Sociology Working Papers

Paper Number 2015-05

Size Matters: The Problems with Counting Protest Events

Michael Biggs

Department of Sociology and St Cross College, University of Oxford

August 2015

Department of Sociology

University of Oxford
Manor Road
Oxford OX1

Size Matters: The Problems with Counting Protest Events[*]

Michael Biggs

michael.biggs@sociology.ox.ac.uk
Department of Sociology, University of Oxford

August 2015

Since the 1970s, catalogs of protest events have been at the heart of research on social movements. Sociologists count the frequency of events to measure how protest changes over time or varies across space, as either the dependent variable or a key independent variable. This measure is disconnected from theory, which conceptualizes protest as collective action—by implication, what should be quantified are actions. Most fundamentally, counting events is inappropriate because it treats an event with ten participants as equal to an event with a million. This paper investigates three forms of protest: demonstrations, strikes, and riots. Their size distributions manifest enormous variation. Most events are small, but a few large events contribute the majority of protesters. When events are aggregated over years, there is no high correlation between event frequency and total participation. Therefore analyses of event frequency are not informative about the causes or consequences of participation *in protest. The fact that the bulk of participation comes from large events has positive implications for the compilation of event catalogs. Rather than fretting about the underreporting of small events, concentrate on recording large ones accurately.*

[*] This paper is possible because other scholars have generously shared their data: Gregg Carter; Ronald Francisco; Doug McAdam, John McCarthy, Susan Olzak, and Sarah Soule; Helen Margetts and Scott Hale. It uses software written by Nicholas J. Cox, Zurab Sajaia, and Yogesh Virkar. Kenneth T. Andrews and Charles Seguin judiciously critiqued an early draft, while Tak Wing Chan and Neil Ketchley supplied refreshment and encouragement. Andrew Gelman's blog persuaded me that measurement is worth taking seriously.

For the study of social movements, the crucial empirical innovation was the catalog of protest events. Such catalogs really originated in the late nineteenth century, when governments—in response to the emerging labor movement—began publishing statistics on strikes (Franzosi 1989). Other forms of protest, however, were not subject to official investigation. American social scientists began collecting their own data on events from the 1960s. In political science, cross-national time series of protest and violence were compiled from the *New York Times Index* (e.g. Rummel 1963; Hibbs 1973). In sociology, Tilly used newspapers to catalog "contentious gatherings" in Britain in the late eighteenth century and early nineteenth (Tilly 1978: Appendix 3; 1995). Tilly's work proved enormously influential. Event catalogs have been at the core of landmark studies of social movements (e.g. McAdam 1982; Kriesi, Koopmans, Dyvendak, and Giugni 1995; Olzak 1992; Tarrow 1989). Large-scale research projects have now compiled national datasets extending over several decades. The United States is covered from 1960 to 1995, using the *New York Times*. As data accumulates, analyses proliferate. Seven leading sociology journals since 2000 have published over forty papers that quantify protest, as either the dependent variable or a key independent variable. To capture how protest changes from year to year or how it varies across cities or states, the standard procedure is to count the frequency of events.[1]

Methodological scrutiny has focused on one problem: the vast majority of protest events are never reported by the news media, and so the frequency of protest is severely underestimated (Barranco and Wisler 1999; Earl, Martin, McCarthy, and Soule 2004; Franzosi 1987; Maney and Oliver 2001; McCarthy, McPhail, and Smith 1996; McCarthy, Titarenko, McPhail, Rafail, and Augustyn 2008; Myers and Caniglia 2004; Oliver and Maney 2000; Oliver and Myers 1999; Ortiz, Myers, Walls, and Diaz 2005). This lacuna leads Myers and collaborators to argue that catalogs compiled from newspapers—especially from a single newspaper—have "very serious flaws" (Ortiz, Myers, Walls, and Diaz 2005: 398; also Myers and Caniglia 2004). Such skepticism is rare. Most social scientists analyzing protest events concur with the conclusions of Earl, Martin, McCarthy, and Soule (2004: 77): "newspaper data does not deviate markedly from accepted standards of quality." Debate over sources of data—about reliability rather than validity—has displaced a more fundamental question. How should protest be quantified?

The frequency of events, I will argue, is an inappropriate variable. It contradicts the statistical property shared by most types of protest: events vary enormously in size. An event

[1] The article consistently refers to the *frequency* of events and the *number* of participants simply to help the reader discriminate between them.

may comprise a single person's action, or it may combine the actions of more than a million participants. Treating such disparate events as equal does not make numerical sense. Such enormous variation in size would not be so serious if it tended to even out when many events are aggregated into time intervals or geographical units. Then the frequency of protest events would correlate highly with the total number of protesters, and the choice between the two would hardly matter. In fact, I will demonstrate that the correlation is low or modest in several datasets, ranging from .10 to .64. Therefore findings from studies which count events do not warrant conclusions about participation in protest.

My critique leads to constructive proposals. Once the extreme variation in event size is appreciated, it is clear that the largest events—relatively few in number—contribute the majority of total participants. That most events go unreported by the media is far less of a problem when we measure total participation, because these small events contribute so little to the total. National series of strikes and demonstrations, for example, are dominated by events of at least 10,000 participants. Therefore we need to focus on the largest events. Effort should concentrate on recording these events properly, by eliminating erroneous duplication and estimating their size more consistently; this effort is feasible because there are so few of them.

This paper develops a quantitative and empirical critique of the use of event frequency as a variable. At the outset, it is worth noting that this variable lacks conceptual or theoretical justification. Conceptualizations of the object of sociological interest refer to actions rather than events. For Tarrow (2011: 7), "[t]he irreducible act that lies at the base of all social movements, protests, rebellions, riots, strike waves, and revolutions is *contentious collective action*." Opp (2009: 38) defines protest as "joint (i.e. collective) action of individuals aimed at achieving their goal or goals by influencing the decisions of a target." Social movements are conceived by Oliver and Myers (2003: 3) as "populations of collective actions." By implication, what should be quantified are actions. Actions create events; events are significant because of the actions they instantiate. The frequency of protest events is only obliquely related to contentious collective action. Action can be measured in different ways (as will be discussed further below). The most basic measure is the number of protest actions, in other words the number of participants in protest.

The number of participants is a crucial variable in theories of how social movements bring about change. DeNardo (1985: 36–37) takes "the disruptiveness of protests, demonstrations, and uprisings to be first and foremost a question of numbers"; thus one key parameter is "the percentage of the population mobilized" (see also Lohmann 1994).

According to Oberschall (1994: 80), "the crucial resource for obtaining the collective good is the number of participants or contributors." Practitioners echo this point. "Remember, in a nonviolent struggle, the only weapon that you're going to have is numbers" (Popovic and Miller 2015: 52). To my knowledge, no theory explicitly identifies the frequency of events as a critical variable. The importance of numbers is reinforced when we consider the orchestration of contentious gatherings. Tilly (1995: 370) postulates that leaders are "maximizing the multiple of four factors: numbers [of participants], commitment, worthiness, and unity." When demonstrators converge on one location (such as the capital city) on a single day, they are clearly maximizing the number of participants. If they wanted to maximize the frequency of events, then they would disperse to different places to perform separate demonstrations at different times. Indeed, the size of the demonstrations would not matter—ten demonstrations of a dozen people would be preferable to a single demonstration of one hundred thousand.

My argument is developed by analyzing several datasets. First is Dynamics of Collective Action in the United States, 1960–1995 (DCAUS), compiled from the *New York Times* (McAdam, McCarthy, Olzak, and Soule n.d.). This has been used in a score of articles. By contrast, the European Protest and Coercion Dataset (EPCD), spanning the years 1980–1995, is unduly neglected (Francisco n.d.). Unlike DCAUS, it is derived from multiple newspapers and newswires. The United Kingdom, as the most similar country to the United States, is selected for comparison. Both datasets encompass heterogeneous events, from prison riots to press conferences, from knee-capping to general strikes. Lumping these together makes little sense. The two datasets also cover a different range of events; routine strikes are included in EPCD but excluded from DCAUS. Therefore I will focus on demonstrations, including marches, rallies, and vigils. In DCAUS, demonstrations contribute half the total number of participants; in EPCD for the United Kingdom, a third of the total. These datasets are usefully compared to a comprehensive catalog of demonstrations in Washington, D.C. in 1982 and 1991, compiled primarily from the records of three police forces (McCarthy, McPhail, and Smith 1996). Although the authors have not made the data available, published tabulations are valuable for revealing the mass of small events that never make the news. Strikes are particularly valuable for my purpose because they are less prone to underreporting and because their size is measured more accurately. The United Kingdom consistently tabulated the size distribution from 1950 to 1984.[2] The United States recorded every strike from 1881

[2] In this period, the size classification pertains to strikes *begun* each year; from 1985, it switches to strikes *in progress* each year. In the latter scheme, some strikes spanned two calendar years and thus were counted twice.

to 1893, which permits analysis of a subset of events in small cities and towns. The final dataset is Carter's catalog of black riots in the United States from 1964 to 1971 (Carter 1986: 220–21). This allows investigation of variation across cities, to complement the time series for strikes and demonstrations.

