Is Police Behavior Getting Worse? Data Selection and the Measurement of Policing Harms

Aurélie Ouss and John Rappaport

ABSTRACT

Public concern about harmful policing is surging. Governments are paying historic amounts for law enforcement liability. Has police behavior changed? Or is society responding differently? Traditional data sources struggle with this question. Common metrics conflate the prevalence and severity of policing harms with the responses of legal actors such as lawyers, judges, and juries. We overcome this problem using a new data source: liability insurance claims. Our data set contains 23 years of claims against roughly 350 law enforcement agencies that contract with a single insurer. We find that, while lawsuits and payouts have trended upward over the past decade, insurance claims have declined. We examine multiple potential explanations. We argue that, in our sample, police behavior is not getting worse; rather, societal responses to policing harms are intensifying. Police litigation is not representative of the broader universe of claims, and adjudicated claims differ systematically from settled ones.

1. INTRODUCTION

The deaths of Michael Brown and Eric Garner in the summer of 2014 fueled widespread public interest in, and media scrutiny of, police be-

havior.[1] In December of that year, Barack Obama created the President's Task Force on 21st Century Policing. In the years since, a steady stream of videos have appeared on the Internet depicting the use of force—sometimes deadly force—by law enforcement officers around the country. Meanwhile, newspapers have reported rising police-related payouts (Elinson and Frosch 2015; Winton 2017). The explanation for these trends remains obscure, however. Is police behavior getting worse (King 2015, 2018) or simply more salient, particularly with the rise of smartphones and body-worn cameras (McLaughlin 2015; Wines and Cohen 2015)? There is some evidence, for example, of increasing militarization of the police, which may in turn correlate with the fatal use of force (Delehanty et al. 2017; Lawson 2018). But it is also possible that growing societal intolerance of police-inflicted harms is driving up media coverage and liability costs without any change in underlying police behavior.

The types of policing data available have limited our ability to discriminate among these explanations and to ascertain the nature of, and trends in, police-inflicted harms more generally (see Harmon 2013; Kane 2007). Of all the challenges facing researchers who study these questions, perhaps foremost is that common data sources—such as lawsuits and payouts—conflate the prevalence and severity of policing harms with the responses of legal actors such as lawyers, judges, and juries. Ideally, researchers would track harmful policing incidents directly. Reliable data of that kind, however, are rarely available, especially at scale. Given this reality, the search is for data that bring researchers closest to the harmful behavior, such as by capturing civilian allegations that an officer has inflicted compensable harm. If the dispute-resolution process is a pyramid, in other words, with jury trials at the peak (Galanter 1990), the base is the set of all potentially compensable policing harms. Researchers interested in understanding harmful police behavior want to be as close to the base as possible.

There is a seldom-used type of data that brings us much of the way

---

1. Querying Google Trends (see Stephens-Davidowitz and Varian 2015) for the phrase "police brutality," for example, shows a jump in search activity in the United States in August 2014. Since then, search activity has remained at levels generally higher than during the preceding decade or more. Figure A1 displays Google search activity for "police brutality" in the United States indexed to all Google searches and normalized to reach 100 in the highest relative search month; the vertical line marks August 2014. Similarly, in searching the US Newspapers collection in LexisNexis for the phrase "police brutality," a spike in monthly coverage appears in August 2014, marked by a vertical line in Figure A2.

there, permitting researchers to make progress in disentangling police be-
havior from the responses it triggers: liability insurance claims. We ex-
ploit this unusual data source in this paper. Police liability insurers open
a claim every time a civilian seeks financial compensation for a police-
inflicted harm, regardless of whether a lawsuit is filed or a payout is ul-
timately made. While there is selection by civilians, sometimes with the
assistance of counsel, regarding when to file claims, a claim reflects the
first step in the dispute-resolution process. Claims data are not filtered
through the additional layers of selection that produce litigation and
payout data and that may bias our understanding of police behavior. By
comparing claims data with these other data types, moreover, we are able
to isolate and describe the latter selection process.

We demonstrate the importance of these points with a novel data set
containing 23 years of police liability claims against roughly 350 small
and midsized police agencies in a midsized state (by population), com-
piled by a single insurer whose identity we have agreed to conceal. In
addition to bringing us closer to the harmful behavior researchers aim
to study, the data set's expansive time horizon—covering 23 years up
through 2017—allows us to measure time trends, which helps contextu-
alize today's contentious policing environment. That our data cover hun-
dreds of agencies, moreover, ensures that the trends we observe are not
driven by idiosyncrasies in a single or small number of jurisdictions. In-
deed, we are aware of no other data set that covers so many agencies for
so many years.[2]

Our data allow us to make three central contributions. First, from
a methodological perspective, we demonstrate the extent to which the
analysis of policing harms depends on the type of data used to measure
those harms. By comparing our overall sample with the subsets of claims
that involve lawsuits or lawsuits with payouts, we show that police litiga-
tion—however ultimately resolved—is not representative of the broader
universe of claims. In fact, using lawsuits or payouts to proxy for policing
harms would have led to opposite conclusions about the direction of po-
lice performance in our sample over the past decade. Furthermore, claims
data allow us to observe settled cases as well as adjudicated ones, so we
can directly test the Priest-Klein hypothesis, which predicts bias in the se-

2. Kane and White (2013, p. 3) present "perhaps the largest study of police miscon-
duct ever conducted in the United States, with a study period spanning from 1975 through
1996," a total of 22 years. They study a single agency, albeit the nation's largest—the
New York City Police Department.

lection of disputes for litigation. We show that, in our data, adjudicated claims differ systematically from settled ones, consistent with the Priest-Klein hypothesis but in contrast to recent empirical work (Helland, Klerman, and Lee 2018). We discuss how these contrasting results may stem from differences in the institutional environments we study.

Second, we present new evidence on the prevalence, composition, cost, and resolution of police liability allegations over time in a stable set of agencies. Consistent with evidence from other sources (Elinson and Frosch 2015; Winton 2017), we find that payouts have trended upward over time, and sharply so in recent years. The rate at which claims pay out and the number of lawsuits have also climbed since 2005. Yet the overall number of claims has declined, possibly in part because of loss-prevention techniques like those detailed in Rappaport (2017). Our data come from a single state and from law enforcement agencies serving small and midsized localities—the average jurisdiction in our data set has roughly 9,000 inhabitants. Although these are not the places researchers of policing most commonly study, they are home to many millions of Americans. According to the 2010 census, around 20 percent of the US population lives in municipalities with fewer than 2,500 residents, and 30 percent lives in municipalities under 50,000 (Ratcliffe et al. 2016). Furthermore, 70 percent of police agencies serve localities with fewer than 10,000 inhabitants (Reaves 2015). Our data therefore shed light on a large and understudied slice of American policing.

Finally, we consider several potential explanations for the opposing trends we observe in the number of claims and resolution of those claims. These include the rise in smartphone and body-worn camera technologies, changes in qualified-immunity doctrine, and deteriorating police-community relations. We argue that the observed trends are most consistent with an intensified societal response to policing harms. Through this exercise, this paper extends the economics-of-crime literature focused on police behavior. A sizable literature studies the effects of policing on crime (for a review, see Levitt and Miles 2006). Other work examines the determinants of officers' productivity (Chandrasekher 2016; Mas 2006) and racial disparities in stops and arrests (see, for example, Anwar and Fang 2006). Economic research into police misconduct, specifically, has emerged only recently. In the last few years, researchers have measured police misconduct either as the primary outcome of interest (Rozema and Schanzenbach 2019; Fryer 2019) or as the consequence of different policy environments (Chandrasekher 2017; Dharmapala, McAdams, and

Rappaport 2019). To our knowledge, our paper is the first to attempt to illuminate the distinction between long-term trends in police behavior and societal responses to it.

The rest of the paper is organized as follows. Section 2 explains how policing data have hamstrung prior research and how insurance claims data allow researchers to tackle questions incompletely explored by prior work. Section 3 describes our claims data in detail. Section 4 illustrates the data's potential by reporting and interpreting results concerning trends in police liability. It then more broadly demonstrates the importance of data selection for policing research. Section 5 discusses the results, using the data to benchmark findings from prior research and outlining the study's limitations. Section 6 concludes.

2. MEASURING POLICE BEHAVIOR

Section 2.1 describes how prior policing research has been limited by its data. Section 2.2 explains why insurance claims data give researchers traction on important but underexamined questions.

2.1. Prior Literature

Research on police behavior has generally relied on a small number of data sources. We review those sources here, noting various shortcomings that have limited the reach of prior work.

2.1.1. Lawsuits. Many studies, especially those focused on financial liability, exploit publicly available records of lawsuits (Eisenberg 1982; Kappeler 2006; Leong 2012; Powell, Meitl, and Worrall 2017; Ross 2000; Worrall and Gutierrez 1999). Because procedural and financial obstacles tend to "filter out marginal or frivolous cases," lawsuits are thought to constitute a high-quality sample of liability disputes (Rozema and Schanzenbach 2019, p. 231). At the same time, legal scholars have long believed that cases selected for litigation are "neither a random nor a representative sample" of the universe of disputes (Priest and Klein 1984, p. 4). Among other concerns, defendants can influence the number of suits filed against them. They can, for example, opt to settle politically embarrassing disputes or high-quality claims they expect to lose. An agency's low lawsuit count, therefore, could signal either responsible policing or an aggressive settlement strategy. In addition, plaintiffs' lawyers filter cases not for legal merit but for expected value, which encompasses

concerns orthogonal to merit such as the personal characteristics of potential clients and the nature of their injuries. Plaintiffs' attorneys may be more likely to file low-quality, high-dollar lawsuits than high-quality, low-dollar suits, which leaves the latter undercounted.

2.1.2. Payouts. Data on public payouts, also familiar to academic research (Iris 2012, 2014; Rozema and Schanzenbach 2019; Schwartz 2014, 2016), may capture some cases that lawsuit data miss. Few jurisdictions make payout data readily accessible, though they are typically covered by public records laws and can be gathered with sufficient effort. Yet payout data are also characterized by significant selection effects, as disputes are filtered by multiple legal actors on the path to resolution. Unsuccessful claims, too, are wholly omitted. As a practical matter, in working with payout data, it can also be difficult to determine what cases localities are counting as police related (Iris 2014).

2.1.3. Civilian Complaints. Another approach is to examine civilian complaints about police behavior (Chandrasekher 2017; Chanin 2016; Kane and White 2009; Lersch and Mieczkowski 2000). Civilian complaints can be valuable for comparing officers in a large agency (see Rozema and Schanzenbach 2019). Yet researchers cannot meaningfully aggregate or compare civilian complaints across agencies because data structures vary substantially (Hickman and Poore 2016). Relatedly, Ba (2018) finds that the physical location of the oversight agency where complaints are signed affects the likelihood that civilians will complete the complaint-filing process. This suggests that small barriers can have real effects on civilian reporting; it is certainly possible that more substantial obstacles have even larger effects, which makes cross-jurisdictional or even intertemporal comparisons difficult to interpret.

