organized according to the different sources of variation available: correlational panel-data studies, natural experiments, international data, and individual-level data. Within each category, I briefly discuss the pros and cons of the method, any existing estimates, and possible future applications.

3.1. Panel Data

A panel-data set has repeated observations over time for a set of individuals, subnational geographic units, or countries. Panel data have both time series and cross-sectional variation, unlike the national time series data. In the study of crime and unemployment, one can easily obtain data series on annual state-level crime and unemployment dating back many decades.

There are three primary benefits of panel data. First, because there are multiple observations per year, one can remove year fixed-effects. Thus, any unobserved shocks that affect the entire country (e.g., changes in the age distribution, national politics, etc.) can be controlled for, even if the shocks are not easily quantified. Similarly, state-fixed effects (and state trends) can also be included in the analysis so that comparisons are not made across states, but only using within-state deviations over time. Again, this allows one to control for differences across places that are not easily quantified. New York and Idaho are clearly very different states, and they differ along so many dimensions that it is likely very difficult to capture the differences fully using typical covariates. With panel data, state-fixed effects eliminate anything consistent about a state over time—only time-varying characteristics need to be taken into account. Although there is nothing explicitly causal about panel data estimates, by eliminating these important sources of omitted variables, one may obtain coefficients that come closer to representing a causal impact.

Second, the high number of degrees of freedom makes it possible to control for a wide range of time-varying factors that might plausibly be linked to both crime and unemployment rates and therefore lead to spurious coefficients in a national time series, e.g., state prison populations, alcohol consumption, and changes in income inequality.

Third, to the extent that there is area-specific variation in crime and unemployment, disaggregating the data to a more local level allows the researcher to make use of variation that is squandered with a national time series.

Panel data also have weaknesses. By including state- and year-fixed effects, only the short-term relationship between the variables will be reflected in the parameter estimates. If there is a high degree of correlation in variables across areas, or over time, there will be little remaining variation with which to identify the coefficients. For the case of unemployment and crime, however, these concerns are not paramount. A second weakness of panel data in the specific context of the questions at hand is that they provide no escape from having to model crudely the different behavioral stories relating unemployment and crime through the use of relatively arbitrary lags in unemployment's effect on crime.

Panel-data analyses using states, counties, metropolitan statistical areas, or cities in the United States have generally obtained relatively consistent estimates of the impact of unemployment on crime. A 1% change in the unemployment rate is typically found to increase property crime by 1–2% contemporaneously but often has no systematic impact on violent crime rates (Lee, 1993; Levitt, 1996, 1997; Raphael and Winter-Ebmer, 2000).[7] Studies that substitute other measures of the labor market conditions at the bottom of the distribution reach similar conclusions (Gould *et al.*, 1998; Machin and Meghir, 2000). The consistency of these results across data sets, included covariates, and degree of aggregation is encouraging, as the national time series data yield results that are much more sensitive to the particulars of the estimation.

I am unaware of any previous panel-data analysis that attempts to relate systematically lags in unemployment to crime rates. Table I presents a preliminary attempt at such an exercise. Using a state-level panel of annual data for the period 1950–1990, I run regressions of the form

$$\text{Crime}_{st} = \beta_1 \text{Unemp}_{st} + \beta_2 \text{Unemp}_{st-1} + X_{st} + \theta_s + \gamma_t + \varepsilon_{st} \tag{1}$$

where s corresponds to states, and t indexes years. Crime reflects logged official crime rates per capita from the Uniform Crime Reports. I use the standard violent and property crime definitions, except that rape is excluded from violent crime since it was not collected in the early part of the sample, and larceny is excluded from property crime because of important changes in definition over the course of the sample. *Unemp* is the state insured unemployment rate (this variable closely approximates the standard unemployment rate but is available for a longer time period at the state level). The unemployment rate is included both contemporaneously and once lagged. X is a vector of covariates including the prisoners per violent crime [once lagged to minimize endogeneity (see Levitt, 1996)], the state execution rate per 1000 prisoners, the state real per capita income, the infant mortality

[7]See Freeman (1995) for a survey of this literature.