The paper begins by reviewing the quantification of protest events in recent literature. The second section conceptualizes various ways of measuring the size of protest events, and outlines the challenge of measuring participation. Size distributions of demonstrations, strikes, and riots are presented in the third section. These distributions are heavy-tailed, meaning that most participants are concentrated in a few huge events. How does this matter? The fourth section compares event frequency and total participation, and shows that they are not strongly associated: conclusions drawn from one cannot be applied to the other. This negative finding is offset by positive implications in the fifth section. The problem of unreported small events diminishes; large events provide a remarkably accurate measure of total participation. The conclusion draws implications for future research.

1. Protest events in the literature

Fifty years ago, Tilly and Rule (1965) wrote at length on *Measuring Political Upheaval*. Their major precedent was strikes, which were quantified in three ways (elegantly explicated by Spielmans 1944): the frequency of events; the total number of participants, workers involved in strikes; and the total number of participant-days, working days lost in strikes. Tilly and Rule (1965: 74) concluded that "[t]he most useful general conception of the magnitude of a political disturbance seems to be the sum of human energy expended in it." This, they argued, was best approximated by participant-days. In the subsequent half century, data on events has accumulated, while consideration of what to quantify has disappeared. A review article from 1989 devotes a single page to this issue (Olzak 1989: 127), and it is not mentioned in a subsequent review (Earl, Martin, McCarthy, and Soule 2004).

How, then, is protest quantified in recent literature? Consider articles published from 2000 to 2014 in seven leading Anglophone journals: *American Journal of Sociology*, *American Sociological Review*, *British Journal of Sociology*, *European Sociological Review*, *Mobilization*, *Social Forces*, and *Social Problems* (following Amenta, Caren, Chiarello, and Su 2010).[3] Articles are selected if they measure protest as either the dependent variable or an

[3] Space limitations preclude consideration of a parallel literature in political science (e.g. Przeworski 2009) which relies on the frequency of strikes, riots, and demonstrations in the Banks' Cross-National Time-Series Data Archive (originating with Rummel 1963).

independent variable or both.[4] This excludes variables defined by the occurrence of protest rather than its magnitude—as for example, the dates at which cities experience rioting, used in event-history analysis.[5] Also excluded are studies that take the event as the unit of observation. Predominantly qualitative articles that define the explanans by graphing a time series are included (e.g. Biggs 2013). The literature search yields 41 articles (Table S1). Most articles construct a time series at the national level, which is usually annual but in a few cases quarterly or monthly. Some articles measure how protest varies across cities or other geographical units. A few combine time and space, so the unit of observation is the city-year for example. There are single instances of other combinations.

The quantification of protest requires two choices. The first is what sort of protest events to cover. Some articles focus on distinct types of protest, such as petitions or demonstrations or strikes or riots. Other articles lump together diverse types of action, from sit-ins to litigation, into a single variable (e.g. Jacobs and Kent 2007). Whether it is conceptually meaningful to aggregate such heterogeneous forms of action is outside the scope of this article. It does, however, have implications for the second choice.

Given a certain definition of events, the second choice is how to quantify them. Protest is almost universally measured by counting the frequency of events occurring in each observation. This variable is used in 83% of articles, and 66% use this measure alone—without alternatives. Four articles use the average number of participants per event (e.g. Soule and Earl 2005). Four articles use the total number of participants in protest events, for example workers involved in strikes (e.g. Checchi and Visser 2005). Two articles on riots combine participation, duration, and severity by using factor analysis to condense several characteristics—the number of people arrested, injured, and killed; the number of buildings set alight; and the duration in days—to a single score (e.g. Myers 2010).[6] Various other variables are confined to a single article.

These two methodological choices are fundamental. Yet most articles fail to justify them. There is methodological discussion, but it concerns reliability rather than validity. Thus an article will claim that the frequency of events reported in the *New York Times* reliably tracks the true frequency of events, while eliding the more fundamental questions: what kinds of

[4] Articles are identified by searching the title and abstract for one of the following keywords: protest(s), demonstrations, strikes, or riots.

[5] If event-history analysis also measures the magnitude of protest as an independent variable, then it qualifies for inclusion.

[6] Factor analysis has a long pedigree for riots (Morgan and Clark 1973; Carter 1986), but has not spread to other sorts of protest events.

protest should be included and how should they be measured? The majority of articles that measure only frequency fail to explain this choice. An exception is Budros' (2011: 444) analysis of petitions from the late eighteenth century, which notes the difficulty of obtaining the number of signatories for each petition. In some cases, data on size has presumably not been collected, as when the source is the *New York Times Index* rather than the original news articles (e.g. Jenkins, Jacobs, and Agnone 2003). The use of event frequency may follow from a heterogeneous definition of events: press conferences and boycotts, for example, seem to lack a common size metric. Conversely, all four articles that measure total participation focus on one type of event (strikes or demonstrations or suicide protest).

2. Conceptualizing and measuring size

Analyzing the frequency of protest events means ignoring their size. Size can be conceived in various ways. The first dimension of size is the number of participants. Some forms of protest exhibit little variation. The ultimate example is suicide protest, where the majority of events comprise a single participant. In a study of five hundred events, the largest involved twelve individuals: a group of monks and nuns in Vietnam in 1975 who killed themselves together (Biggs 2005a: 188). By contrast, participation varies appreciably for most types of protest, such as hunger strikes, demonstrations, and riots. Some events extend significantly over time; strikes can continue for months, and occupations for years. Duration is the second dimension of size. It creates the number of participant-days as a two-dimensional measure. In strike statistics, this is the number of working-days lost: the total number of days that every striker was out on strike. A rather different way of conceptualizing size is severity or disruptiveness. The severity of a riot, for example, can be measured by the number of fatalities or the number of properties destroyed. Severity is specific to the type of protest.

Whether measured by participants, participant-days, or severity, events are naturally aggregated over time and space to construct quantitative variables, such as the total number of demonstrators per month or the total number of riot fatalities per year. Note that the total number of participants is not identical to the total number of individual people who protested in the period, because some protested more than once. Because we are interested in contentious collective action, it is right to count *actions*: a worker who goes on strike twice in the course of a year properly contributes two actions to the annual total. One point to emphasize is that aggregation entails taking the sum and not computing the average. The literature sometimes takes the average number of protesters per event as a measure of protest, but this does not serve to quantify participation. (The same objection applies to average duration.) A simple numerical example clarifies this point. One city has two demonstrations, with 100 and 100,000 participants. Another city has three demonstrations, two of 100 participants and one of 100,000. Measuring average participation implies that there is 50% more protest in the former city than in the latter. In fact, of course, there is marginally more protest in the latter city. If protesters wanted to maximize average rather than total participation, then they would never hold small events.

Aggregate measures are often used to trace change over a long period, when population grew appreciably, or to compare across spatial units varying in population. Total participants and total participant-days are naturally divided by the population of potential protesters, to create proportional measures of propensity and intensity respectively. Propensity neatly corresponds to the individual's continuous-time hazard rate of protesting. It is common usage to also use population as the denominator for event frequency (e.g. Rosenfeld 2006); or, equivalently, to specify this as the exposure term in Poisson or negative binomial regression (e.g. Inclán 2009). Events per person, however, does not seem a meaningful metric. It is also not mathematically justified. For a given propensity to protest, double the population means double the total number of participants (and hence double the total number of participant-days). But we would not expect twice as many events, because that would implausibly imply no change in the average number of participants per event. Presumably event frequency and average size would each increase by a factor of between 1 and 2, with their product equal to 2.

This paper focuses on the first and most basic dimension of size: the number of participants. The two-dimensional measure of size, participant-days, is pertinent for strikes but is not relevant for demonstrations and is not empirically measurable for riots. Therefore it is not considered further here. Also omitted are measures of severity. My argument for the importance of size also implies the need to compare different measures of size, but there is no space to undertake this here. This paper also ignores the problem, noted above, of aggregating disparate types of protest. Conceptually this would seem to require measuring the cost (or impact) of different types of protest actions, arraying them on a spectrum from signing an email petition to setting oneself on fire. This problem will be avoided here by analyzing different types of protest separately.

Measuring the number of participants is often challenging. At the reliable end of the spectrum are strikes. Data are compiled by specialized officials, who are not involved in the dispute. They can use the records of firms, and of trade unions if they provide strike pay. Because the point of a strike is to inflict costs on the employer, neither side has reason to exaggerate or minimize the number of strikers, at least not in statistics published long after the event is over. At the opposite extreme of reliability are riots. Aside from the question of how exactly to define what counts as participation, the nature of a riot—fluid, dispersed, and furtive—hinders numerical estimation. The number of arrests provides a rough proxy for participation, though this depends not just on the number of rioters but also the tactics and capabilities of the police.

Estimating the number of demonstrators falls somewhere between riots and strikes. A proper estimate can be calculated from the dimensions of the gathering place and the physical density of demonstrators (McPhail and McCarthy 2004).[7] Such calculations began to be produced by the U.S. Park Police in Washington, D.C. in the mid 1970s. More usually, however, size has to be taken from the guesstimates of police or reporters—who are usually unsympathetic—or of organizers—who naturally exaggerate. Media sources can select which of these estimates to report, in accordance with their sympathy or antipathy to the protesters' cause (Mann 1974).[8] A notorious example of conflicting estimates was the "Million Man March" in 1995. The organizers naturally claimed that it lived up to its name, while the Park Police calculated 400,000. The ensuing controversy led Congress to forbid Federal police forces from estimating crowd size.