In addition, civilian-complaints data are frequently unavailable or unreliable, especially for smaller agencies. Hickman and Poore (2016, p. 472) find "many examples of erratic reporting" in their audit of data submitted by small agencies (100–249 officers) to the Law Enforcement Management and Administrative Statistics (LEMAS) survey. Many agencies in their study, moreover, both large and small, purge their records under record-retention laws after as little as 3 years. Civilian complaints and insurance claims also capture different types of police behaviors. Where good civilian-complaints data are available, they tend to feature allegations of failure to provide service and discourtesy (Ba 2018; New York City Civilian Complaint Review Board 2017). These are important

for understanding police-community relations but less helpful for studying the types of incidents that trigger the regulatory apparatus of civil liability. Indeed, there are disincentives to filing civilian complaints over incidents that cause financially compensable harm, as documents submitted in connection with the complaint may become admissible in any later litigation (Rozema and Schanzenbach 2019).

2.1.4. Civilian Deaths.   In recent years, journalists and activists have compiled databases that collect information on officer-involved shootings and civilian deaths. Researchers have begun to use these data sources, which are close to comprehensive, in peer-reviewed work (Campbell, Nix, and Maguire 2017; Jennings and Rubato 2017; Legewie and Fagan 2016; Mesic et al. 2018; Nicholson-Crotty, Nicholson-Crotty, and Fernandez 2017; Ross 2015). Yet the most inclusive database, Mapping Police Violence, contains only 6 years of data. For measuring police misconduct, moreover, these databases are simultaneously underinclusive, as they capture only the most extreme incidents, which are statistically rare, and overinclusive, as they include some legally justified killings that may not count as misconduct by reasonable standards, such as cases of officer self-defense.

2.2. Insurance Claims Data

Insurance claims data allow us to tackle some questions incompletely explored by prior work. A challenge in policing research is that most common data sources—particularly lawsuits and payouts—conflate the prevalence and severity of policing harms with the responses of legal actors such as lawyers, judges, and juries. Using insurance claims data to study policing does not eliminate these selection effects, as there is still selection into claims, but it does reduce their influence, as claims lie closest to the base of the dispute-resolution pyramid. In particular, we can track disputes beginning when they are initially filed and before claim processing begins. Defendants cannot control the number of claims they incur by changing their settlement behavior. The plaintiffs' bar also plays a smaller role in filtering claims, as claims are easier than lawsuits to file without legal assistance.[3] Cutting through these layers of selection allows

---

3. There are no systematic data on the rate of representation of insurance claimants. As one point of reference, in the past several years, the rate was around 20–30 percent in our sample.

researchers to disaggregate the responses of legal actors such as lawyers, judges, and juries from the prevalence and severity of policing harms.

In addition, because the insurance mechanism requires spreading risk, each insurer typically covers many law enforcement agencies. This mitigates the concern that observable patterns in the data might reflect the idiosyncrasies of one agency or a small number of agencies. Data for covered agencies, moreover, are collected consistently across time and place and are well preserved, which makes them valuable for studying time trends and—while not the focus of this paper—cross-jurisdictional questions as well.

Finally, insurance data give researchers an unprecedented view of policing outside the nation's largest cities. Empirical research rarely covers smaller agencies,[4] which generates an oft-noted big-city bias in the literature (see Barrett, Haberfeld, and Walker 2009; Crow and Adrion 2011; Terrill, Leinfelt, and Kwak 2008; Walker and Katz 2008). This scholarly bias mirrors a similar skew in national media coverage. In recent years, just three cities—New York, Chicago, and Los Angeles—generated between half and two-thirds of all stories about police misconduct in the Associated Press and the 50 most widely circulated newspapers, despite employing less than 12 percent of sworn officers nationwide and accounting for only 4 percent of police-related civilian deaths.[5]

Big-city bias may warp our understanding of policing harms to the extent that big-city policing differs from policing elsewhere. The United States has over 12,000 local police departments. Roughly half of them

---

4. Indeed, it is not uncommon for studies to exclude agencies that serve fewer than 100,000 residents (Parker et al. 2005; Smith 2004). The few exceptions consist principally of descriptions of individual jurisdictions (Garcia and Cao 2005; Lytle and Randa 2015; Payne, Berg, and Sun 2005; but see Liederbach [2005, 2007], which examine multiple jurisdictions). Yet some evidence suggests that smaller agencies differ meaningfully from larger ones in structure, professionalization, and personnel (Falcone, Wells, and Weisheit 2002; Weisheit, Wells, and Falcone 1995). Smaller departments, for example, often lack big-city protocols for investigating policing harms, including shootings of civilians (Toner and Rutecki 2018).

5. In the LexisNexis collection Major US Newspapers, which contains the top 50 US papers by circulation, a search for the phrase "police misconduct" on October 11, 2017, returned 431 stories published during the immediately preceding year. After filtering the results for duplicates and relevance, 66 percent of the stories concerned the three largest cities. Similarly, a search of the LexisNexis Associated Press collection on the same day generated 167 results. The top three cities accounted for 54 percent of the filtered results. The 10 largest cities generated 70 percent and 58 percent of the stories in the search results, respectively. Employment figures are from 2013, as reported in Reaves (2015). Data on police-related civilian deaths are from Mapping Police Violence (2019).

employ 10 or fewer officers (Reaves 2015). Around 78 percent of the country's sworn officers serve populations of fewer than a million residents, 66 percent serve populations below 500,000, and 46 percent serve populations smaller than 100,000 (Reaves 2015).[6] Nor does the lion's share of crime and policing happen in the biggest cities (Fox 2017). Back-of-the-envelope calculations suggest that roughly two-thirds of Federal Bureau of Investigation index crimes, arrests, and police-related fatalities occur outside the 100 largest cities. (For perspective, Fremont, California; Garland, Texas; and Hialeah, Florida, make the 100-largest-cities list.)[7] If anything, these numbers may increase over time if experts are right that "the periphery," rather than the urban core, "remains the dominant, and fastest growing, part of the American landscape" (Kotkin and Berger 2017; Berger and Kotkin 2017).

Only two studies, to our knowledge, rely on insurance-claims data. McCoy (2010) exploits claims data from 1974 to 1984 from a defunct insurer that covered counties nationwide. She uses the data to assess the impact of the US Supreme Court's 1978 decision holding municipalities amenable to suit under federal civil rights law. Using on a before/after comparison, she finds that the number of claims and their outcomes did not materially change but that the number of wrongful-death lawsuits increased. Closest to our endeavor, Ross and Bodapati (2006) analyze 15 years of law enforcement liability claims data—from 1985 to 1999—from the Michigan Municipal Risk Management Authority. They present certain aggregate information over time, such as the total cost of claims for all agencies and the percentage of claims that were litigated. Their data, however, are limited in several important respects. First, Ross and Bodapati cannot verify whether the same agencies were covered throughout their observation period, which makes it difficult to interpret time trends. Second, they cannot determine, in general, whether their claims relate to law enforcement or to detention functions. In addition, their data set includes employment claims, and their liability claims are domi-

---

6. The figures look even starker if one counts agencies rather than officers: 94 percent of police departments serve fewer than 50,000 people, 70 percent serve fewer than 10,000 people, and 40 percent serve fewer than 2,500 people (Reaves 2015).

7. Data for the crime and arrest figures are from the Federal Bureau of Investigation's Uniform Crime Reports for 2016 as compiled in Kaplan (2017). Crimes are reported offenses; arrests are clearances by arrest. Where cities did not report, they were excluded from the city counts. Data for the police-related civilian-deaths calculation are from Mapping Police Violence (2019) for January 2013 through November 2017 and from 2016 US Census estimates.

nated by auto accidents. Finally, their observation period ends almost 20 years ago, so they cannot speak to recent trends, which hold substantial interest given current policy debates.

Before describing our data set, we offer two caveats about claims data generally. First, even though claims data capture incidents that other data miss, they do not count every incident involving policing harm—there is selection into claims. Second, because the largest cities in the country, which may be some of the most racially diverse, tend to self-insure, claims data may be unhelpful for studying discriminatory policing unless they are drawn from states with diverse populations (or large populations of people of color) outside their biggest cities.

3. INSTITUTIONAL BACKGROUND AND DATA

We illustrate the points made in Section 2 by using insurance claims to study trends in police liability in our sample over the past 2 decades. Our data set consists of claims data provided by a single insurer in a midsized state under an agreement to preserve the insurer's anonymity. The data cover roughly 350 law enforcement agencies over a period of 23 years. This is the largest and most detailed data set of police insurance claims ever employed in academic research. We begin with a brief background on the insurer. We then describe the data set in detail.

3.1. Institutional Background

Police liability insurance is common nationwide (see Rappaport [2017] for details). Providers include both for-profit carriers and intergovernmental risk pools. The latter are organizations formed by groups of local governmental entities to finance a risk, typically by pooling or sharing that risk. The exact number of localities that purchase police liability insurance is unknown but is thought to be a majority share. The market breakdown is also unknown. As a very general matter, however, small localities tend to join pools, midsized entities divide between pools and commercial carriers, and the largest localities self-insure. Even self-insured localities, though, often purchase excess coverage on the market.

The insurer in our study furnishes third-party liability coverage to its policyholders on a variety of lines, including automobile insurance, employers' liability, and workers' compensation. We focus on the coverage it provides for liability triggered by police operations, which is part of a

comprehensive liability-coverage package. Roughly 350 agencies in one midsized state obtain police liability coverage from the insurer. Consistent with the general pattern nationwide, a handful of the largest cities in the state, fewer than 10, self-insure and are not included in this total. The insurer captures the vast majority of the state's municipal police departments, however, corresponding to about 60 percent of the state's population. The covered agencies tend to be small, with most having 10 or fewer full-time officers.

The insurer's police liability coverage is of the typical sort described in Rappaport (2017). It reaches nearly all losses stemming from sworn police activities that cause personal injury or property damage. The localities and their officers all are covered; for convenience, though, we generally refer to claims as running against agencies. Deliberate violations of the law and criminal acts are excluded from coverage, though the insurer will typically defend an officer suspected of such conduct under a reservation of rights; the insurer will also defend the locality. Acts committed outside an officer's official capacity are excluded as well.[8]

A claim, for our purposes, represents a file the insurer opens on demand by an individual (the claimant) for compensation from a covered agency for a police-related harm. Most often, claimants make demands to the localities or agencies, which report them to the insurer. On rare occasions, claimants submit demands directly to the insurer, which then coordinates with the locality, or the insurer opens a claim in anticipation of a demand, which is typically followed by an actual demand. The claims-making process is consistent across localities. To our knowledge, the principal change in the process over time has been increasing reliance on the Internet as a substitute for in-person claims submission.

The claims process is independent from litigation in the sense that claimants are not required to initiate or exhaust the claims process before suing, nor are they required to sue at all to receive compensation. In other words, some claimants never sue, some sue without first engaging in out-of-court negotiation with the insurer, and some sue after negotiation fails. If the process is started when the claimant initiates a lawsuit, this is also considered a claim for our purposes, which means that every

---

8. The deliberate-and-criminal-acts exclusion would apply to incidents in which there is no arguable legal justification for the officer's act, such as the sexual assault of a detainee. Even then, many insurers would defend the officer, subject to a reservation of rights, unless and until he is convicted of a crime (Rappaport 2017). The insurer informed us that it very rarely invokes the exclusion, which is consistent with general industry practice (Rappaport 2017).

lawsuit against a covered agency for which the insurer provides coverage will be included in our data set.