Table I. The Relationship Between Unemployment and Crime in State-Level Panel Data, 1950–1990[a]

Variable	Property crime per 100,000 residents			Violent crime per 100,000 residents		
Insured unemployment rate						
Contemporaneous	.027	.014	.018	.005	−.013	−.001
	(.008)	(.005)	(.006)	(.014)	(.010)	(.009)
Once-lagged	.004	.001	−.000	−.012	−.022	−.010
	(.008)	(.006)	(.006)	(.014)	(.010)	(.009)
Executions/1000 prisoners	.028	.008	.028	.023	.011	.028
	(.014)	(.010)	(.013)	(.025)	(.018)	(.019)
ln(prisoners/100,000 residents)						
(once lagged)	.050	−.142	−.076	.025	−.072	−.129
	(.025)	(.035)	(.027)	(.038)	(.039)	(.036)
Real per capita income (*1000)	.021	−.012	−.005	.040	−.011	−.000
	(.007)	(.007)	(.008)	(.012)	(.010)	(.011)
% Black	−1.71	3.08	7.03	2.19	3.83	2.42
	(.53)	(1.31)	(.96)	(1.01)	(1.62)	(1.27)
% urban	−2.66	−2.40	−4.70	−3.06	−3.38	−3.85
	(.25)	(.51)	(.63)	(.46)	(1.17)	(1.26)
% 0–24 yr olds	1.51	1.09	.91	.88	−.78	−.79
	(.71)	(1.25)	(1.16)	(1.33)	(1.45)	(1.09)
% 25–44 yr olds	−1.59	−5.34	−6.33	−.00	−1.03	−6.14
	(.70)	(1.33)	(1.28)	(1.33)	(1.98)	(1.55)
						(1.51)
Adj. R^2	.938	.972	.965	.914	.969	.964
State trends?	No	Yes	No	No	Yes	No
State–decade interactions?	No	No	Yes	No	No	Yes

[a]The dependent variable is the property crime rate (excluding larceny) per 100,000 residents. Larceny is excluded because of changes in its definition over the sample period. Rape is excluded because data are unavailable for the early part of the sample. Data are annual, state-level observations for the period 1950–199. Data for 1971 are missing. The regressions are weighted by the state's share of the nations's population. State-fixed effects and year dummies are included in all specifications, except where they are redundant. The number of observations is 1903 in the property crime regressions and 1889 in the violent crime regressions.

rate per 100,000 live births, the percentage Black, the percentage urban, the percentage age 0–24, and the percentage age 25–44. Fixed effects for both years and states are included, meaning that the identification of the model comes solely through within-state changes over time. In some specifications, state–decade interactions or state trends are also added. All of the regressions are estimated using weighted least squares, with weights determined by state population.

Results for property crime are presented in the columns 2–4 in Table I. The coefficients on contemporaneous and once-lagged unemployment rates are fairly consistent across these three specifications (in contrast to many of the other covariates, which are quite sensitive to specification). The robustness of these coefficients is encouraging given that the source of variation differs substantially across columns. The second column is a straightforward panel regression with year dummies and state-fixed effects. Given the long time series (40 years), the assumption that other factors remains constant is suspect. The third column introduces state–decade interactions, so only variation within a state and decade is used to identify the parameters. The fourth column includes state-specific linear time trends. A 1% increase in the unemployment rate is associated with a 1.4–2.7% increase in property crime. The lagged unemployment rate does not have a statistically significant impact on property crime. This pattern of coefficients is quite different from those of either Greenberg (2001) or Cantor and Land (1985), casting doubt on a causal interpretation of those authors' results.

None of the unemployment variables are ever statistically significant in the violent crime regressions. The sign of the coefficients flips across specification. Once again, it is important to note that the pattern of coefficients that emerges in the national time series data are absent in a panel-data analysis.