Conflicting estimates of size may seem to pose an insuperable problem. It is attenuated, however, if we conceive size multiplicatively rather than additively, on a logarithmic rather than linear scale. This point was originally made by Richardson (1948: 523) for the size of wars measured by fatalities; these too are estimated with considerable error (and also have a heavy-tailed distribution, to anticipate the next section). What really matters is the order of magnitude, the difference between one power of ten and the next. To return to the example of the Million Man March, the difference in estimates is considerable: the Park Police's figure is 60% less. Taking the logarithm to the base ten, the competing estimates are 5.6 and 6.0, on a scale beginning at 0 (for a single demonstrator). Thus transformed, the difference shrinks. To put this another way, even the lower estimate makes the Million Man March one of the very largest demonstrations in the United States. Thinking of size in terms of orders of magnitude is intuitive, for commentators often describe size in this way—referring to "hundreds" or "thousands" of participants, for example.

---

[7] By the same principle, participants in a march can be estimated by taking one fixed observation point and sampling the number passing per minute; adding a second point enables the estimation of error bounds (Yip et al. 2010).

[8] McCarthy, McPhail, Smith, and Crishock (1998: 123) compare the number of participants reported by newspapers with the number anticipated by the organizers when applying for a permit, and find a correlation of .71. Ideally we would discover the correlation with the actual number of participants.

3. Size distributions

We can now investigate the number of participants in demonstrations, strikes, and riots. Table 1 presents summary statistics for five datasets. A crude way of gauging variation is to compare the maximum to the median. Another is to divide the standard deviation by the mean, which yields the coefficient of variation. Following mean and variance, the third and fourth moments of the distribution are skewness and kurtosis. The latter is significant because it indicates the heaviness of the tail of the distribution. Table 1 presents *L*-kurtosis, one of the *L*-moments which are derived from order statistics (Hosking 1990). Unlike the conventional measure of kurtosis, this does not increase with sample size and it is less sensitive to extreme values. *L*-kurtosis can range from -.25 to 1; for a normal distribution it is .12, and for an exponential distribution it is .17. (The Gini index will be discussed later.) A heavy tail may be defined as one that is heavier than the exponential distribution, meaning that there is a nontrivial probability of extremely large values. Heavy tails are less familiar to our statistical intuition, nurtured on the normal distribution—the name reveals its hegemony—and developed by experience with thin-tailed distributions such as age and years of education.[9]

[Table 1 about here]

For demonstrations, a single estimate of size is reported in the majority of events in the United States and the United Kingdom (63% in DCAUS and 71% in EPCD). The remainder are described approximately, as by orders of magnitude. EPCD employs the sensible convention of coding "hundreds" as 300, "thousands" as 3,000, and so on.[10] I apply this to DCAUS.[11] In both countries, the typical or median size was in the hundreds, while the maximum was in the hundreds of thousands. These distributions underestimate the degree of variation, because the smallest events were much less likely to be reported. It is therefore valuable to compare the exceptionally comprehensive record of demonstrations in Washington, D.C. (McCarthy, McPhail, and Smith 1996: 484, 488). Here the number of participants is usually the number anticipated by organizers when applying for a permit to demonstrate; sometimes it is the number as subsequently amended by the police. The median size was two dozen. The largest demonstrations exceeded the median by over four orders of magnitude.

---

9 Empirical variables enountered in introductory and intermediate statistical courses are invariably thin-tailed. Personal income would be a potential exception, but survey data given to students (like the U.S. General Social Survey) truncate its heavy tail. Resnick (2010) provides a useful primer on heavy-tailed distributions.

10 Koopmans (1995: 261) codes "hundreds" as 500, "thousands" as 5,000, and so on. This is too generous.

11 In the smaller categories, I take the midpoint: thus 5 for 2–9, 30 for 10–49, and 75 for 50–99. I exclude 15 demonstrations in the U.S. dataset that lack information on the number of participants.

Data on strikes in the United Kingdom cover those lasting for at least a day and involving at least ten workers, along with small or brief strikes if at least a hundred working days were lost. Size is measured by the number of workers involved.[12] Published tabulations use broad size intervals, but fortunately it is possible to identify each individual strike reaching 10,000 workers (see Appendix S1). The median was under a hundred, which is comparable to demonstrations when they are measured comprehensively. The maximum was larger, because a strike does not require the physical assembly of participants in a single place.

For riots, size is proxied by the number of arrests. (Events where no one was arrested are omitted.) The median number of arrests was 17. The maximum—a riot in Washington, D.C. in April 1968—was greater by over two orders of magnitude. From retrospective surveys in three cities, it is estimated that between one in six and one in three rioters were arrested (Fogelson 1971: 36–37). By implication, participation in the largest riots was an order of magnitude smaller than participation in the largest demonstrations. This makes sense given that a major demonstration in a capital city always attracts people from elsewhere, whereas the vast majority of rioters are local. Because a riot is geographically circumscribed, we can also consider the number of arrests in relation to the population of potential rioters, in this case the number of nonwhites in the city.[13] The median was 0.1%. The maximum was greater by an order of magnitude, at 4%. Dividing by population significantly reduces variance (as manifested in the coefficient of variation) and kurtosis. Note that the data on riots comprise fewer events than the other datasets. In principle, of course, the maximum will increase with the number of observations, and markedly so for a heavy-tailed distribution. In practice, though, it is hard to conceive of a much greater maximum number of arrests in one city.

All these distributions exhibit enormous variation in size. The standard deviation always exceeded the mean. The least variation occurred in the distribution of riot arrests relative to population; even then the maximum was forty times the median. The ratio of maximum to median exceeded 20,000 for participants in demonstrations—when not filtered through news media—and in strikes. The typical size is profoundly misleading.

---

[12] There is a distinction between workers directly involved and indirectly involved. The latter are involuntarily thrown out of employment by a strike in their firm. If the explanandum is protest, then we ideally should consider only workers directly involved, for they have chosen to participate. The tabulation of U.K. size distributions, however, does not separate workers indirectly involved. These contributed only a tiny fraction of the total, 0.15%.

[13] Cities are defined as urban places in the 1960 Census (Haines and ICPSR 2010: DS63; U.S. Census 1960: table 21). It is not possible to match 38 riots (6.4%), most commonly because the place was too small for the Census to tabulate population by race.

A heavy tail implies "mass-count disparity" (Crovella 2001): most of the total size comes from a small number of huge events. The literature on protest occasionally notes this feature (Franzosi 1989: 352; Koopmans 1995: 251; Rucht and Neidhardt 1998: 76–77), but its importance has not been fully appreciated. Figure 1 illustrates the disparity by comparing two cumulative distributions (after Feitelson 2006). One is the distribution of events by size. The other, shifted to the right, is the distribution of participants by size of event. The horizontal scale must be logarithmic, of course, to encompass the extreme range. Discontinuities in the graphs for demonstrations in the United States and the United Kingdom are due to size ranges (like "thousands") being translated into a number.

[Figure 1 about here]

Comparison of the three graphs for demonstrations reveals the gap left by underreporting. In D.C., demonstrations with no more than 25 participants accounted for over half the events. When events were filtered through the news media, the total contained far fewer of these tiny events: 5% in the United Kingdom, 7% in the United States. Paradoxically, however, the distribution of participation shows that when small events were comprehensively recorded—demonstrations in D.C. and strikes—they still contributed almost nothing to the total. Total participation was dominated by the largest events, which constituted a tiny fraction of all events. In D.C., only 1% of demonstrations had more than 10,000 participants, but they accounted for 69% of the total demonstrators. For strikes, only 0.4% involved as many as 10,000 workers, but they accounted for 56% of the total workers involved. In only 1.7% of riots did arrests reach 1,000, but they accounted for just over half of the total arrests. For riots, as before, we can adjust for population. One in ten riots led to the arrest of at least 1% of the nonwhite population, and these riots accounted for half of the total percentage of nonwhites arrested.

Mass-count disparity is visualized as the gap between the two cumulative distributions. It can be captured in a familiar statistic, the Gini index of inequality, shown in Table 1.[14] For demonstrations, the index ranged from .88 to .94; it was highest for D.C. because more small events were included.[15] For strikes the index was .87. These figures indicate extreme inequality: for comparison, the distribution of wealth in the contemporary United States has a Gini index of about .8. The index for arrests was .84. When arrests are denominated by

[14] The Gini index is also the *L*-moment equivalent of the coefficient of variation: mean divided by *L*-scale.

[15] The index for D.C. is calculated from binned data which requires interpolation, and so is less reliable than the others.

population, the index was .69. The latter is still higher than the index for the distribution of income (before taxes and transfers) in the contemporary United States, about .5.