3.2. Data

The insurer initially provided us a raw, claims-level data set of all police liability claims filed between 1992 and June 30, 2018, containing a total of 3,353 claims. Unlike in Ross and Bodapati (2006), police-related claims that fall under separate employment or automobile policies are not included, nor are workers' compensation claims. It is worth noting at the outset, however, that not all of our claims allege what, colloquially, might be called police misconduct. As Kane and White (2013) discuss at length, there is no expert consensus on how to define this term. Merely negligent infliction of personal or property harm, for example, may violate state tort law but probably not the US Constitution. Is this misconduct? It is behavior society would prefer not occur, and it can cause serious harm (including death). Yet it lacks the element of willfulness some might expect from police misconduct. For this reason, when referring to claims generically, we use terms like "policing harms" or "police-inflicted harms," reserving the term "misconduct" for discussing prior literature using the term.

Each claim in the data set has a unique identifier and is associated with a specific law enforcement agency. For each claim, we observe a brief description of the underlying incident (the loss description). For claims that are closed, we see whether a lawsuit was filed—and, if so, how it was resolved—as well as payouts, expenses, and deductible payments. We also observe the dates on which the incident allegedly occurred, the claim was reported to the insurer, and the insurer formally opened and closed the claim. In our temporal analyses, we focus on the reporting date.

The loss-description field characterizes the incident, not the legal arguments claimants make. If, for example, a claimant alleges that an officer used excessive force in arresting her and argues that the locality is liable because it failed to train the officer adequately, the loss description will reflect the underlying alleged use of force but typically will not mention the failure-to-train allegation. What we refer to as payouts includes money paid to the claimant under either a settlement or a judgment. It also includes the claimant's attorney fees, whether taken from the claimant's recovery—such as in a contingent-fee arrangement—or separately delineated, which is rare. When claims have multiple claimants, we include a single observation per claim, and we sum payouts made to all of

the claimants. Our data are therefore at the incident level and not at the claimant level.

The earliest reported claims are from 1992 but, because of data-quality concerns, we exclude the 333 claims reported before 1995. For analyses involving claims' outcomes, we limit our sample to closed-claims data, as open-claims data are necessarily incomplete. And because many claims from 2016 through 2018 remain open, we also exclude from these analyses 218 closed claims from those years. There are only eight open claims from 2015 and earlier, however, which means our pre-2016 data are essentially complete (Figure A3).

We also exclude certain types of claims against unusual agencies. First, we exclude 127 claims against a couple dozen multijurisdictional, specialized task forces focused on problems like drugs or violent crime. These agencies differ in kind from general services police departments. Among other things, oversight responsibility is spread among numerous jurisdictions, some of which are entities not otherwise covered by the insurer. In contrast, we do include in our main analyses claims against a few multijurisdictional, general services police departments. These agencies are formed by small clusters of neighboring municipalities, all of which are covered by the insurer, simply to achieve economies of scale. Second, we exclude 36 claims against agencies that did not exist or did not obtain coverage from the insurer throughout the entire observation period.[9] Our findings are qualitatively similar when we relax any of the above-mentioned restrictions.

This winnowing process left us with 2,858 claims across 23 years and 2,590 closed claims for which we can see claim outcomes across 21 years.[10] To study the substantive composition of claims, we developed a coding scheme consisting of the following eight salient types of police-related harms plus one residual category, which we applied to claims on the basis of the loss-description field:

9. We exclude claims from agencies that opened or closed or that obtained or lost coverage between 2007 and 2015, the period for which we have such information. Only 1.6 percent of claims between 2007 and 2015 and 1.26 percent of claims between 1995 and 2007 are from such agencies. Our results are very similar when we include these claims. There may be a small number of additional claims from agencies that opened or closed or that obtained or lost coverage between 1995 and 2007, but we cannot identify them. Given the stability of our results regardless of whether we include the claims from agencies that changed status between 2007 and 2015, we think their exclusion (were it possible) would not change our results substantially.

10. Note that there were multiple reasons to drop some claims. Our findings are qualitatively similar when we include any of the dropped claims.

Force: excessive force

Sexual: sexual misconduct (for example, sexual assault, coerced sexual activity)

Property Harm: property damage or loss (for example, damage to door on entry, items lost during arrest, animals injured, car damage from pursuit)

Seizure of Person: illegal seizure of a person (for example, false arrest, unlawful detention)

Seizure of Property: illegal seizure of property (for example, illegal forfeiture, improper towing, impoundment of vehicle)

Search: illegal search

Discrimination: discriminatory policing (including profiling based on race, sex, religion, ethnicity, or national origin)

Negligence: negligence resulting in personal injury (for example, inadvertent police dog bite, accidental discharge of weapon)

Other: for example, unspecified violation of constitutional rights, failure to investigate, defamation

We also flagged any claim that involved a civilian fatality or vehicular pursuit. Multiple claim types can be assigned to a single claim. Claims that do not fit into any of our categories are classified as "other."[11] The most common allegations in claims classified as "other" include failure to intervene or investigate, defamation, and misuse of confidential information.

Table 1 presents descriptive statistics on our closed-claims data, overall and by type of claim. Of the 2,590 closed claims, 32 percent involve a lawsuit, and 30 percent lead to some payout being made, though these figures vary greatly across claim type. Claims for property damage or loss are most common, followed by use of force. When successful (that is, when resulting in some payout), claims for discriminatory policing, use of force, and sexual misconduct result in the highest average payouts. Overall, mean payouts for successful claims are $23,700 (median, $700).

The data set contains an average of 123 claims per year across the covered agencies. In recent years, the agencies experienced an average of fewer than five claims per year per 100 officers. These seemingly small numbers may suggest that policing harms are less common (at least in

---

11. A pair of law-student research assistants manually and independently coded each claim. When they agreed on the coding, we adopted their coding decisions; when they disagreed, one of us made the ultimate determination of how to code the claim.

Table 1. Summary Statistics for Claims

Type of Claim	Closed Claims (N)	Lawsuit (%)	Payout (%)	Payouts (Real 2015 $)		
				Mean	SD	Median
Force	595	42	23	103,600	252,000	15,300
Sexual	26	35	38	100,700	156,300	44,000
Property harm	839	6	54	1,200	5,400	400
Seizure of person	334	54	26	32,400	61,900	6,900
Seizure of property	224	48	19	1,700	3,500	200
Search	88	56	35	45,000	72,800	9,200
Discrimination	191	64	6	106,500	279,400	16,600
Negligence	130	21	32	13,300	53,500	2,300
Other	390	40	10	19,300	35,900	3,600
Overall	2,590	32	30	23,700	116,000	700

Note. Payouts are for claims that led to positive payouts.

some places) than is thought, that many harms go unreported, or some combination of the two.

4. PRINCIPAL RESULTS

We report our main findings in three parts. In Section 4.1, we present time-trend evidence for the number and substantive composition of claims in our sample and for claim outcomes. In Section 4.2, we consider potential explanations for these trends. In Section 4.3, we use our data to demonstrate the importance of data selection in studying the police.

4.1. Trends in Frequency, Composition, and Resolution of Claims against Police Agencies

Our data, which span more than 2 decades, allow us to explore time trends in the number, type, and resolution of claims against police agencies in our sample. Motivating this analysis—as noted above—is that public interest in police brutality and media coverage have spiked in recent years (Figures A1 and A2). Is this because police behavior has worsened over time? Or are legal responses intensifying, possibly because people are simply paying more attention? Our data can help us get traction on these questions.

We ask first whether the total number of claims against the police in our sample has increased over time. Figure 1 plots the number of claims,
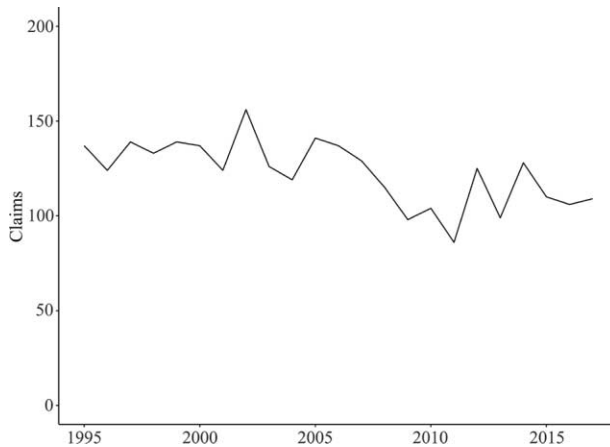
Figure 1. Claims against covered agencies

including open claims, against all covered agencies reported each year between 1995 and 2017. There is no apparent increase in recent years to explain the heightened interest in police behavior. In fact, the number of claims has decreased over time, from an average of 135 claims per year in the 1990s to 127 claims per year in the 2000s to 106 claims per year in the 2010s. Nor do we see any uptick in claims to match the spike in interest in the summer of 2014. Note that the downward trend in claims cannot be explained by a declining population in the localities served by agencies in our data, as the population was steadily rising while claims declined. Indeed, when we plot claims per capita (not reported), the trend looks very similar to Figure 1.[12]

Trends in the aggregate number of claims may mask variation in the substantive types of claims that citizens bring. If claims for property damage are going down, for example, while use-of-force claims are rising, the appearance of overall stability (or improvement) may be deceiving. The graphs in Figure 2 break out claims reported by substantive type, with different scales on the Y-axes across the graphs. While we do see year-to-year variation in the substantive composition of claims, there are no

12. Nor can the decline in claims be explained by changes in the claims-submission process over time. As noted in Section 3.1, the principal change has been increasing reliance on the Internet as a substitute for in-person submission, which would tend to increase the number of incidents reported in more recent years.
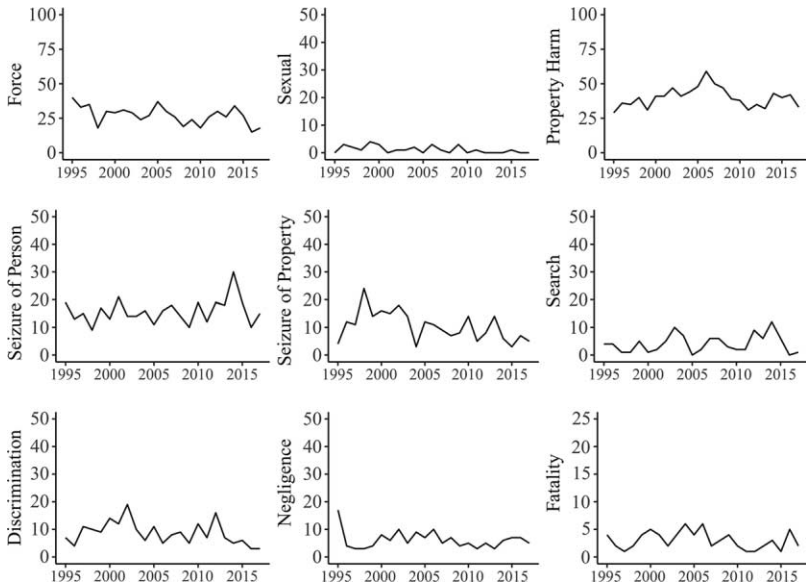
Figure 2. Claims by type

obvious trends and no clear increases in recent years for the most seri-
ous allegations—in particular, those involving fatalities, excessive force,
or seizure of a person (that is, false arrest). Overall, Figures 1 and 2 show
that increased scrutiny of the police has occurred in a context in which
claims were, if anything, decreasing.