Although the other variables in the model are of only secondary interest, they provide some insight into the strengths and weaknesses of panel-data analysis. For those variables that exhibit little year-to-year, state-specific variation (percentage Black, percentage urban, and the age variables), the coefficients are extremely sensitive to specification and estimated imprecisely. This is because virtually all of the signal is absorbed in the state and year dummies. Note also that many of the coefficients change substantially when one moves from a single state-fixed effect over the 40-year period to state–decade interactions. For instance, the impact of the prison population on crime goes from an implausible positive sign to a negative coefficient with magnitudes similar to those obtained using panel data on more recent samples (Marvell and Moody, 1994; Levitt, 1996). These results suggest that a single state-fixed effect does not effectively capture unobserved differences across states over such a long time period.

3.2. Natural Experiments/Instrumental Variable Estimates

Both correlational panel-data estimates and aggregate time series analysis have the drawback that the estimated parameters do not have an explicitly causal interpretation. For instance, if having a criminal record makes it difficult to find a job, or high crime rates in an area drive away businesses or customers, then crime may cause unemployment, rather than vice versa. On the other hand, there may be factors that lead the correlation between unemployment and crime to be understated. For instance, if alcohol consumption falls when the economy is doing poorly, and alcohol use leads to crime, then failing to account for the alcohol–crime link may lead to a spurious result in which more unemployment leads to less crime. In that case, however, it is the changing alcohol consumption, not the unemployment, which is playing the causal role. From the perspective of public policy, only a causal relationship is of interest.

Absent randomized experiments, "natural" experiment or instrumental variables analysis provides a possible means of isolating causal effects. The key to the natural experiment approach is to identify an "exogenous" source of variation in the variable of interest, in this case unemployment rates. In other words, an instrument is needed that affects unemployment but has no impact on crime, except through the change in unemployment rates. If such a variable can be found, then the instrumental variables estimate will be free of the spurious influences that potentially bias OLS and other correlational estimates.

There are three major weaknesses of the natural experiment approach.[8] First, in many natural experiments, the instrumental variable is only weakly related to the variable of interest (see, e.g., Levitt, 1997), leading parameter estimates to lack precision and robustness. Second, the critical assumption underlying the validity of this approach—the exogeneity of the instrument—is not directly testable.[9] Thus, the researcher must provide auxiliary arguments for the exogeneity of the instruments, rather than relying solely on statistical arguments. Third, the estimates generated by instrumental variables analysis may not be easily generalizable out of the particular setting in which they arise. For instance, if the natural experiment is the permanent closing of a steel mill in a small town (Black and Sanders, 2000), this might lead to an increase in unemployment that is dominated by very long spells, substantial out-migration, and permanent underemployment for those who eventually find jobs. Such a plant closing may entail very different impacts on crime than a short-term increase in unemployment associated with the national business cycle.

I am aware of only one study that provides instrumental variables estimates of the link between unemployment and crime.[10] Raphael and

[8] For an exhaustive treatment of the strengths, weaknesses, and limitations on the interpretation of instrumental variables estimates, see Cameron and Heckman (1999).

[9] Although when there is more than one instrument, one can perform a test of the exogeneity of the overidentifying restrictions, under the null hypothesis that one of the instruments is valid.

[10] Although neither investigates the competing channels that are the focus of the papers in this symposium, there is nothing preventing an extension of their analyses in that direction.