Thus far I have referred generically to heavy-tailed distributions without considering the shape of the tail. The archetypal heavy tail is the power law, where the probability of an event of size x is proportional to $x^{-\alpha}$. In a pioneering study, Richardson (1948) argued that the size of wars, measured by fatalities, followed a power law. The same distribution has been identified in terrorist and insurgent attacks (Clauset, Young, and Gleditsch 2007; Bohorquez et al. 2009). According to Biggs (2005b), strikes in two cities in the late 19th century followed a power law, with α estimated as 1.9–2.0 (for $x \geq$ 100–150). Discriminating a power law from other heavy-tailed distributions—such as the lognormal and the power law with exponential cutoff—is empirically demanding: only a small fraction of events comprise the tail, and so the total number of observations must be very large (Clauset, Shalizi, and Newman 2009). Therefore it makes sense to concentrate on strikes. Figure 2 plots the complementary cumulative distribution, with both axes on a logarithmic scale. (Figure S1 depicts the other datasets.) Following the method of Virkar and Clauset (2014), the power law with the best fit has α of 2.2, starting at 1,000–2,499 workers. It appears on the graph as a straight diagonal line. Clearly this power law does not describe the very upper tail of the distribution; it predicts too few huge strikes. But alternative heavy-tailed distributions (fit to the same tail, starting at 1,000–2,499) are inferior.[16] The graph also serves to illustrates the significance of a heavy tail. Recall that heavy means heavier than the exponential distribution. The graph shows the best-fitting exponential distribution, which resembles an inverted J. Such a distribution would predict many more medium-sized strikes and fewer large ones; a strike involving more than a hundred thousand workers would be vanishingly rare. A heavy tail, by contrast, reflects the fact that huge events—more than four orders of magnitude greater than the median, in this case—can occur.

[Figure 2 about here]

Such huge events occur in datasets that cover a population of many millions over decades; these datasets are the staple of sociological analysis. Does the size distribution differ for small populations? Answering this question requires comprehensive data on small events in minor places, which rules out the media as a source. The best candidate is strikes in the United States between 1881 and 1893 (see Appendix S2). At that time almost all employers

[16] The hypothesis that strikes follow a power law is rejected with $p = .003$. But the power law is superior ($p < .001$) to the power law with exponential cutoff, the lognormal, the stretched exponential, and the exponential distribution.

were confined to a single location, with the major exception of railroads. Table 2 shows the size distribution for strikes in Illinois outside the industrial metropolis of Chicago (Cook county, to be precise). A few large railroad strikes are omitted because the Commissioner did not specify all their locations; the period ends before the massive Pullman railroad strike in 1894. The size distribution exhibits much less variation than the large national datasets. The coefficient of variation is 2.8, compared to 27 for strikes in the United Kingdom. Mass-count disparity is still significant: only 5% of strikes involved at least 1,000 workers, but these accounted for almost half (47%) the total number of participants. Although the largest strike fell short of 10,000 workers, note that the number of events is relatively small. Sampling the same number of events from strikes in the United Kingdom, we would expect only two (0.4%) to reach 10,000. Focusing on a particular location further reduces variation in the size of events. Table 2 shows strikes in Peoria, the state's second city, which had a working-class population of about 8,000. The largest strike involved 600 workers. With so few events, the upper tail of the distribution cannot be estimated; a much larger strike could have been possible. Even so, the Gini index shows greater inequality than is found in the distribution of income.

[Table 2 about here]

4. Event frequency and total participation

Given the enormous variation in the size of events like demonstrations and strikes, they cannot be treated as equal. Yet sociologists usually choose to count the frequency of events in each year or in each city. My argument is that the variable of interest should instead be *total participation* in those events. This incidentally avoids another problem of counting events which has not been acknowledged in the literature.[17] How is one event to be demarcated from another? In principle, this should be straightforward when dealing with a contentious gathering which is characterized by continuity and contiguity of action, like a march. In practice, however, event catalogs often treat multiple gatherings in different locations as one event. Thus vigils and processions in 21 cities on "National Free Sharon Kowalski Day" (to support a disabled woman whose lesbian partner was denied access by her family) become a single event in DCAUS. Why not count 21 events? EPCD likewise classifies the annual Loyalist Orange parades throughout Northern Ireland on July 12 as a single event (sometimes entering a particularly large or contentious parade as a second event). Why not consider the parades in each town or city as separate events? These coding decisions apparently reflect the

[17] An exception is Shorter and Tilly (1974: 353), but this is buried in an appendix.

level of detail provided by news reports. The problem of demarcation vanishes when we measure total participation in a time interval. When a newspaper reported "scores of Orange demonstrations in which an estimated 100,000 people took part" (*The Times* of London, July 13, 1990), whether we treat this as one event or scores makes no difference—either way, total participation increases by 100,000.

Whatever the theoretical, conceptual, and empirical reasons for preferring total participation over event frequency, it is still necessary to ask whether it actually makes a significant difference in practice. One might expect a high correlation between frequency and participation especially in an annual national time series. After all, there are many protest events in each year: on average, for example, over 200 demonstrations in the United States, and over 2,000 strikes in the United Kingdom. When so many events are aggregated, we might suspect that size differences would tend to average out.

Figure 3 traces both time series for demonstrations in the United States. Event frequency peaks in the mid 1960s, declines to the mid 1970s, and then continues at a low level. Total participation spikes in 1969 and peaks in 1982. The two variables follow a completely different trajectory. The scatterplots in Figure 4 portray the association between event frequency and total participation; the linear regression line is dashed. For demonstrations in the United States, the correlation coefficient is only .10. The correlation is somewhat higher for demonstrations in the United Kingdom, albeit for only 16 years. An estimated correlation coefficient automatically increases with fewer observations; in the United States, the correlation for each 16-year subperiod (1960–1975 to 1980–1995) averages .19. For strikes in the United Kingdom the correlation is low. For strikes in Illinois outside Chicago (not graphed), the correlation is .46, again for a much shorter period. For these time series, the population of potential protesters did not increase sufficiently to require adjustment.[18]

[Figure 3 and Figure 4 about here]

Riots are mapped on to the 460 urban places with a population of at least 25,000 and a nonwhite population of at least 1,000 in 1960. Few events are distributed across many observations, whereas the time series have many events distributed across few observations. The maximum number of riots was 14, in Washington, D.C. Half of the cities experienced no riots, and so the regression line (in Figure 4) is anchored at the bottom left-hand corner. The correlation between total arrests and riot frequency is only modest. But it is really necessary to adjust for the tremendous variation in the size of cities. Total participation obviously scales

[18] The U.S. population grew by 45%, the U.K. employed labor force by 1%, the Illinois non-agricultural occupied population (Cook county cannot be separated) by 75%.

with the population at risk. As argued above, event frequency must scale to population raised to a fractional power; the square root is used here.[19] Thus denominated, riot frequency and total arrests reach the highest correlation, .64. This means that frequency predicts only 41% of the variation in per capita arrests.[20]

In sum, when events like riots and demonstrations are aggregated over time or space, the frequency of events is only minimally or modestly associated with total participation in those events. The divergence between event frequency and total participation partly reflects the heavy-tailed size distribution of events: the occurrence of a huge event significantly increases total participation while only incrementing event frequency. Furthermore, the four annual time series reveal negative correlations between event frequency and average event size; in years with many events, events tended to be smaller. Whether this is coincidental or reflects a more general pattern must await research on other datasets.

These findings do not prove that total participation always diverges from event frequency. For events with minimal variation in size (exemplified by suicide protest), total participation will be practically the same as event frequency. For the staple tactics of social movements, however, there is no justification to assume a high correlation over time or across spatial units. Similar disassociation is evident, for example, for demonstrations in Belarus in the 1990s, aggregated by quarter (Titarenko, McCarthy, McPhail, and Augustyn 2001: 137). It follows that the findings from multivariate analysis using event frequency—as either the dependent variable or an independent variable—cannot be assumed to apply to total participation. Those analyses therefore do not license conclusions such as "X increases protest" or "protest has no effect on Y"—instead, they show only that "X increases protest frequency" or "protest frequency has no effect on Y." Whether "X increases protest participation" or "protest participation has no effect on Y" will be known only when total participation is analyzed.

[19] Recall that the number of events and the average number of participants per event must each scale to a fractional power of population, with the two fractions adding to one—in order for the total number of participants to scale with population. There is no reason to suppose that either event frequency or average participation increases faster than the other. The most parsimonious assumption is that each scales with population raised to the power of .5.

[20] Some analyses use the logarithm of event frequency (e.g. Jenkins, Jacobs, and Agnone 2003. Its correlation with the logarithm of total participation is barely higher: the correlation coefficients are .23 for U.S. demonstrations, .32 for U.K. demonstrations, .29 for U.K. strikes, and .48 for Illinois strikes. The same calculation cannot be performed for riots because half the cities experienced no riot and zero has no logarithm.

5. The dominance of large events

My argument thus far has been negative. Fortunately an equally important positive lesson can be drawn. As we know, most events are not reported by the news media and so are omitted from event catalogs. In Washington, D.C., newspapers reported only one in ten demonstrations. Worse, the extent of this underreporting will fluctuate over time—depending, for example, on the significance of other news—in ways that cannot be ascertained. Hence the argument that newspapers are a seriously flawed source of data (Ortiz, Myers, Walls, and Diaz 2005; Myers and Caniglia 2004). They are indeed flawed if one wishes to count the frequency of events. But the problem evaporates if one wants to trace total participation over time. After all, the probability of an event being reported increases with the number of participants. Newspapers reported only 3% of the demonstrations involving 2–25 participants in D.C., but they reported 52% of the demonstrations with 10,001–100,000, and both the demonstrations exceeding 100,000 (McCarthy, McPhail, and Smith 1996: 488).