When we look at payouts, however, the trends are anything but
stable, as shown in Figures 3 and 4. Figure 3*A* plots the total payouts
(in real 2015 dollars) made by all agencies for all claim types. Consistent
with media reports (Elinson and Frosch 2015; Winton 2017), the trend
is clearly upward, sharply so in 2014 and 2015. Between 2013 and 2015
alone, total annual payouts rose almost 10-fold, from $400,000 to $3.97
million. Multiple factors could be driving this trend in total payouts: a
rising claimant win rate (that is, an increasing share of claims leading to
payouts), rising average payouts per claim, or some combination of the
two. Figure 3*B* and *C* show that both factors are at play—claimants are
winning more often and taking home more dollars per successful claim
(that is, per claim with a positive payout). And, as with total payouts, av-
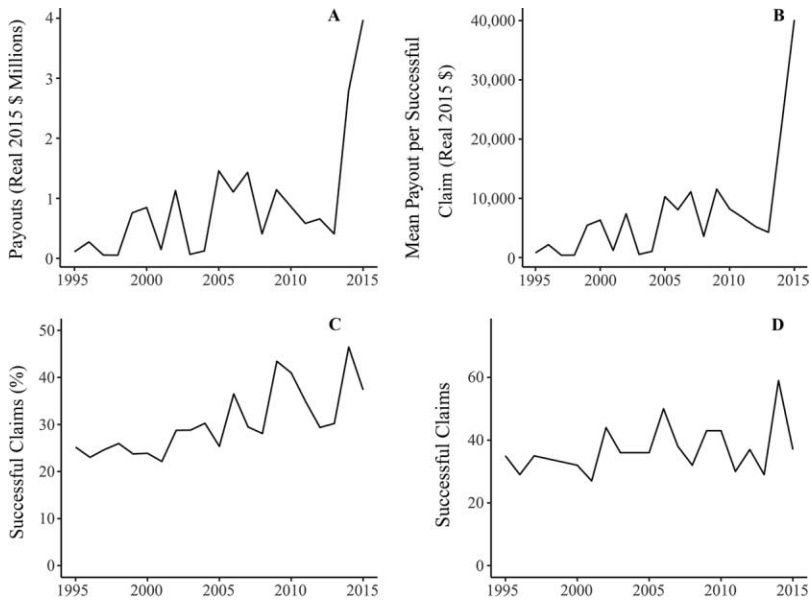erage payouts shoot rapidly upward in 2014 and 2015, rising from about

Figure 3. Claim outcomes

$4,000 to about $40,000 per claim.[13] Finally, Figure 3*D*, which plots the number (rather than percentage) of successful claims, shows that the rising win rate is partly, but not entirely, an artifact of the declining total number of claims. The number of successful claims increases throughout our observation period, from an average of 33 per year in the 1990s to 37 per year in the 2000s to 39 per year in the 2010s.

Particularly given the stark increase in recent years, it is possible that outliers are driving the increase in average payouts. Given that most claims are fairly small, a few multimillion-dollar claims involving fatalities, for example, could pull the averages up. Figure 4, however, rules out this explanation. Figure 4 presents the median, first-, and third-quartile

13. We do not include graphs showing total and mean payouts by claim type because, given that only 30 percent of claims pay out, the data are too sparse. Our analysis of the data does suggest, however, that the trends in total and mean payouts are driven largely by the trends in payouts for claims involving fatalities or excessive force. It is possible that some of the increase in observed payouts is attributable to changes in claimants' attorney fees rather than compensatory payments to the claimant. Because the division of payouts is typically a matter between claimants and their attorneys, our data do not permit us to speak to this point with precision. In the minority of claims in which separate payments were made to claimants and their counsel, however, we do not observe any upward trend in the attorney's share, though the number of observations is too small to support any firm conclusions.
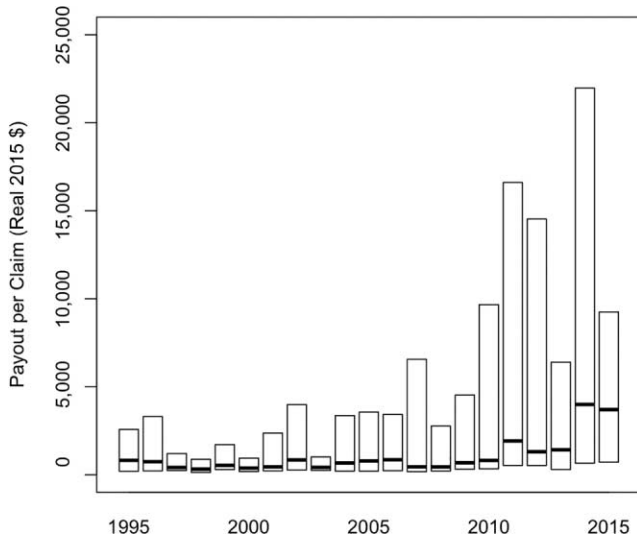
Figure 4. Median payouts per successful claim

payouts, which are also climbing during our observation period. Most of that increase occurred during the last 5 years of the sample. Between 2013 and 2015, the median payout per claim rose from $1,400 to $3,700. This is all the more significant given that the number of successful claims is rising too; it would be reasonable to think that low-dollar claims might be driving the latter trend, which would drive median payouts down.

Figures 1–4, then, leave us with the following puzzle: over the past 2 decades, the overall number of claims in our sample is stable or declining. This might suggest that the police are involved in a shrinking number of harmful incidents—in other words, that police behavior is improving in our sample. Yet there is also some evidence suggesting that police behavior may be getting worse. The number of successful claims is climbing. Similarly, payouts—total, mean, and median—are rising and spike sharply upward in recent years. In Section 4.2, we consider multiple potential explanations for this collection of trends.

4.2. Mechanisms and Interpretation

Several phenomena could explain the juxtaposition of falling claims with a rising claimant win rate and rising payouts. We consider three candidate explanations in turn: claimants' selection of claims may have improved; policing harms may have become more severe, even if less frequent; and societal responses to policing harms may have grown more

punitive. While our data do not permit a definitive answer, we argue that the available evidence favors the last explanation.

4.2.1. Improved Claim Selection.   The first potential explanation is that, during our observation period, claimants became more selective in bringing claims. This could have occurred for a variety of reasons, though two stand out from policing debates: more video evidence of encounters between the police and citizens (hereafter, police-citizen encounters) and changes in qualified-immunity doctrine. We consider each separately, as each generates slightly different predictions.

The introduction of body-worn cameras and the rising popularity of smartphones have improved available evidence of some police-citizen encounters. Prospective claimants with low-quality claims may have become less likely to file knowing that video footage would tend to defeat their allegations. Claimants with high-quality claims, in contrast, possess better evidence than in the past. This improves their odds of winning and increases expected damages per win, assuming video evidence elicits emotional responses from jurors by making harms more vivid. If this explanation is correct, we would expect to see several patterns in the data. Coincident with the introduction of these video-recording technologies, the number of claims should decline, as low-quality claims are withheld; claimants' win rate should rise, as average claim quality increases; and average payouts should rise, as average claim quality and the emotional appeal of claimants' evidence both increase. The explanation is consistent with multiple patterns in total payouts—whether total payouts increase depends on whether the effects of higher average payouts swamp the effects of reduced frequency of claims. This explanation, in other words, is consistent with the rising trend in total payouts but does not require it.

Although we observe all three of these trends over our sample period, the timing of pertinent events allows us to reject this explanation: the observed trends occur too early relative to the popularization of these video-recording technologies. According to the insurer, as of 2016 fewer than 10 of the 350 agencies in our sample were using body-worn cameras. Similarly, very few, if any, claims during our sample period involved video evidence from smartphones.[14] The rise of these technologies may

14. The insurer's reported experience on this point is consistent with other evidence regarding the popularization of filming the police with smartphones. For example, the American Civil Liberties Union released its Mobile Justice smartphone application, which facilitates civilian recording of the police, shortly before the end of our sample period. It is worth noting, too, that our claims are organized by the date on which they were

change patterns of police liability in the future, but it cannot explain the trends we examine here.

Developments in legal doctrine may also have altered claim selection. In particular, changes in qualified-immunity doctrine, which shields government employees from suit as long as they discharge their duties reasonably, may have hindered plaintiffs in many types of litigation against the police (Baude 2018). By effectively raising plaintiffs' burden of proof, these changes may have reduced the expected value of lawsuits and thus the threat value of lawsuits to claimants who hoped to settle claims without litigation. Prospective claimants, therefore, may have become less likely to file marginal claims. Unlike the new video technologies, however, the expansion of qualified-immunity protections would have had no effect on the quality or emotional appeal of evidence for claims ultimately brought. In other words, claims that were filed notwithstanding the heightened barriers to suit were of no greater value after than before the doctrinal shifts. If this explanation is correct, we would expect that (1) the number of claims and (2) lawsuits should gradually decline, as prospective claimants withhold marginal claims and suits; (3) claimants' win rate and (4) average payouts should gradually rise, as average claim quality increases; and (5) total payouts should decline, as marginal claims are withheld and the value of successful claims is unchanged.

There are three reasons to reject this explanation. First, the number of lawsuits did not gradually decline. In Figure 5, we plot the number of lawsuits over time alongside the number of successful lawsuits (that is, lawsuits with positive payouts) and the overall number of claims. We also include smoothed trend lines. Contrary to expectations, the number of lawsuits has been steadily climbing since 2005.[15]

Second, the large spike in average payouts in 2014 (Figure 3*B*) is also inconsistent with expectations. As stated, qualified-immunity doctrine is

---

reported to the insurer. Our data set contains only six excessive-force claims regarding incidents that allegedly occurred (and thus could have been filmed) in 2015. Among other things, these claims could not be driving the recent jump in median damages we observe in Figure 4.

15. Figure 5 also suggests that number of claims is a better measure than number of lawsuits of the incidence of policing harms. If we assume that the number of claims proxies for the number of policing harms, there are several reasons why the number of lawsuits might rise and fall (and rise again) during the observation period. The behavior of the plaintiffs' bar, which has historically shifted its attention among litigation areas, could explain this trend. But if we assume instead that the number of lawsuits proxies for the number of policing harms, it is significantly more difficult to explain why the number of claims moves as it does.
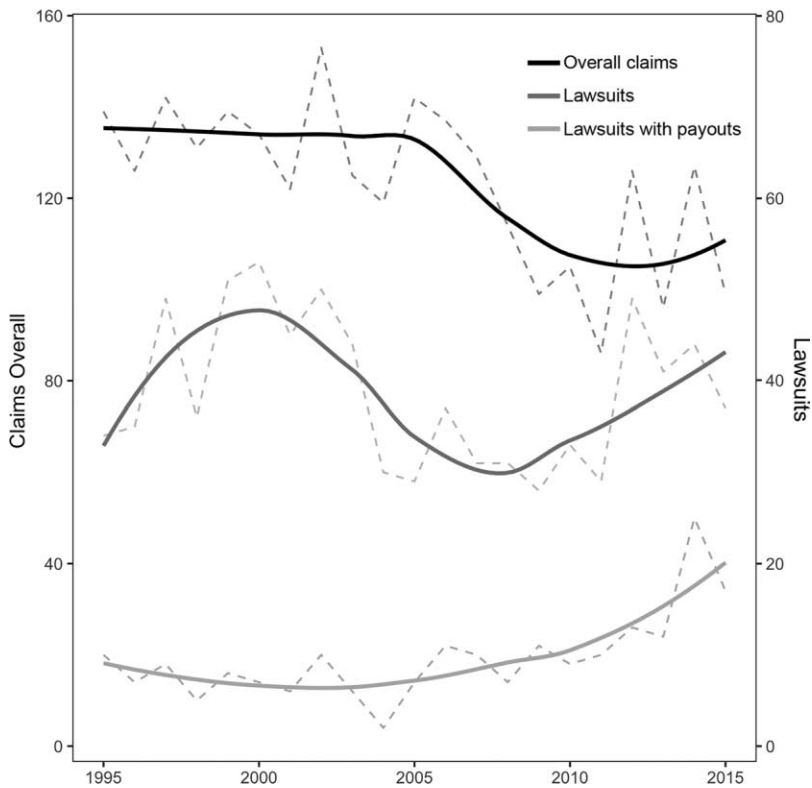
Figure 5. Number of lawsuits and claims

generally understood to have gradually raised the bar for plaintiffs' success. No single judicial decision since 1995 stands out as a game changer (see Baude [2018, pp. 88–90] for a comprehensive list of the Supreme Court's qualified-immunity decisions). Finally, and once again contrary to expectations, total payouts are clearly trending upward (Figure 3A).