Winter-Ebmer (2000) use two types of instruments for unemployment in a state-level panel data set: the closing of military bases and shocks to oil prices. Raphael and Winter-Ebmer (2000) demonstrate that military base closings over the last two decades have induced substantial variation across states in unemployment. They argue that the base closings are otherwise unrelated to crime, once other observable factors such as the imprisonment rate, demographic composition, alcohol consumption, and percentage in poverty have been controlled for. Raphael and Winter-Ebmer (2000) also show that fluctuations in oil prices have very different effects on employment rates across states, and they again argue that the primary impact of such shocks on crime is operating through changes in unemployment rates. Raphael and Winter-Ebmer (2000) find that instrumenting for unemployment rates using military contracts and oil shocks leads to more negative estimates of the impact of unemployment on property crime (each percentage point increase in unemployment leads to a 3.9% increase in property crime). Neither OLS nor 2SLS leads to significant effects on unemployment on violent crime. There are justifiable reasons to be skeptical of the results of this paper on both empirical grounds (it is surprising that the 2SLS estimates are larger than the OLS estimates for property crime—most stories would suggest that the coefficient should have shrunk) and from a theoretical perspective (the assumption that the instruments are exogenous here, as usual, can be challenged). Nonetheless, the natural experiments the authors examine represent a different source of identification and thus are a useful contribution to the literature.

3.3. International Data

Cross-country data provide a very different window onto the link between unemployment and crime. There are three primary attractions of international data. First, there is an enormous amount of variation across countries in both their short-run and their long-run unemployment rates. For instance, in 1998, Switzerland had an unemployment rate of 3.6%, Denmark of 5.5%, Italy of 12.3%, and Spain of 18.8%. Second, cross-country studies utilize a completely different source of variation than either national time series or panel-data studies within the United States. To the extent that similar results are obtained across different data sources, our confidence in the robustness of the results is enhanced. The third benefit of international data is that the coefficients from a cross-country regression can reasonably be interpreted as a long-run relationship between unemployment and crime. The obvious drawbacks of international data are the possibility of noncomparability of data and definitions across countries (Neapolitan, 1997), as well as the standard omitted variable concerns that arise in cross-sectional regressions. Use of a panel of international observations may partly alleviate the latter of these concerns.

I am personally unaware of any systematic studies of the unemployment–crime link using international data. This is a surprising omission in the literature given that there have been a number of papers written about the link between crime and income inequality, income, and poverty using international data (Krohn, 1978; Stack, 1984; Soares, 1999).

3.4. Individual-Level Data

Given that Cantor and Land (1985) and Greenberg (2001) are attempting to differentiate between the opportunity and the motivational channels for unemployment affecting crime, individual-level data would seem to be the most plausible starting point for the analysis.[11] As Greenberg himself notes, when one wants to test a theory formulated for individuals, it is preferable to obtain data for individuals.

The stumbling block to conducting such individual-level analyses is the paucity of longitudinal data sets with information on criminal activity. The National Longitudinal Study of Youth (NLSY) is the individual-level data set most frequently used in this respect, although it suffers from the great drawback of having direct information on criminal activity in only 1 year.[12] The National Longitudinal Study of Adolescent Health, which has an extensive set of questions on criminal activity in each interview, may eventually yield more convincing evidence (see, e.g., Mocan and Rees, 1999).

Even absent a comprehensive longitudinal, individual-level data set, there is the possibility of making headway on this problem. For instance, researchers have linked arrest and conviction data to official records of state parole offices (Lott, 1992) or to unemployment insurance agencies (Grogger, 1995; Kling, 1999). Although the focus of these previous studies has been to determine the impact of arrest or conviction on future earnings, this same methodology could be used to analyze the relationship between an individual's employment status and the timing of arrests. For instance, Kling (1999) demonstrates that among convicted criminals, employment rates and wages begin to fall 2 to 3 years before their conviction. Because Kling has only the conviction date and not the date that the crime is committed, it is impossible to determine precisely how employment status affects

[11] Although as Cantor and Land (2001) argue, their initial paper was an attempt to capture not only a *direct* effect of unemployment on crime, but also a *contextual* effect. Nonetheless, if one wants to differentiate between these two effects, individual-level data are critical.

[12] One can ascertain whether the respondent is interviewed in prison in a given year, providing some indirect evidence on criminality.

crime rates. With a slightly different data set, however, such an exercise would be feasible.