Moreover, mass-count disparity has a surprising implication: when aggregated over time or space, total participation can be predicted by looking only at the large events. What counts as "large" depends on the actual size distribution, of course. We can use the threshold of 10,000 for strikers and demonstrators in the national datasets, and 1,000 for strikers in Illinois and for arrests. Table 3 reports the correlation coefficients. Aggregated by year, the correlations are almost perfect. The result for strikes is especially compelling, because these data do not omit small events. Figure 5 details the total number of workers involved in U.K. strikes reaching various size thresholds. Remarkably, to trace how the total number of strikers fluctuated from year to year, it is possible to ignore 99.6% of events; just track the total number of workers involved in strikes of at least 10,000. (This understates the *absolute* level of participation, of course, but it accurately captures relative change.) Increasing the threshold to the next order of magnitude hardly alters the graph. Even confining attention to strikes involving at least a million workers—only nine events!—captures the salient peaks.

[Table 3 and Figure 5 about here]

Given the dominance of large events, it may be tempting to count their frequency rather than to sum their participants. McAdam and Su (2002), for example, count the frequency of events exceeding 10,000 participants. Unfortunately, a heavy tail implies that large events also manifest mass-count disparity: most are moderately large, while a few are massive. Table 3 compares the correlation between total participation and the frequency of large events. The correlation is consistently lower, and it is greatly inferior for U.K. demonstrations.

The problem of underreporting, which has preoccupied the literature on protest events, shrinks when mass-count disparity is understood. For analysis of protest over time at the national level, the problem has been greatly overstated. Total participation is dominated by large events—and large events are most likely to be reported. Since we should measure annual fluctuations in the total number of protesters, we can focus on the number of protesters in the large events. This reassurance does not hold for the cross-sectional analysis of protesters denominated by population, because a national newspaper may not bother reporting a *relatively* large event in a small city.

Mass-count disparity does, however, highlight another issue. Because large events are so important, errors affecting them have severe repercussions. Both DCAUS and EPCD erroneously duplicate some demonstrations with hundreds of thousands of participants. One duplicate, for example, increases the year's total participants by 55%! Coding errors are inevitable in assembling a large dataset from news reports. Happily, mass-count disparity means that validation can focus on large events, which constitute only a tiny fraction of all events. In DCAUS, for example, only 4% of demonstrations reached 10,000 participants. Aside from errors, uncertainty in estimating the size of large events is also a significant concern. When estimates are wildly discrepant—if, for example, the higher is more than double the lower—the decision on which size to use will make a difference. When collecting data, therefore, it will be crucial to record varying estimates and their provenance. Estimates derived from spatial calculation (area occupied and density of demonstrators) are obviously preferable. Even without the benefit of such calculation, one could consistently select estimates with the same direction of bias, for example using the optimistic figures provided by the organizers.

Conclusion

The study of social movements requires the quantification of protest. What explains protest? What does protest explain? Whether we treat protest as an effect and seek its causes, or treat protest as a cause and seek its effects, we need to differentiate less protest from more. This basic distinction between less and more is required even for explanations that are pursued with qualitative evidence rather than statistical analysis; an historical narrative typically graphs the time series of protest. The pioneers of protest event analysis devoted considerable effort to thinking about what should be measured and how (e.g. Tilly and Rule 1965). As data accumulated, thanks to their efforts, these deeper methodological questions faded from view. The ready availability of event frequency became its own justification.[21]

My argument against event frequency has several components. Sociologists agree that the object of investigation is collective action. Therefore we should measure some property of actions—such as, most basically, their number. The total number of participants is specified explicitly in several theories of social movements, whereas event frequency is not. In practice, movements act as if the number of participants matters. Conceptually, it makes no sense to count the protest events because they are not equal in size. A strike or demonstration involving ten participants is not equal to one involving half a million. Yet the standard variable embodies an extraordinary assumption: two events represent twice as much as protest as one event, even if the two events are each attended by only ten participants whereas the single event attracts a million people. That assumption underpins the findings of most quantitative work on protest.

Empirical findings highlight the perils of ignoring size. Measured by participants, the size distribution of events—like demonstrations, strikes, and riots—does not range modestly around a typical value. Such protest events have a heavy tail: most events are small, but a few are massive. Even when we confine attention to events in one small city (like strikes in Peoria) or divide participants by population (as with riots), there is pronounced inequality in size. How far does this generalization hold? It does not apply to every form of protest; an exception is suicide protest. It may not apply in some social contexts.[22] But the generalization does hold for the phenomena that are the staple of sociological analyses of protest events.[23] For the purpose of analysis, sociologists usually aggregate protest events over time intervals (like years) or across geographical units (like cities). Under aggregation, there is no high correlation between event frequency and total participation. Four time series yield correlation coefficients from .10 to .46; with city as the unit of observation, the coefficient does not exceed .64. Perhaps other datasets will reveal higher correlations, but this will need to be

---

[21] The related enterprise of quantifying protest through population surveys followed a similar trajectory. The pioneers carefully formulated survey questions for particular theoretical purposes; the proliferation of survey data subsequently came to define the object of investigation (Biggs 2015).

[22] One candidate could be demonstrations under authoritarian regimes, where only a few dissidents are willing to risk participating. But such regimes also generate, on rare occasions, truly massive gatherings, as witnessed in East Germany in 1989 and Egypt in 2011.

[23] Yet another example is provided by petitions hosted by official websites in the U.K. and U.S. (Margetts et al. 2015: ch. 3). Out of 32,873 petitions submitted to the U.K. Parliament from 2011 to 2015, only 0.77% reached 10,000 signatures, but these accounted for 70% of the total signatures (calculated from data kindly supplied by the authors).

demonstrated. As it stands, there is no justification for treating the frequency of events as a tolerable approximation for the total number of participants. In short, then, analyses which take the number of protest events as dependent or independent variable cannot contribute to the understanding of participation in protest—the phenomenon that is most theoretically pertinent.

The problems identified by this paper point to a solution. With a heavy-tailed distribution, the weight of the distribution is concentrated at the top. Most protesters participate in large events. While the literature frets about the underreporting of small events, this matters far less if we wish—as theory suggests—to measure participation. What matters most are the largest events. The national series of demonstrations and strikes analyzed here are dominated by events involving at least 10,000 participants. It is important to emphasize that "large" is relative rather than absolute, and so it varies with the population of potential protesters. In Peoria in the 1880s, for example, a large strike was one that involved hundreds of workers. When aggregated over time or space, the total number of participants is very highly correlated with the total number in the largest events. By implication, attention should focus on the accurate recording of these events. One issue is the estimation of size. Varying estimates of the same event establish lower and upper bounds, which can be used to check the sensitivity of empirical findings. For contemporary events, the huge volume of photographic and video evidence can surely be exploited to refine estimates of crowd size.[24] Another issue is the validation of data. The duplication of large events in datasets—identified here in both DCAUS and EPCD—can severely distort empirical results. Fortunately, of course, there are only a small number of large events, and so it is feasible to check them thoroughly.

If my argument fails to convince, I hope that it will provoke the proponents of event frequency to justify their method. Do not assume that event frequency is the most natural—let alone the only—variable. Specify a theory that makes the frequency of events, irrespective of size, crucial. If the variable is constructed from news sources that report only a tiny fraction of events, then give grounds for choosing a variable that omits most of what it ostensibly measures. (Total participation, by contrast, is less affected by the omission of small events.) Explicate the criteria used to demarcate one event from another. (Total participation, by contrast, does not depend on demarcation.) If adjusting for population size, explain the choice of scaling parameter; recall that event frequency cannot increase in proportion to population.

[24] This will be facilitated by imagery from aerial unmanned vehicles (Choi-Fitzpatrick, Juskauskas, and Sabur 2015).

My argument has two further implications for future research. I have focused on one measure of size: the number of participants. Arguably this variable is most appropriate when explaining the origins of protest, as it measures the number of individual decisions to participate. When using protest as an explanation for subsequent outcomes, however, participant-days—incorporating the second dimension, duration—could be preferable. There are alternative measures of size, such as severity (e.g. Carter 1986), which also deserve further examination. The point of my argument is not to impose one single variable for all theoretical purposes, but to encourage the investigation of size.

The size distribution of protest events has implications beyond method. A significant task for theory is to explain why most protest events are small while a few are huge. One simple explanation is that this distribution reflects the division between populations of potential protesters. The population of cities, for example, follows a power law with α of 2 (Rozenfeld, Rybski, Gabaix, and Makse 2011). Thus the distribution of riot arrests is partly a function of the distribution of nonwhites; the heavy tail is significantly diminished when arrests are divided by population. Evaluating this explanation for other types of protest is more complicated, because it is hard to conceive of pre-existing population divisions for strikes—though the distribution of workers by industry or occupation could be considered—or for demonstrations. A second explanation is that the size distribution reveals something about the *process* generating protest events. Biggs (2005b) argues for positive feedback: for an event in progress, the larger it becomes, the more likely people are to join it. This kind of process—synonymous with cumulative advantage or preferential attachment—is often used to explain heavy-tailed size distributions (e.g. Seguin in press). Other models of the generative process are also possible. For violent attacks, a model of the fusion and fission of insurgent groups predicts the distribution of fatalities (Johnson et al. 2006; Clauset and Weigel 2010). Modeling the size of protest events is a promising avenue for future research.

References

Amenta, Edwin, Neal Caren, Elizabeth Chiarello, and Yang Su. 2010. "The Political Consequences of Social Movements." *Annual Review of Sociology* 36:287–307.

Barranco, José and Wisler, Dominique. 1999. "Validity and Systematicity of Newspaper Data in Event Analysis." *European Sociological Review* 15(3):301–22.