4.2.2. Increasing Severity of Harm. A second potential explanation for the trends presented in Section 4.1 is that, while the frequency of policing harms decreased during our observation period, the severity of harms increased. It could be that police are interacting with civilians less often but that interactions are more volatile because police-community relations are deteriorating and the police are increasingly militarized (Delehanty et al. 2017; Lawson 2018). If this is the case, the severity of claims should increase, either across or within claim types, especially after the events of 2014, and average payouts should rise, as harms are increasingly severe.

This explanation could be consistent with multiple patterns in total payouts or claimants' win rate.

Our data do not permit us to reject this explanation with certainty, but we offer several reasons for skepticism. First, external evidence casts doubt on the motivating premise. While survey findings do suggest that confidence in the police declined between 2013 and 2015 (from 57 to 52 percent nationwide), during our sample period more generally, public confidence was fairly stable (ranging between 52 and 64 percent) (Jones 2015). There is no sign, in other words, of a persistent decline in police-community relations. Nor does existing evidence support the assumption that a deteriorating relationship would aggravate policing harms. Multiple studies have found that, if anything, the police backed off, rather than leaned in, after Ferguson (Morgan and Pally 2016; Shjarback et al. 2017). Over three-quarters of police officers surveyed in 2016, for example, indicated that Ferguson and other high-profile events have made them more reluctant to use force when it is appropriate (Morin et al. 2017).

Second, we find no evidence that the severity of claims increased, either across or within claim types. As an initial matter, recall that Figure 2 shows no clear trends over time in the substantive composition of claims. It is not the case, for example, that claims for property damage are consistently growing less common while excessive-force claims are growing more so. If claims are becoming more severe, it is happening within type in unobservable ways. One way to approach the question is to examine claims involving fatalities, many (though not all) of which concern severe examples of use of force. If within-type severity is increasing, we might expect fatalities to increase as well. Yet Figure 2 confirms that the number of fatalities is not rising over time. Our data on this point are consistent with external, national data sources that show a high, but stable, number of police-related civilian deaths each year (Mapping Police Violence 2019). Furthermore, in Figure A4, we show that there is no increase in the number of claims whose descriptions mention officer-involved shootings (identified by searching the text of the loss-description field) or vehicular pursuits.

4.2.3. Changing Societal Attitudes. Our final potential explanation for the trends we observe, and the one we find most compelling, is that society grew increasingly intolerant of harmful policing, especially after the spike in public scrutiny and media coverage precipitated by the events of 2014 (Figures A1 and A2). Juries may have become more inclined to is-

sue large damage awards on finding liability. They may also have grown increasingly willing to believe plaintiffs and disbelieve police officers in close cases or to deny police the benefit of the doubt. This would be consistent with other evidence that public behavior—such as calls to 911—changes markedly in response to highly publicized cases of police violence (Desmond, Papachristos, and Kirk 2016).

Of course, civil juries resolve a very small number of claims; most claims are settled out of court. But negotiation takes place in the jury's shadow (Galanter 1990). Multiple empirical studies confirm the influence of anticipated jury verdicts on settlement decisions by insurance claims adjusters (Ross 1970; Metzloff 1991). For example, Metzloff (1991, p. 88), writing in the medical malpractice setting, concludes, "While the number of litigated outcomes is modest, the jury is a real and present influence." Our conversations with the insurer, moreover, suggest that these effects matter in its business. For example, one claims adjuster reported that, when assessing claims, adjusters consider how a jury would respond even though the probability of a jury trial is very small.

Several empirical patterns would be consistent with this explanation: the severity of claims should not increase, as police behavior is not worsening; claimants' win rate should rise, as jurors increasingly favor plaintiffs in close cases; and average payouts should increase, especially around 2014, as jurors return higher damages awards on finding liability. No firm predictions about the number of claims or total payouts are warranted.

Our data bear out all of these predictions. As discussed above, we find no evidence of increasing claim severity. Figure 3C, meanwhile, shows a rising claimant win rate and rising average payouts with a spike in 2014 (Figure 3B). There is also external evidence consistent with this explanation. In Philadelphia, for example, the average damages for police shooting cases—the severity of which cannot rise over time—more than tripled between 2010 and 2014 (Elinson and Frosch 2015). And while anecdotal in nature, we think it not insignificant that insiders on all sides of the debate—insurers, plaintiffs' lawyers, city attorneys, and law enforcement experts—agree that "current scrutiny of law enforcement" and "shifting attitudes toward police" are affecting the direction and magnitude of case outcomes (Elinson and Frosch [2015, p. A8]; see also Winton [2017, pp. A1, A14], who notes, "Jurors are now less likely to give law enforcement the benefit of the doubt and more likely to award larger sums to plaintiffs, driving up the cost of judgments and emboldening attorneys to seek larger settlements during negotiations, experts said").

4.3. The Importance of Data Selection

4.3.1. Claims Data versus Lawsuit Data: Substantive Composition and Payout Magnitude.   Given that insurance claims data may be difficult to obtain, it is worth examining the extent to which other data sources might substitute. Figure 5 suggests reasons for caution on this point, as the trends in claims, lawsuits, and lawsuits with payouts differ markedly. Indeed, if, like prior researchers, we had used lawsuits or lawsuits with payouts as a proxy for police-inflicted harms over the past decade or so, we would have reached the opposite conclusion: we would have concluded, that is, that policing harms were increasing rather than decreasing. In this section, we investigate whether these data sources—claims, lawsuits, and lawsuits with payouts—differ along additional dimensions. We find that they do.

We first compare the universe of all claims with the subsets of claims that are commonly employed in academic research. In Figure 6A, we present the distribution of types of claims for all claims (1995–2015) and the 32 percent of claims that involve lawsuits. Use of force and illegal seizure of a person generate the majority of lawsuits, while property claims are much more common overall. Lawsuit data, in other words, drastically undercount the incidence of property harms. Payouts also vary between claims that do and do not go to court (Figure 6C). Average payouts are significantly higher for excessive-force and sexual misconduct claims that involve a lawsuit, for example. This does not mean that lawsuits cause higher payouts because of selection. It does mean, however, that estimates of average payouts calculated using lawsuit data will be biased upward, as we discuss in more detail below. Claims that go to court differ significantly from those that do not in both type and severity.

Many papers that study police liability match lawsuit data—which rarely includes settlement amounts—with payout data (Iris 2012, 2014; Rozema and Schanzenbach 2019; Schwartz 2014, 2016). Requiring both a lawsuit and a positive payout narrows the sample to only 8 percent of the claims in our data set, as Table 2 shows. As Figure 6 demonstrates, this small subset of claims differs from the larger sample in both substantive composition (Figure 6B) and average payouts made (Figure 6D). Again, use of force and illegal seizure of a person are overrepresented and property harms underrepresented relative to the overall sample. The distortion is even greater when considering average payouts. Without any $0 claims, and with presumably very few low-value claims because of the
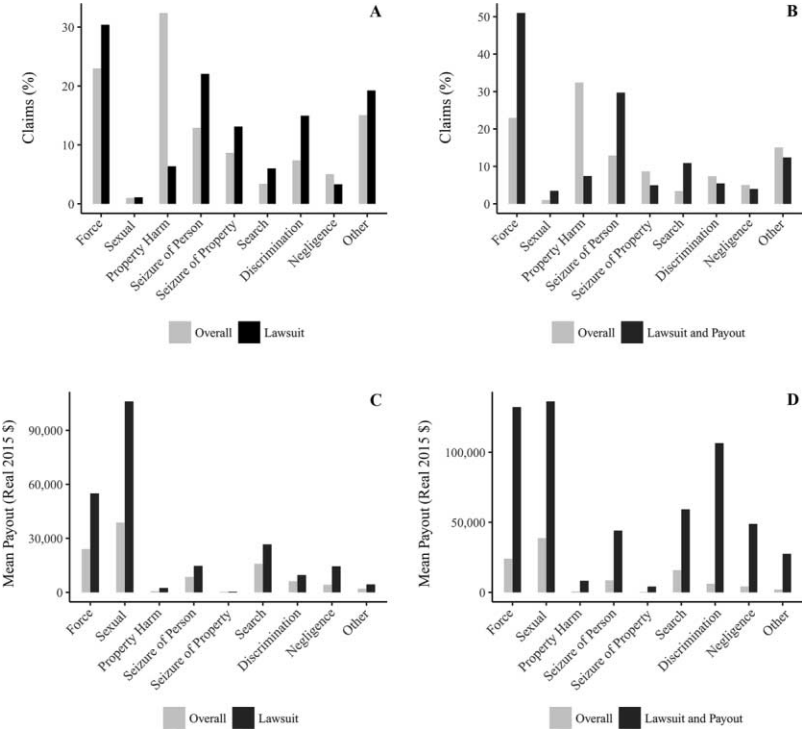
Figure 6. Comparison of claims

costs of litigation, the average payout value in this subset is more than 10 times what it is in the overall sample of claims. Furthermore, because many claims involve neither a lawsuit nor a payout, even data containing all lawsuits (regardless of payout) and all payouts (regardless of lawsuit) would miss a large fraction of allegations: on average, for overall claims, 46 percent. And this is not the case only for minor infractions: for example, 58 percent of sexual misconduct and 53 percent of excessive-force claims lead neither to a lawsuit nor to a payout.[16]

One might argue that property harms, the least egregious among the claim types, are not worth much attention. If we disregard these minor claims, do claims data still differ systematically from the other data types?

---

16. Some of these claims may be frivolous, and so focusing on those that make it through the legal system selects for claims with legal merit. Given our results in Section 4.2, however, another possible interpretation is that, in earlier periods, the public was more trusting of the police and tolerant of the harms they cause.

Table 2. Claims by Type

	Payout		No Payout			
Type of Claim	Lawsuit (%)	No Lawsuit (%)	Lawsuit (%)	No Lawsuit (%)	Claims (N)	Fraction of All Claims
Force	17	6	25	53	595	.23
Sexual	31	15	12	58	26	.01
Property harm	2	5	5	42	839	.32
Seizure of person	18	9	36	38	334	.13
Seizure of property	5	15	44	38	224	.09
Search	26	11	32	35	88	.03
Discrimination	6	0	58	36	191	.07
Negligence	7	26	15	55	130	.05
Other	7	4	34	56	390	.15
Overall	8	22	24	46	2,590	

In unreported analysis, we reproduce several of our key figures after dropping property claims from the data set. We find that, when we omit property claims, the substantive composition of claims does converge for claims and lawsuits (as well as lawsuits with payouts, though to a lesser extent). At the same time, the average magnitudes of payouts across claim types continue to differ substantially when we compare claims to lawsuits or lawsuits with payouts. Likewise, the aggregate time trends in number of claims, lawsuits, and lawsuits with payouts remain distinct when we omit property claims.