A very different analytical approach to individual-level analysis involves ethnography or surveys. For instance, Wright and Decker (1994) interview a large number of burglars. Many of the burglars that they interview say that they commit burglaries only when they do not have sufficient cash to meet current expenses. That provides some evidence of a link between unemployment and crime at the individual level. However, it is also true that 17 of the 95 burglars who report committing their crimes primarily to raise money were employed. Other offenders report that having legitimate jobs that allowed them to enter homes (e.g., repair or delivery jobs) enhanced their ability to burgle by allowing them to identify promising targets.

Although I am unaware of any existing prisoner surveys that ask in great detail about the employment history of inmates or precise motivations that led them to commit their crimes, such information would provide yet another window onto the issues at hand.

4. CONCLUSION

Debate on matters of empirical technique are critical to the advancement of the field. The proper use of statistical methods is an important input into understanding criminological questions. For that reason, the papers in this symposium are valuable contributions. My own view, however, is that excessive focus on empirical methods—particularly ones that seem ill-suited to answering the questions posed—can also prove to be a distraction. National time series analysis is just one approach to untangling the link between unemployment and crime. In my opinion, aggregate data are far from the best tool for the job: it wastes local variation in unemployment and crime, does not allow for a wide range of covariates, and yields coefficients that do not have a causal interpretation. The usefulness of aggregate time series estimates in this realm is further called into question by the fact that state-level panel data reveal a very different pattern of coefficients with respect to the link between unemployment and crime.

While other approaches that I propose—more disaggregated panel data approaches, natural experiments, international data, and individual-level analysis—each has its own weaknesses, collectively they are likely to shed far more light on the issues at hand than national time series data. Surprisingly, however, there has been relatively little previous research on the subject utilizing those alternative strategies. Correcting this oversight presents an exciting opportunity for future research.

ACKNOWLEDGMENTS

Richard Freeman, Edward Glaeser, Austan Goolsbee, Lars Hansen, and Jeff Kling provided insightful comments on this research. Financial support from the National Science Foundation is gratefully acknowledged.

REFERENCES

Black, D., and Sanders, S. (2000). When work disappears: A preliminary investigation into the relationship between job destruction and crime, Unpublished manuscript.

Blumstein, A., Cohen, J., Roth, J., and Visher, C. (eds.) (1986). *Criminal Careers and "Career Criminals*," National Academy of Sciences, Washington, DC.

Cameron, S., and Heckman, J. (1999). The economics and econometrics of active labor market programs. In Ashenfelter, O., and Card, D. (eds.), *Handbook of Labor Economics, Vol. 3*, Elsevier, Amsterdam, Chap. 31.

Cantor, D., and Land, K. (1985). Unemployment and crime rates in the post-World War II United States: A theoretical and empirical analysis. *Am. Sociol. Rev.* 50: 317–332.

Cantor, D., and Land, K. (2001). Unemployment and crime rate fluctuations: A comment on Greenberg, *J. Quant. Criminol.* 17: 329–342.

Chiricos, T. (1987). Rates of crime and unemployment: An analysis of aggregate research evidence. *Soc. Problems* 34: 187–211.

Cohen, L., and Land, K. (1987). Age structure and crime: Symmetry vs asymmetry and the projection of crime rates through the 1990s, *Am.. Sociol. Rev.* 52: 170–183.

Freeman, R. (1983). Crime and unemployment. In Wilson, J. Q. (ed.), *Crime and Public Policy*, ICS Press, San Francisco.

Freeman, R. (1995). The labor market. In Wilson, J. Q., and Petersilia, J. (eds.), *Crime*, ICS Press, San Francisco.

Goring, C. (1913). *The English Convict*, Patterson Smith, Montclair, NJ.

Greenberg, D. (2001). Time series analysis of crime rates, *J. Quant. Criminol.* 17: 291–327.

Grogger, J. (1995). The effect of arrests on the employment and earnings of young men. *Q. J. Econ.* 110: 51–71.