Biggs, Michael. 2005a. "Dying without Killing: Self-Immolations, 1963-2002." In Diego Gambetta (ed.), *Making Sense of Suicide Missions*. Oxford University Press.

Biggs, Michael. 2005b. "Strikes as Forest Fires: Chicago and Paris in the Late 19th Century." *American Journal of Sociology* 110(6):1684–14.

Biggs, Michael. 2013. "How Repertoires Evolve: The Diffusion of Suicide Protest in the Twentieth Century." *Mobilization* 18(4):407–28.

Biggs, Michael. 2015. "Has Protest Increased Since the 1970s? How a Survey Question Can Construct a Spurious Trend." *British Journal of Sociology* 66(1):141–62.

Bohorquez, Juan Camilo, Sean Gourley, Alexander R. Dixon, Michael Spagat, and Neil F. Johnson. 2009. "Common Ecology Quantifies Human Insurgency." *Nature* 461:911–14.

Budros, Art. 2011. "Explaining the First Emancipation: Social Movements and Abolition in the U.S. North, 1776–1804." *Mobilization* 16(4):439–54.

Carter, Gregg Lee. 1986. "The 1960s Black Riots Revisited: City Level Explanations of Their Severity." *Sociological Inquiry* 56(2):210–28.

Checchi, Daniele and Visser, Jelle. 2005. "Pattern Persistence in European Trade Union Density: A Longitudinal Analysis 1950–1996." *European Sociological Review* 21(1):1–21.

Choi-Fitzpatrick, Austin and Juskauskas, Tautvydas. 2015. "Up in the Air: Applying the Jacobs Crowd Formula to Drone Imagery." *Procedia Engineering* 107: 273–81.

Clauset, Aaron, Cosma Rohilla Shalizi, and M.E.J. Newman. 2009. "Power-Law Distributions in Empirical Data." *SIAM Review* 51(4):661–703.

Clauset, Aaron and Wiegel, Frederik W. 2010. "A Generalized Aggregation-Disintegration Model for the Frequency of Severe Terrorist Attacks." *Journal of Conflict Resolution* 54(1):179–97.

Clauset, Aaron, Maxwell Young, and Kristian Skrede Gleditsch. 2007. "On the Frequency of Severe Terrorist Events." *Journal of Conflict Resolution* 51(1):1–31.

Crovella, Mark E. 2001. "Performance Evaluation with Heavy Tailed Distributions." In Dror G. Feitelson and Larry Rudolph (eds), *Job Scheduling Strategies for Parallel Processing: 7th International Workshop*. Berlin and London: Springer Verlag.

DeNardo, James. 1985. *Power in Numbers: The Political Strategy of Protest and Rebellion*. Princeton University Press.

Earl, Jennifer, Andrew Martin, John D. McCarthy, and Sarah A. Soule. 2004. "The Use of Newspaper Data in the Study of Collective Action." *Annual Review of Sociology* 30:65–80.

Feitelson, Dror G. 2006. "Metrics for Mass-Count Disparity." In *Proceedings of the 14th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS '06)*.

Fogelson, Robert M. 1971. *Violence as Protest: A Study of Riots and Ghettos*. New York: Doubleday.

Francisco, Ronald A. n.d. European Protest and Coercion Data, 1980–1995 (computer file). http://web.ku.edu/~ronfrand/data/index.html.

Franzosi, Roberto. 1987. “The Press as a Source of Socio-Historical Data: Issues in the Methodology of Data Collection from Newspapers.” *Historical Methods* 20(1):5–15.

Franzosi, Roberto. 1989. “One Hundred Years of Strike Statistics: Methodological and Theoretical Issues in Quantitative Strike Research.” *Industrial and Labor Relations Review* 42(3):348–62.

Haines, Michael R. and Inter-university Consortium for Political and Social Research. 2010. *Historical, Demographic, Economic, and Social Data: The United States, 1790–2002*. ICPSR 02896-v3 (computer file).

Hibbs, Douglas A., Jr. 1973. *Mass Political Violence: A Cross-National Causal Analysis*. John Wiley and Sons.

Hosking, J.R.M. 1990. “*L*-Moments: Analysis and Estimation of Distributions Using Linear Combinations of Order Statistics.” *Journal of the Royal Statistical Society, Series B* 52:105–124.

Inclán, María. 2009. “Sliding Doors of Opportunity: Zapatistas and Their Cycle of Protest.” *Mobilization* 14(1):85–106.

Jacobs, David and Kent, Stephanie L. 2007. “The Determinants of Executions since 1951: How Politics, Protests, Public Opinion, and Social Divisions Shape Capital Punishment.” *Social Problems* 54(3):297–318.

Jenkins, J. Craig, David Jacobs, and Jon Agnone. 2003. “Political Opportunities and African-American Protest, 1948–1997.” *American Journal of Sociology* 109(2):277-303.

Johnson, Neil F., Mike Spagat, Jorge A Restrepo, Oscar Becerra, Juan Camilo Bohorquez, Nicolas Suarez, Elvira Maria Restrepo, and Roberto Zarama. 2006. “Universal Patterns Underlying Ongoing Wars and Terrorism.” http://xxx.lanl.gov/abs/physics/0605035.

Koopmans, Ruud. 1995. *Democracy From Below: New Social Movements and the Political System in West Germany*. Boulder, Col.: Westview.

Kriesi, Hanspeter, Ruud Koopmans, Jan Willem Dyvendak, and Marco G. Giugni. 1995. *New Social Movements in Western Europe: A Comparative Analysis*. University of Minnesota Press.

Lohmann, Susanne. 1994. “The Dynamics of Informational Cascades: The Monday Demonstrations in Leipzig, East Germany, 1989–91.” *World Politics* 47(1):42–101.

Maney, Gregory M. and Oliver, Pamela E. 2001. “Finding Collective Events: Sources, Searches, Timing.” *Sociological Methods and Research* 30(2):131–69.

Mann, Leon. 1974. “Counting the Crowd: Effects of Editorial Policy on Estimates.” *Journalism Quarterly* 51(2): 278–85.

Margetts, Helen, Peter John, Scott Hale, and Taha Yasseri. 2015. *Political Turbulence: How Social Media Shape Collective Action*. Princeton University Press.

McAdam, Doug. 1982. *Political Process and the Development of Black Insurgency*. University of Chicago Press.

McAdam, Doug and Su, Yang. 2002. “The War at Home: Antiwar Protests and Congressional Voting, 1965 to 1973.” *American Sociological Review* 67(5):696–721.

McAdam, Doug, John D. McCarthy, Susan Olzak, and Sarah A. Soule. n.d. Dynamics of Collective Action in the United States (computer file). http:www.stanford.edu/group/collectiveaction/.

McCarthy, John D., Clark McPhail, and Jackie Smith. 1996. "Images of Protest: Dimensions of Selection Bias in Media Coverage of Washington Demonstrations, 1982 and 1991." *American Sociological Review* 61(3):478–99.

McCarthy, John D., Clark McPhail, Jackie Smith, and Louis J. Crishock. 1998. "Electronic and Print Media Representations of Washington D.C. Demonstrations, 1982 and 1991: A Demography of Description Bias." Pp. 113–30 in Dieter Rucht, Ruud Koopmans, and Friedhelm Neidhardt (eds), *Acts of Dissent: New Developments in the Study of Protest*. Berlin: Edition Sigma.

McCarthy, John D., Larissa Titarenko, Clark McPhail, Patrick S. Rafail, and Boguslaw Augustyn. 2008. "Assessing Stability in the Patterns of Selection Bias in Newspaper Coverage of Protest During the Transition from Communism in Belarus." *Mobilization* 13(2):127–46.

McPhail, Clark and John McCarthy. 2004. "Who Counts and How: Estimating the Size of Protests." *Contexts* 3(3):12–18.

Morgan, William R. and Clark, Terry Nichols. 1973. "The Causes of Racial Disorders: A Grievance-Level Explanation." *American Sociological Review* 38(5):611–24.

Myers, Daniel J. 2010. "Violent Protest and Heterogeneous Diffusion Processes: The Spread of U.S. Racial Rioting from 1964 to 1971." *Mobilization* 15(3):289–321.

Myers, Daniel J. and Caniglia, Beth Schaefer. 2004. "All the Rioting That's Fit to Print: Selection Effects in National Newspaper Coverage of Civil Disorders, 1968–1969." *American Sociological Review* 69(4):519–43.

Oberschall, Anthony R. 1994. "Rational Choice in Collective Protests." *Rationality and Society* 6(1):79–100.

Oliver, Pamela E. and Maney, Gregory. 2000. "Political Processes and Local Newspaper Coverage of Protest Events: From Selection Bias to Triadic Interactions." *American Journal of Sociology* 106(2):463–505.

Oliver, Pamela E. and Myers, Daniel J. 1999. "How Events Enter the Public Sphere: Conflict, Location, and Sponsorship in Local Newspaper Coverage of Public Events." *American Journal of Sociology* 105(1):38–87.

Olzak, Susan. 1989. "Analysis of Events in the Study of Collective Action." *Annual Review of Sociology* 15:119–41.

Olzak, Susan. 1992. *The Dynamics of Ethnic Competition and Conflict*. Stanford University Press.

Opp, Karl-Dieter. 2009. *Theories of Political Protest and Social Movements: A Multidisciplinary Introduction, Critique, and Synthesis*. Abingdon: Routledge.