In addition, even though property claims result in low average payouts and involve no physical harm to civilians, they do shed light on civilian interactions with the police. Each property claim represents an occasion on which a civilian believes a law enforcement officer has damaged or destroyed his property. Even a relatively low-value claim—a damaged door or window, a broken smartphone, or a scratched vehicle, for example—may reflect a substantial monetary shock, especially for poorer households. Moreover, frequent property damage by officers may be interpreted as a sign of disrespect that can breed distrust in the community. Research consistently finds that people form views about police legitimacy on the basis of not only their own interactions with officers but also what they hear from their neighbors and friends (Mazerolle et al. 2013). Yet precisely because property claims are underrepresented in other data sources, researchers likely undercount the frequency of property damage

and all that follows from it. Claims data helpfully bring these incidents to light.

4.3.2. Selection of Disputes for Adjudication. Even if lawsuits and lawsuits with payouts are not representative of the broader universe of claims, one recent study suggests that researchers might use adjudicated claims to proxy for the whole. Adjudicated claims are the subset of claims that are resolved by a dispositive motion or trial verdict—all of which necessarily involve lawsuits—regardless of whether the judgment favored the plaintiff or the defendant. In recent empirical work using New York litigation data from contingent-fee cases, Helland, Klerman, and Lee (2018) find that the average judgment amount for adjudicated claims is very close to the average settlement, and the distributions of settlements and adjudicated damages are quite similar. This is a surprising finding that runs contrary to the influential Priest-Klein hypothesis, which holds that "the disputes selected for litigation (as opposed to settlement) will constitute neither a random nor a representative sample of the set of all disputes" (Priest and Klein 1984, p. 4).

To see whether the finding of Helland, Klerman, and Lee (2018) holds in our distinct institutional environment, we replicate the comparison they draw between adjudicated and settled claims (see Helland et al. [2017] for a detailed description of their data), using the same definitions of these terms as in their paper. Adjudicated claims, as stated, include claims that are resolved by motion or trial verdict. Settled claims include all claims that are not adjudicated but that result in some payout being made, regardless of whether a lawsuit was filed—in other words, all negotiated payments, not just those that resolve a lawsuit. Following Helland, Klerman, and Lee (2018), our specification of settled claims does not include unadjudicated claims that lead to no payout for whatever reason, such as claims that were voluntarily dismissed or abandoned by claimants. There are no $0 settlements, that is. Settled claims also exclude claims for which the insurer denied coverage.

In Figure 7, we compare the substantive composition and average payouts for settled and adjudicated claims. The substantive composition of claims varies between the adjudicated and settled samples (Figure 7A). Knowing what share of adjudicated claims involve the seizure of property, for example, does not allow one to estimate the share of settled claims of that same type. It does not appear that Helland, Klerman, and Lee (2018) were able to disaggregate their data by substantive
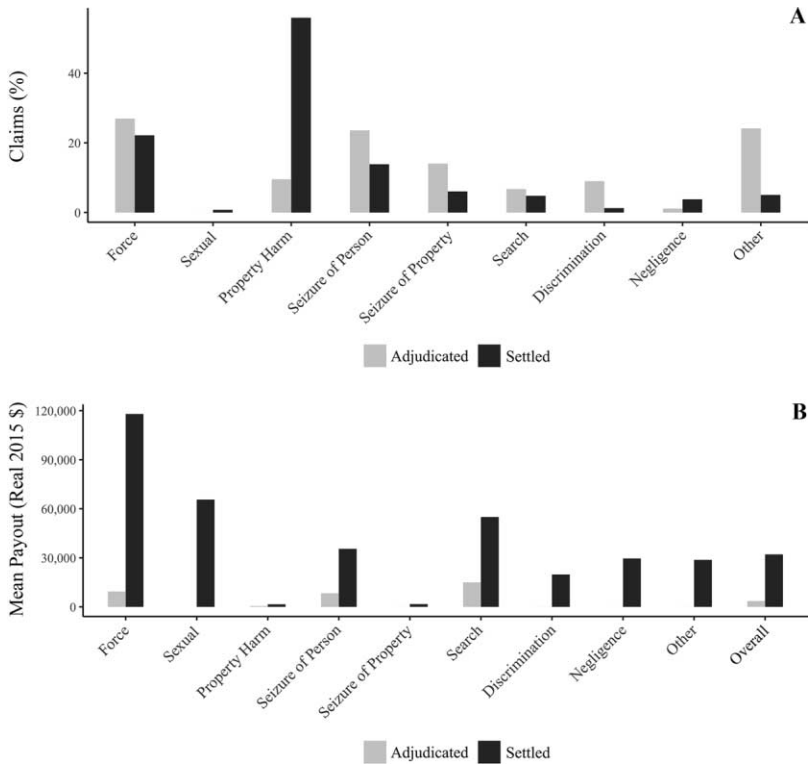
Figure 7. Comparison of adjudicated and settled claims

case type, so it is unclear whether their data would have resembled ours in this respect. Likewise, we do not replicate the payout-related findings of Helland, Klerman, and Lee (2018): the average judgment amount in our data ($3,562) is an order of magnitude lower than the average settlement ($32,072; Table 3), and the difference is statistically significant ($p < .001$). At least in our sample, then, a researcher without access to settlement data cannot estimate the average settlement value by reference to the average adjudicated judgment. The same is true for claims of different substantive types, such as excessive force (Figure 7B). The extent to which these results will hold in other police liability data sets is unclear.

The contrast between our results and those of Helland, Klerman, and Lee (2018) may reflect differences in the institutional environments we study. Our data contain exclusively claims against public law enforcement agencies; theirs encompass contingent-fee cases involving personal

injury, property damage, and wrongful death, as well as condemnation and change of grade (Helland et al. 2017). Although, as noted, Helland, Klerman, and Lee (2018) do not disaggregate their data by case type, their contingent-fee cases very likely include a significant number of tort suits against private defendants. It may be that the defendants in our data set more frequently face asymmetric stakes in litigation. Public defendants, for example, may incur greater reputational costs from adverse judgments and, as repeat players, have incentives to play for rules that will govern future disputes (Galanter 1974). Where defendants have more at stake than plaintiffs from adverse judgments, defendants can be expected to settle all but the cases they are likeliest to win, which drives down the expected plaintiff win rate (Priest and Klein 1984) and potentially pushes apart the average judgment and average settlement amounts. Indeed, in the data in Helland, Klerman, and Lee (2018), plaintiffs won 29 percent of adjudicated cases; in ours, plaintiffs win only 3.9 percent (authors' calculations, not reported). If we were to assume a 29 percent win rate in our own data and assume (perhaps unrealistically) that average judgments would be the same for those 29 percent as they are for the 3.9 percent, then the overall average judgment amount would be around $26,000—substantially closer to the average settlement amount of $32,072.

We also observe that the difference between the average judgment and average settlement is highly sensitive to how we specify claims as settled. As explained above, Helland, Klerman, and Lee (2018) exclude all unadjudicated claims that result in no payout (Settled 1 in Table 3), which constitute 13 percent of their cases. They acknowledge that "such cases could be classified as settlements for $0" but argue that they "probably reflect nuisance suits the defendant refused to settle or cases where investigation or discovery revealed that the plaintiff's case was very weak" (Helland, Klerman, and Lee 2018, p. 146). If, however, we recode unadjudicated, $0 claims as settled for $0 (Settled 2 in Table 3)—which seems to us an equally plausible way of interpreting the data—the average judgment ($3,562) and average settlement ($13,574) grow closer, but the difference between them remains statistically significant ($p = .006$).

Similarly, Helland, Klerman, and Lee (2018) had the somewhat unusual opportunity—as we do—to observe settlements for which no lawsuit was ever filed. In many settings, researchers will never learn of such disputes. With this in mind, if we restrict the universe of claims to those in which a lawsuit was filed—treating settlement as settlement of a law-

Table 3. Adjudicated and Settled Claims

Type of Claim	Percentage of Claims				Mean Payouts (Real 2015 $)			
	Adjudicated	Settled 1	Settled 2	Settled 3	Adjudicated	Settled 1	Settled 2	Settled 3
Force	27	22	22	56	9,416	117,952	49,193	141,208
Sexual	0	1	1	2		65,602	21,867	65,602
Property harm	10	56	42	4	483	1,553	882	21,301
Seizure of person	24	14	14	34	8,334	35,500	15,019	42,372
Seizure of property	14	6	6	6	0	1,663	713	4,434
Search	7	5	4	11	14,981	54,950	26,101	69,910
Discrimination	8	1	7	4	0	19,795	1,596	19,795
Negligence	1	4	6	3	0	29,594	8,376	95,178
Other	24	5	10	11	0	28,732	6,315	39,371
Overall (N)	177	397	938	124	3,562	32,072	13,574	95,307

Note. The Settled 1 column follows Helland, Klerman, and Lee (2018) and includes only unadjudicated claims with positive payouts. The Settled 2 column includes all unadjudicated claims, regardless of whether payouts were made. The Settled 3 column includes only unadjudicated claims in which a lawsuit was filed. Percentage values do not sum to 100 because some claims have multiple claim types. Data are from claims opened after 2005, when information about claim resolution is available.

suit (Settled 3 in Table 3)—the average judgment ($3,562) and average settlement ($95,307) instead diverge even further ($p < .001$). Table 3 compares the composition and average payout figures for adjudicated claims with those of settled claims defined in these three different ways.

5. DISCUSSION

In Section 5.1, we comment on some implications of our research. We then review the principal limitations of our data and analysis in Section 5.2.

5.1. Research Implications

Three implications of our analysis deserve mention. First, in our sample, as we move up the pyramid from claims to lawsuits to payouts, the composition of cases changes substantially. Claims that involve litigation differ systematically from claims that do not, in terms of both type and severity, and adjudicated claims differ systematically from settled claims. All of this is consistent with the Priest-Klein hypothesis but is in contrast to the recent empirical contribution by Helland, Klerman, and Lee (2018) to the Priest-Klein hypothesis literature. More research is warranted, for example to explore whether there are particular institutional environments in which the Priest-Klein hypothesis is more or less likely to hold. Future research can assess whether selection effects will manifest similarly, in terms of direction and magnitude, even in other policing data sets. For now, we recommend caution when drawing inferences from litigation data about the composition or severity of policing harms.

Second, and relatedly, having demonstrated the value of insurance claims data to the academic study of policing, we hope future researchers will pursue this data source and additional insurers will consider making their data available for study. By analogy, researchers have leveraged to great effect the wider availability of aggregate data on closed claims for medical malpractice (Black et al. 2005; Vidmar et al. 2005). Given the social salience of contemporary American policing, the value of wider access to claims data could be significant. It would shed light not only on other jurisdictions but also on the behavior of other insurers.