Gould, E., Weinberg, B., and Mustard, D. (1998). Crime rates and local labor market opportunities in the United States: 1979–1991. Working paper.

Kling, J. (1999). The effect of prison sentence length on the subsequent employment and earnings of criminal defendants. Woodrow Wilson School Economics Discussion Paper 208.

Krohn, M. (1978). A Durkheimian analysis of international crime rates. *Soc. Forces* 57: 654–670.

Land, K., McCall, P., and Cohen, L. (1990). Structural covariates of homicide rates: Are there any invariances across time and social space, *Am. J. Sociol.* 95: 922–963.

Lee, D. (1993). An empirical investigation of the economic incentives for criminal behavior. Unpublished manuscript.

Levitt, S. (1996). The effect of prison population size on crime rates: Evidence from prison overcrowding litigation. *Q. J. Econ.* 111: 319–351.

Levitt, S. (1997). Using electoral cycles in police hiring to estimate the effect of police on crime, *Amer. Econ. Rev.* 87(3): 270–290.

Levitt, S. (1998). Juvenile crime and punishment. *J. Polit. Econ.* 106: 1156–1185.

Levitt, S. (1999). The exaggerated role of changing age structure in explaining aggregate crime changes. *Criminology* 37: 581–597.

Lott, J. (1992). An attempt at measuring the total monetary penalty from drug convictions: The importance of an individual's reputation. *J. Legal Stud.* 21: 159–187.

Lynch, J. (1995). Crime in international perspective. In Wilson, J. Q., and Petersilia, J. (eds.), *Crime*, ICS Press, San Francisco.

Machin, S., and Meghir, C. (2000). Crime and economic incentives. Unpublished manuscript.

Marvell, T., and Moody, C. (1991). Age structure and crime rates: The conflicting evidence. *J. Quant. Criminol.* 7(3): 237–273.

Marvell, T., and Moody, C. (1994). Prison population growth and crime reduction, *J. Quant. Criminol.* 10: 109–140.

Marvell, T., and Moody, C. (1997). The impact of prison growth on homicide, *Homicide Stud.* 1: 205–233.

Mocan, N., and Rees, D. (1999). Economic conditions, deterrence, and juvenile crime: Evidence from micro data. Unpublished manuscript.

Neapolitan, J. (1997). *Cross-National Crime: A Research Review and Sourcebook*, Greenwood Press, Connecticut.

O'Brien, R. (1989). Relative cohort size, and age-specific crime rates: An age-period-relative-cohort-size model. *Criminology* 27(1): 57–77.

Quetelet, A. (1831). *Research on the Propensity for Crime at Different Ages* (Sylvester, S. F., 1984), Anderson, Cincinnati.

Raphael, S., and Winter-Ebmer, R. (2000). Identifying the effect of unemployment of crime. Unpublished manuscript.

Soares, R. (1999). Development, crime, and punishment: Accounting for the international differences in crime rates. Unpublished paper. Department of Economics, University of Chicago, Chicago.

Spelman, W. (1994). *Criminal Incapacitation*, Plenum, New York.

Stack, S. (1984). Income inequality and property crime: A cross-national analysis of relative deprivation theory. *Criminology* 22: 229–257.

Steffensmeier, D., and Harer, M. (1987). Is the crime rate really falling? An "aging" U.S. population and its impact on the nation's crime rate, 1980–1984. *J. Res. Crime Delinq.* 24: 23–48.

Steffensmeier, D., Streifel, C., and Harer, M. (1987). Relative cohort size and youth crime in the United States, 1953–1984. *Am. Sociol. Rev.* 52: 702–710.

Wilson, J. Q. (1983). *Thinking About Crime*, Basic Books, New York.

Wilson, J. Q., and Herrnstein, R. (1985). *Crime and Human Nature*, Simon and Schuster, New York.

Wright, R., and Decker, S. (1994). *Burglars on the Job*, Northeastern University Press, Boston.