Ortiz, David G., Daniel J. Myers, Eugene N. Walls, and Maria-Elena D. Diaz. 2005. "Where Do We Stand with Newspaper Data?" *Mobilization* 10(3):397–419.

Popovic, Srdja and Miller, Matthew. 2015. *Blueprint for Revolution: How to Use Rice Pudding, Lego Men, and Other Non-Violent Techniques to Galvanise Communities, Overthrow Dictators, or Simply Change the World*. Melbourne: Scribe.

Przeworski, Adam. 2009. "Conquered or Granted? A History of Suffrage Extensions." *British Journal of Political Science* 39(2):291–321.

Resnick, Sidney I. 2010. *Heavy-Tail Phenomena: Probabilistic and Statistical Modeling*. New York: Springer.

Richardson, Lewis F. 1948. "Variation of the Frequency of Fatal Quarrels with Magnitude." *Journal of the American Statistical Association* 43:523–46.

Rozenfeld, Hernán D., Diego Rybski, Xavier Gabaix, and Hernán A. Makse. 2011. "The Area and Population of Cities: New Insights from a Different Perspective on Cities." *American Economic Review* 101(5):2205–25.

Rucht, Dieter and Neidhardt, Friedhelm. 1998. "Methodological Issues in Collecting Protest Event Data: Units of Analysis, Sources and Sampling, Coding Problems." Pp. 65–89 in Dieter Rucht, Ruud Koopmans, and Friedhelm Neidhardt (eds), *Acts of Dissent: New Developments in the Study of Protest*. Berlin: Edition Sigma.

Rummel, Rudolph J. 1963. "Dimensions of Conflict Behavior Within and Between Nations." *General Systems Yearbook* 8:1–50.

Seguin, Charles. In press. "Cascades of Coverage: Dynamics of Media Attention to Social Movement Organizations." *Social Forces*.

Shorter, Edward and Tilly, Charles. 1974. *Strikes in France, 1830-1968*. Cambridge University Press.

Soule, Sarah A. and Earl, Jennifer. 2005. "A Movement Society Evaluated: Collective Protest in the United States." *Mobilization* 10(3):345–64.

Spielmans, John V. 1944. "Strike Profiles." *Journal of Political Economy* 52(4):319–39.

Tarrow, Sidney. 1989. *Democracy and Disorder: Protest and Politics in Italy, 1965–1975*. Oxford: Clarendon Press.

Tarrow, Sydney. 2011. *Power in Movement: Social Movements, Collective Action, and Politics*, 3rd ed. Cambridge University Press.

Tilly, Charles. 1978. *From Mobilization to Revolution*. New York: McGraw-Hill.

Tilly, Charles. 1995. *Popular Contention in Great Britain, 1758–1834*. Harvard University Press.

Tilly, Charles. 2008. *Contentious Performances*. Cambridge University Press.

Tilly, Charles and Rule, James B. 1965. *Measuring Political Upheaval*. Centre of International Studies, Princeton University.

Titarenko, Larissa, John D. McCarthy, Clark McPhail, and Boguslaw Augustyn. 2001. "The Interaction of State Repression, Protest Form and Protest Sponsor Strength During the Transition from Communism in Minsk, Belarus, 1990–1995." *Mobilization* 6(2):129–50.

U.S. Bureau of the Census. 1960. *Census of Population, 1960*, Vol. 1, *Characteristics of the Population*, parts 2–52. Washington, D.C.: Government Printing Office.

Virkar, Yogesh and Clauset, Aaron. 2014. "Power-Law Distributions in Binned Empirical Data." *Annals of Applied Statistics* 8(1):89–119.

Yip, Paul S. F. et al. 2010. "Estimation of the Number of People in a Demonstration." *Australian and New Zealand Journal of Statistics* 52(1):17–26.

Table 1. Size distributions of protest events

	Demonstrations in USA	*Demonstrations in UK*	*Demonstrations in DC*	*Strikes in UK: workers involved*	*Riots in USA: arrests*	*Riots in USA: arrests / nonwhites*
Number of events	7,878	1,047	3,065	78,378	593	555
Minimum	1	1	2	<10	1	0.0001%
Median	201	300	23	70	17	0.1033%
Mean	2,824	6,693	929	628	116	0.3220%
Maximum	600,000	300,000	500,000	1,750,000	7,772	4.0307%
Standard deviation	19,061	25,695	10,931	16,996	565	0.5266%
Coefficient of variation	6.7	3.8	11.8	27.0	4.9	1.6
L-skew	.88	.82		.87	.82	.54
L-kurtosis	.78	.67		.79	.69	.28
Gini index	.91	.88	.94	.87	.84	.69

Table 2. Size distributions of protest events

	Strikes in Illinois outside Chicago: workers involved	*Strikes in Peoria: workers involved*
Number of events	551	29
Minimum	2	4
Median	80	40
Mean	250	95
Maximum	8,929	600
Standard deviation	696	145
Coefficient of variation	2.8	1.5
L-skew	.68	.57
L-kurtosis	.51	.33
Gini index	.72	.66

Table 3. Total participation and large events

	Demonstrations in USA by year	*Demonstrations in UK by year*	*Strikes in UK by year: workers involved*	*Strikes in Illinois outside Chicago by year: workers involved*	*Riots in USA by city: arrests*
Threshold defining large events	10,000	10,000	10,000	1,000	1,000
Large events / all events	4.1%	11.3%	0.39%	5.1%	1.8%
Correlation between total participation and—					
participation in large events	.98	.99	.98	.93	.97
frequency of large events	.76	.45	.68	.83	.87