Third, in the immediate term, we can use our data to generate multipliers that estimate the numerical relationships among claims, lawsuits, and lawsuits with payouts. As an illustration, we then use these multi-

pliers to benchmark prior studies. A number of recent papers, which focus on civil rights cases, match lawsuit data with payout data and report case counts and average payout figures (see, for example, Iris 2012, 2014; Rozema and Schanzenbach 2019). Using the multipliers generated by different cuts of our own data, we show how those estimates might be adjusted to reflect differences among data sources. In principle, future researchers could use these multipliers in similar ways on other data sets.

For this benchmarking exercise, to mirror the prior literature we limit our data to civil rights claims. We approximate this category by removing from our data set property damage, negligence, and unclassified (other) claims. In Table 4, we report the number of civil rights claims and average payout per claim, in 2015 dollars, for three different splits of the data: all claims, claims involving a lawsuit, and claims involving a lawsuit and a positive payout. In Table 5, we present ratios of claim counts and average payout. The idea is to estimate the extent to which data on lawsuits with positive payouts may lead researchers to understate the number of claims for compensation and overstate average payouts relative to data on all claims or all lawsuits.

The values in Tables 4 and 5 indicate that, had we used data on claims that led to both a lawsuit and a positive payout, like much of the existing literature, we would have painted a very different picture of the volume and consequences of demands for compensation from the police. There are many fewer claims that result in lawsuits with payouts, or even just in lawsuits, than there are overall; conversely, average payouts are much larger for claims with lawsuits and payouts than for all lawsuits or all claims. For example, we would have undercounted by a factor of 8.3 (= 1/.12 in the first column of Table 5) the number of demands for compensation from the police. At the same time, we would have overstated the average payout per demand by a factor of 7.6 had we used data on lawsuits that led to positive payouts instead of data on all claims. Similar calculations can be used to compare lawsuits with payouts with all lawsuits.

Ratios like these could inform how we interpret other research. For example, using data from Chicago on civil rights lawsuits that paid out between 2006 and 2012, Iris (2014) reports an average payout figure of $112,335. If the ratios of payouts are similar in Chicago and our context, then the average payout including lawsuits that did not result in payouts would be closer to $30,400 and more like $14,800 including claims that never ripened into lawsuits. Similarly, Rozema and Schanzenbach (2019)

Table 4. Civil Rights Claims, Lawsuits, and Average Payouts

	Claims	Mean Payout (Real 2015 $)
All claims	1,278	12,970
Claims with lawsuits	592	26,440
Claims with lawsuits and positive payouts	159	98,440

Table 5. Ratios of Civil Rights Claims and Payouts

	Claim Counts	Payouts
Claims with lawsuits and positive payouts/all claims	.12	7.6
Claims with lawsuits and positive payouts/claims with lawsuits	.27	3.7

find that, between 2009 and 2014, the average payout for civil rights lawsuits in Chicago was $329,322 for lawsuits that led to positive payouts. We might then predict that, for all lawsuits (all claims), the average payout would be closer to $89,000 ($43,000).

We emphasize that these ratios are offered principally to illustrate a concept. It is entirely possible that different ratios characterize different settings. Indeed, Iris (2014) suggests that, in Chicago, very few civil rights claims are brought outside of court. Perhaps residents of smaller, more tight-knit communities—like many of those in our data set—are more hesitant to file lawsuits or seek large recoveries that could burden a modest tax base (Black 1972). Even within a single jurisdiction, different insurers might generate different ratios as well, depending on factors including settlement strategy. Indeed, we find that, even in our own data set, the ratios change over time. We calculated number-of-claims and average payout ratios for different 5-year time windows in our data (results available from the authors on request). In the most recent 5-year window, the average payout ratios are smaller than in earlier years: payouts for civil rights claims with lawsuits and positive payouts are 2.5 times higher than for all civil rights claims with lawsuits and 4.4 times higher than payouts for all civil rights claims. Even these smaller ratios, however, indicate large differences in the magnitude of payouts depending on the data set employed.

5.2. Limitations

The data we employ are rich and closer to the base of the pyramid, so to speak, than lawsuit or payout data. That said, there remain important limitations to our data that should be kept in mind when interpreting or extending our results.

First, we have data only on incidents in which an individual demanded, or was anticipated to demand, financial compensation for a covered police-related harm. This means our data do not include what might be called victimless police malfeasance, such as an officer using a controlled substance while on duty. Nor do they capture police behavior like discourtesy that, while offensive or undesirable in some respects—and while not victimless—creates no plausible claim of compensable harm. Claims data also omit misconduct that is clearly excluded under the insurer's policy, such as acts committed outside an officer's official capacity.

Second, we observe only incidents that were reported to the insurer. Localities may resolve some covered police-related incidents without reporting them, possibly in a perceived effort to control their premiums. Our data set would not capture these occurrences. These incidents, however, are likely rare and minor in magnitude. Policyholders are contractually obligated to report each covered occurrence to the insurer promptly. Failing to report even a minor occurrence, which could result in the denial of coverage, is risky because seemingly small claims can escalate. Premiums are experience rated, but they are not sensitive enough to respond meaningfully to any particular police liability claim. The insurer's experience-modification adjustment is calculated across several years of claims, over all lines of liability—not just policing—and caps the impact of large losses. Moreover, many localities carry an aggregate deductible, which creates an additional incentive to report: if the locality is going to pay out, it might as well get credit for the payment against its deductible so that coverage kicks in sooner for future claims. The insurer also assists policyholders in evaluating and processing claims, executing adequate releases, and dealing with the tax implications of payments, complications localities would need to handle independently if they declined to report claims.

Third, all of our data come from a single insurer in a single midsized state, which we have agreed not to identify, which raises questions about external validity. Part of what we are measuring is the behavior of this

particular insurer (see Schwartz, forthcoming). The decline in claims we observe, for example, may stem in part from the insurer's efforts to reduce policing harms (and therefore claims), efforts that other insurers may not match. The state we study, moreover, is relatively racially homogeneous, especially outside of the largest municipalities, several of which self-insure and thus are not included in our data. It is possible that trends in policing differ in more heterogeneous settings. Kane and White (2013) also suggest the possibility of regional variation in the types and prevalence of police misconduct. At the same time, most law enforcement agencies in the United States are small, and the populations in small-town and rural communities are heavily white—78 percent, compared with 64 percent in the country overall and 44 percent in urban areas (Housing Assistance Council 2012). Nor are policing harms confined to big cities. Recall that roughly two-thirds of police-related fatalities occur outside the nation's 100 largest cities; relatedly, roughly half of all individuals killed by the police are white (Mapping Police Violence 2019). There are also distinct methodological advantages to working with a single state. State-level effects—including state law and other state characteristics—are held constant. And where, as here, a single entity collects all of the data, data-consistency concerns are greatly reduced.

Fourth, we cannot directly measure the rate of police-inflicted harms; there is selection into claims. Insurance claims, that is, reflect the behavior of claimants and (in some instances) their lawyers as well as that of the police. The decision whether to make a demand depends on factors including the putative claimant's perception of whether the grievance is worth pursuing and her knowledge of the law and legal institutions. The relatively stable trend we observe in the number of claims per year, for instance, could reflect stability in police behavior. We would see the same trend, however, if policing harms increased (decreased) but selection into claims decreased (increased). In Section 4.2.1, we attempted to identify and rule out the most obvious explanations for why selection into claims may have changed over time, and we are not aware of any other well-motivated explanation. But we cannot rule out the possibility. Nevertheless, nearly every familiar data source, especially when it comes to crime and victimization, suffers from this shortfall.[17]

Finally, insurance claims data may be proprietary and thus difficult

---

17. One possible exception is field observation data (see, for example, Mastrofski, Snipes, and Supina 1996; Terrill and Reisig 2003). Field observation presents its own difficulties, however. It is difficult to scale and prone to reactivity bias.

to obtain. Some civilian-complaints data, in contrast, and all lawsuit data are in the public record. Still, especially for longitudinal or cross-jurisdictional purposes, these other data sources are not as readily accessible as it may seem. As we discussed in Section 2.1, civilian-complaints data may be unavailable or unreliable especially in smaller cities, and many jurisdictions purge these data after a few years. Gathering lawsuit data on police misconduct across any significant span of time or space is also challenging and prone to measurement error. Rozema and Schanzenbach (2019) describe the process, which involves key-word searches followed by manual review, one jurisdiction at a time.

6. CONCLUSION

Motivated by sharply heightened public scrutiny of the police in recent years, we examined a new data source, insurance claims data, in an effort to disentangle trends in police behavior from the responses of legal actors such as lawyers, judges, and juries. We found that, over time, the number and substantive composition of claims have held relatively stable in our sample, while the claimant win rate, total and average payouts, and (more recently) number of lawsuits all have increased. After considering several potential explanations, we interpret these trends to suggest that the incidence of policing harms has not been rising in our sample, nor has the substantive nature of those harms been changing in observable ways. Instead, society has likely grown increasingly intolerant of policing harms, which has led to more claimant victories and larger recoveries as claims are settled in the jury's shadow. The basic methodological takeaway is that existing data sources capture vastly different types and amounts of harmful policing.

The policy implications of our findings are less clear. Rising payouts per claim, we have argued, seem to signal increasing societal awareness of, and concern about, harmful policing. This could be an optimistic story about the power of protest and public influence campaigns if rising payouts are leading agencies to pursue productive reforms. But it also points to the critical role of the insurer, which, by converting payouts to premiums, mediates the relationship between payouts and incentives for reform. Some insurers, for example, may pin premiums tightly to payouts, which creates the sort of behavioral incentives proponents of civil liability imagine. Others, however, may prioritize rate stability as a service

to their policyholders, which raises the possibility of moral hazard, as rising payouts fail to translate into rising costs for policyholders. Further research might investigate the relationship between police liability claims experience and subsequent premiums and, even more important, between premiums and subsequent police behavior.