Figure 1. Cumulative size distributions: mass-count disparity

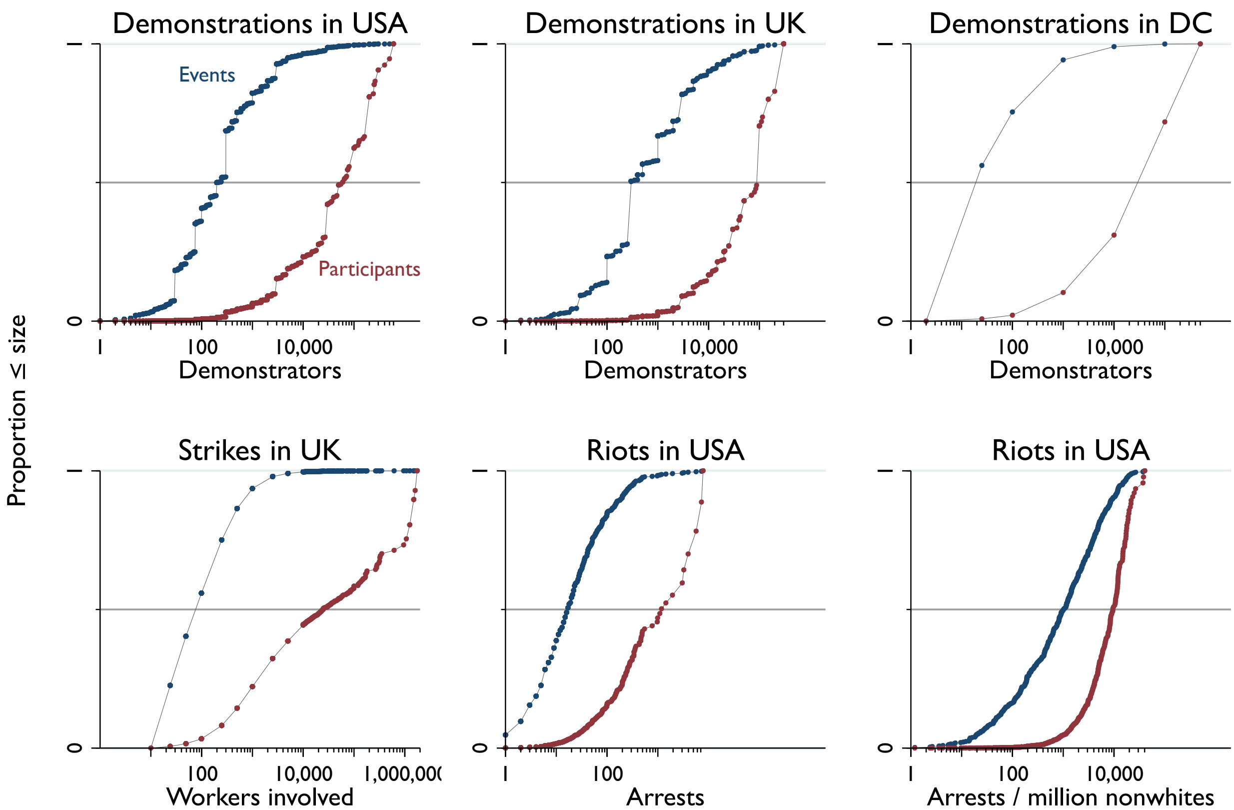

Figure 2. Strikes in UK: complementary cumulative size distribution

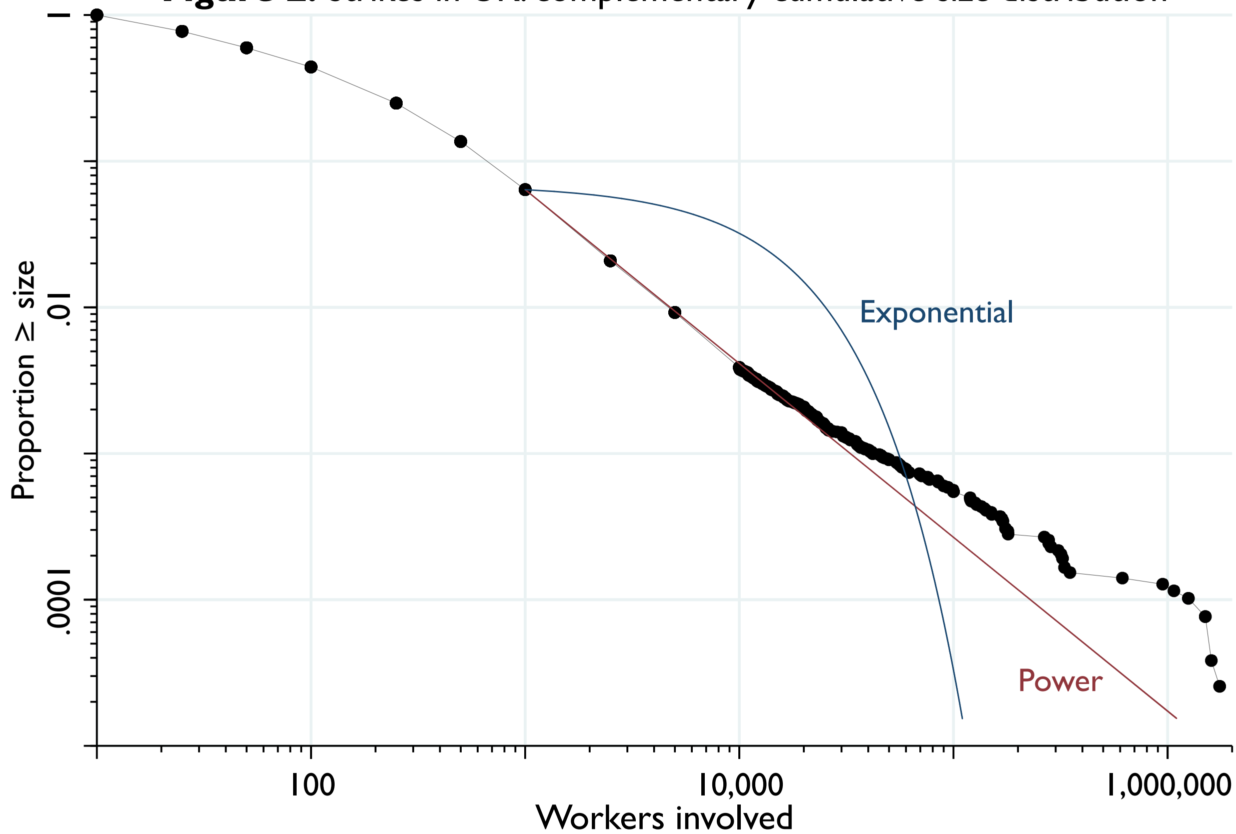

Figure 3. Demonstrations in USA: total participation and event frequency

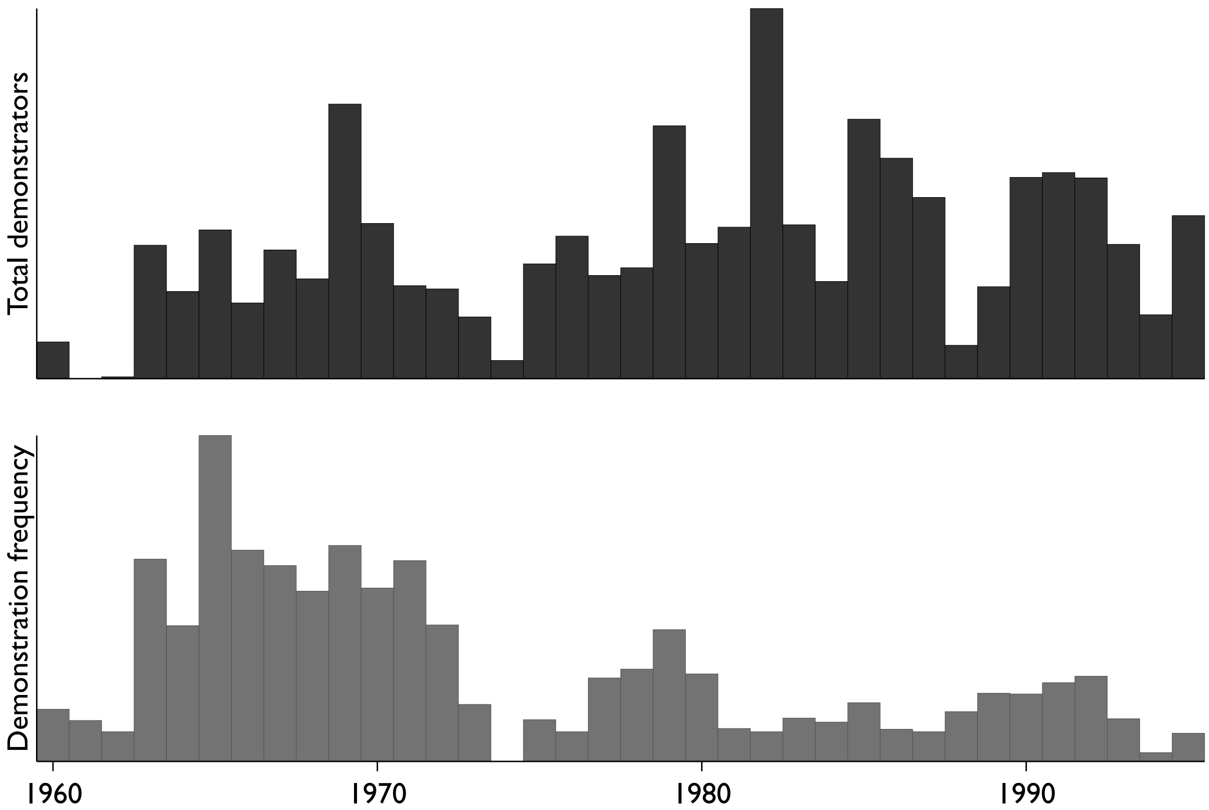

Figure 4. Event frequency and total participation

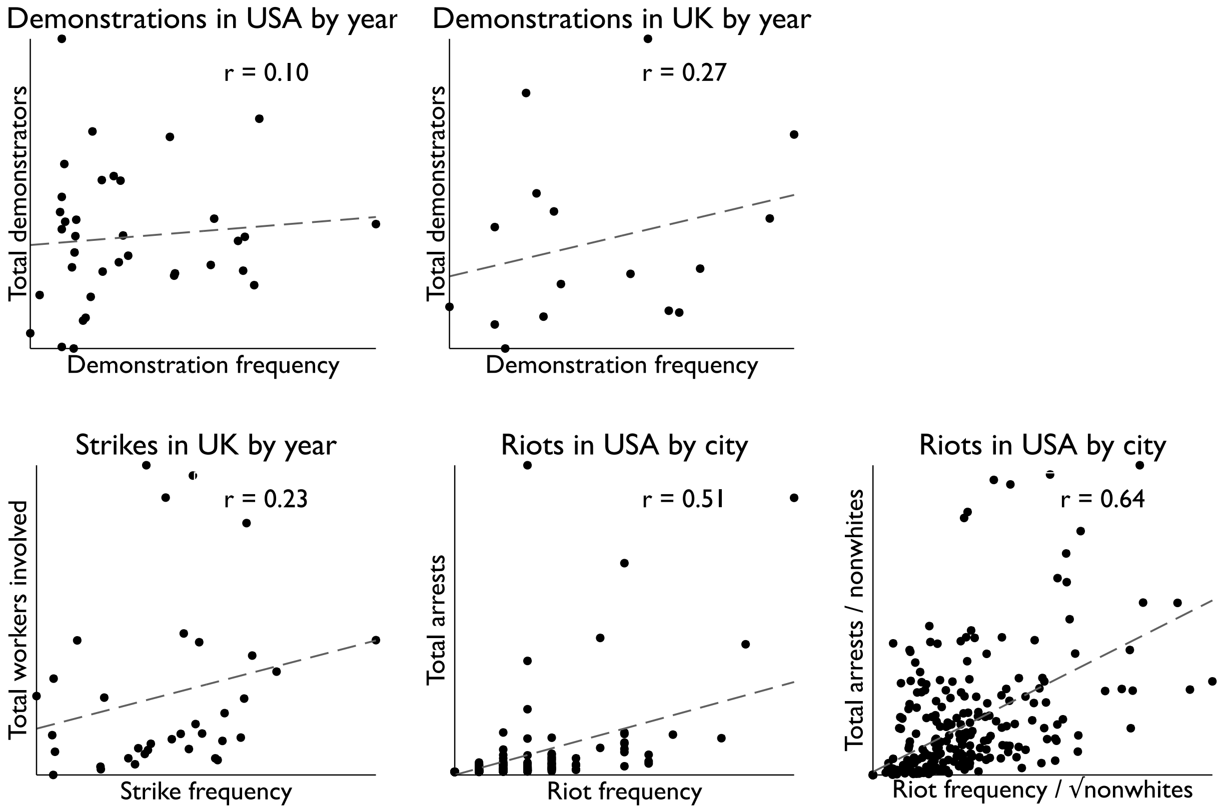

Figure 5. Strikes in UK: workers involved in events of varying size