We conclude with two parting disclaimers. First, our message is not that policing is self-improving—just give it time, and the incidence of harm will fall. While our findings do not bear on the question, the decline in claims may have resulted from policy efforts on all sides of the issues, including but not limited to the insurer's loss-prevention initiatives. Falling crime and arrest rates may have helped as well. Second, we do not suggest that, because claims are declining, they are necessarily low enough. They could be declining, that is, but still be higher than the socially optimal level. To dramatize the point, homicide rates have been falling for many centuries, if not millennia, and "today we may be living in the most peaceable era in our species' existence" (Pinker 2011, p. xxi). But few who are living in high-crime neighborhoods in America's cities would report that all is well. Progress does not warrant complacency.
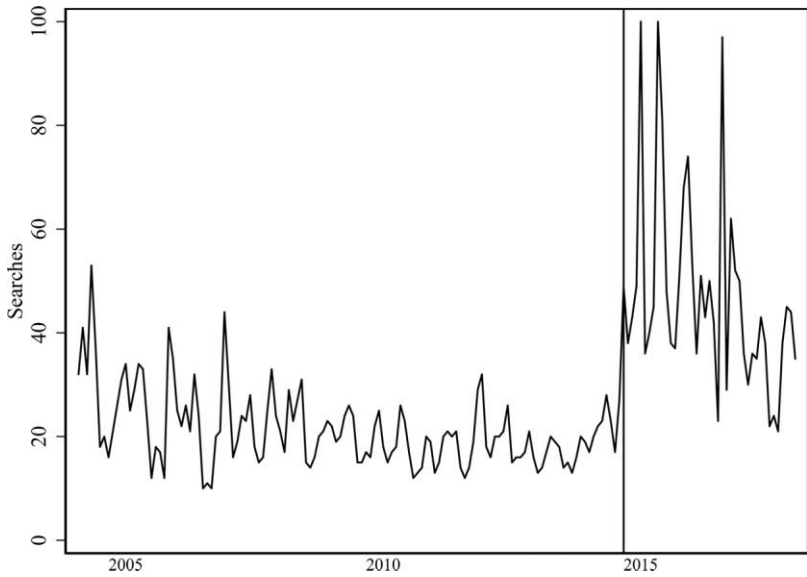
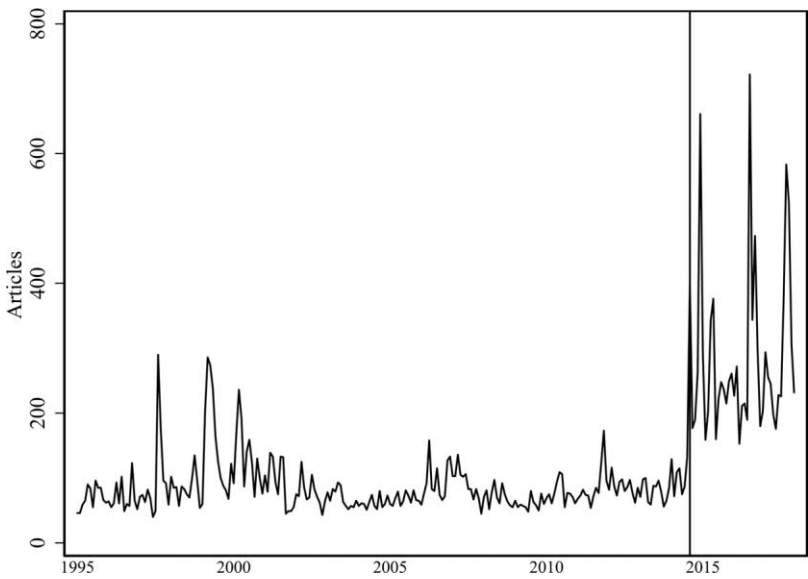**Figure A1.** Monthly Google searches for "police brutality"

Figure A2. Monthly newspaper articles mentioning police brutality
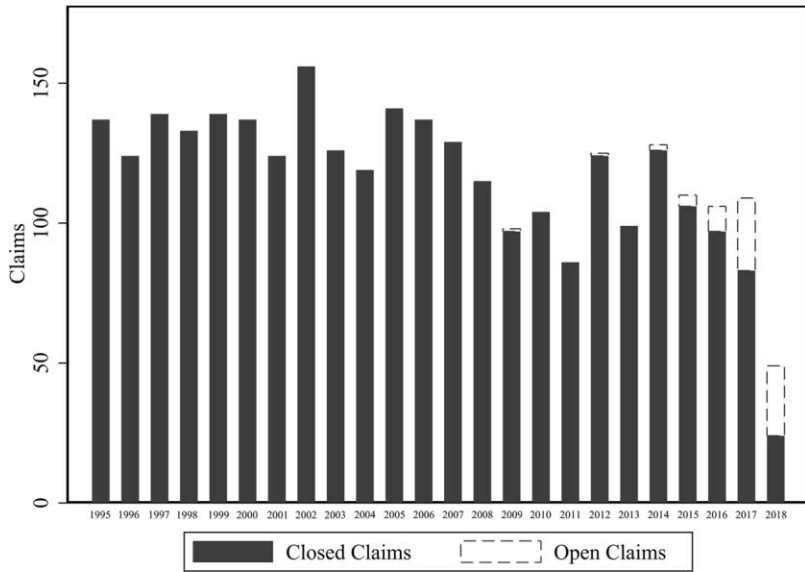
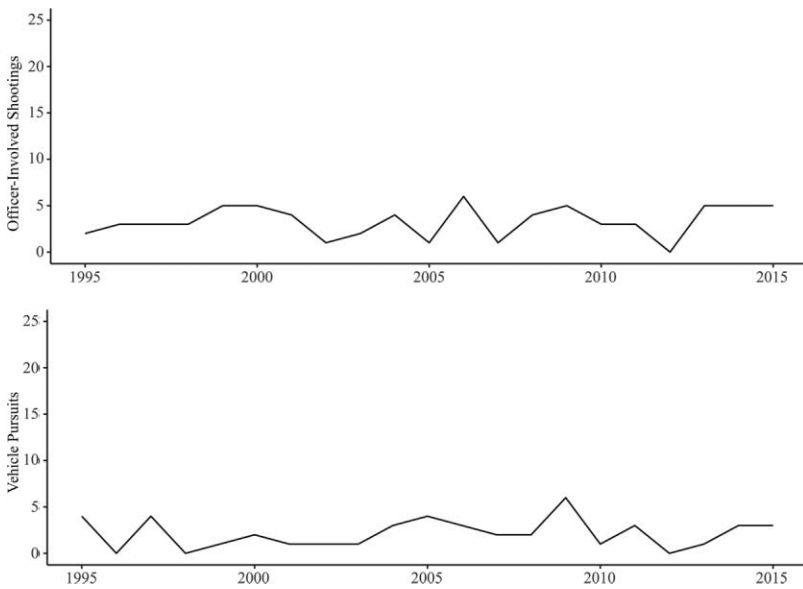**Figure A3.** Open and closed claims by year

Figure A4. Claims mentioning officer-involved shootings or vehicle pursuits

REFERENCES

Anwar, Shamena, and Hanming Fang. 2006. An Alternative Test of Racial Prejudice in Motor Vehicle Searches: Theory and Evidence. *American Economic Review* 96:127–51.

Ba, Bocar A. 2018. Going the Extra Mile: The Cost of Complaint Filing, Accountability, and Law Enforcement Outcomes in Chicago. Working paper. University of Chicago, Harris School of Public Policy Studies, Chicago.

Barrett, Kevin J., Maria (Maki) Haberfeld, and Michael C. Walker. 2009. A Comparative Study of the Attitudes of Urban, Suburban, and Rural Police Officers in New Jersey regarding the Use of Force. *Crime, Law, and Social Change* 52:159–79.

Baude, William. 2018. Is Qualified Immunity Unlawful? *California Law Review* 106:45–90.

Berger, Alan M., and Joel Kotkin, eds. 2017. *Infinite Suburbia*. With Celina Balderas Guzmán. New York: Princeton Architectural Press.

Black, Bernard, Charles Silver, David A. Hyman, and William M. Sage. 2005. Stability, Not Crisis: Medical Malpractice Claim Outcomes in Texas, 1988–2002. *Journal of Empirical Legal Studies* 2:207–59.

Black, Donald J. 1972. The Boundaries of Legal Sociology. *Yale Law Journal* 81:1086–1100.

Campbell, Bradley A., Justin Nix, and Edward R. Maguire. 2017. Is the Number of Citizens Fatally Shot by Police Increasing in the Post-Ferguson Era? *Crime and Delinquency* 64:398–420.

Chandrasekher, Andrea Cann. 2016. The Effect of Police Slowdowns on Crime. *American Law and Economics Review* 18:385–437.

———. 2017. Police Labor Unrest and Lengthy Contract Negotiations: Does Police Misconduct Increase with Time Spent out of Contract? Working paper. University of California, School of Law, Davis.

Chanin, Joshua. 2016. Evaluating Section 14141: An Empirical Review of Pattern or Practice Police Misconduct Reform. *Ohio State Journal of Criminal Law* 14:67–112.

Crow, Matthew S., and Brittany Adrion. 2011. Focal Concerns and Police Use of Force: Examining the Factors Associated with Taser Use. *Police Quarterly* 14:366–87.

Delehanty, Casey, Jack Mewhirter, Ryan Welch, and Jason Wilks. 2017. Militarization and Police Violence: The Case of the 1033 Program. *Research and Politics* 4(2):1–7.

Desmond, Matthew, Andrew V. Papachristos, and Davis S. Kirk. 2016. Police Violence and Citizen Crime Reporting in the Black Community. *American Sociological Review* 81:857–76.

Dharmapala, Dhammika, Richard H. McAdams, and John Rappaport. 2019. Collective Bargaining Rights and Police Misconduct: Evidence from Florida.

Working paper. University of Chicago Law School, Chicago.

Eisenberg, Theodore. 1982. Section 1983: Doctrinal Foundations and an Empirical Study. *Cornell Law Review* 67:482–556.

Elinson, Zusha, and Dan Frosch. 2015. Police Misconduct Costs Soar. *Wall Street Journal,* July 15, pp. A1, A8.

Falcone, David N., L. Edward Wells, and Ralph A. Weisheit. 2002. The Small-Town Police Department. *Policing: An International Journal* 25:371–84.

Fox, Justin. 2017. Violent Crime Isn't Just a Big-City Thing. *Bloomberg,* September 25. https://www.bloomberg.com/view/articles/2017-09-25/violent-crime-isn-t-just-a-big-city-thing.

Fryer, Roland G., Jr. 2019. An Empirical Analysis of Racial Differences in Police Use of Force. *Journal of Political Economy* 127:1210–61.

Galanter, Marc. 1974. Why the "Haves" Come Out Ahead: Speculations on the Limits of Legal Change. *Law and Society Review* 9:95–160.

———. 1990. The Civil Jury as Regulator of the Litigation Process. *University of Chicago Legal Forum,* pp. 201–71.

Garcia, Venessa, and Liqun Cao. 2005. Race and Satisfaction with the Police in a Small City. *Journal of Criminal Justice* 33:191–99.

Harmon, Rachel. 2013. Why Do We (Still) Lack Data on Policing? *Marquette Law Review* 97:1119–46.

Helland, Eric, Daniel Klerman, Brendan Dowling, and Alexander Kappner. 2017. Contingent Fee Litigation in New York City. *Vanderbilt Law Review* 70:1971–92.

Helland, Eric, Daniel Klerman, and Yoon-Ho Alex Lee. 2018. Maybe There Is No Bias in the Selection of Disputes for Litigation. *Journal of Institutional and Theoretical Economics* 174:143–70.

Hickman, Matthew J., and Jane E. Poore. 2016. National Data on Citizen Complaints about Police Use of Force: Data Quality Concerns and the Potential (Mis)Use of Statistical Evidence to Address Police Agency Conduct. *Criminal Justice Policy Review* 27:455–79.

Housing Assistance Council. 2012. *Race and Ethnicity in Rural America*. Rural Research Brief. Washington, DC: Housing Assistance Council. http://www.ruralhome.org/storage/research_notes/rrn-race-and-ethnicity-web.pdf.

Iris, Mark. 2012. Illegal Searches in Chicago: The Outcomes of 42 U.S.C. § 1983 Litigation. *Saint Louis University Public Law Review* 32:123–45.

———. 2014. Your Tax Dollars at Work! Chicago Police Lawsuit Payments: How Much, and for What? *Virginia Journal of Criminal Law* 2:25–61.

Jennings, Jay T., and Meghan E. Rubato. 2017. Preventing the Use of Deadly Force: The Relationship between Police Agency Policies and Rates of Officer-Involved Gun Deaths. *Public Administration Review* 77:217–26.

Jones, Jeffrey M. 2015. In U.S., Confidence in Police Lowest in 22 Years. *Gallup,* June 19. https://news.gallup.com/poll/183704/confidence-police-lowest-years

.aspx.

Kane, Robert J. 2007. Collect and Release Data on Coercive Police Actions. *Criminology and Public Policy* 6:773–80.

Kane, Robert J., and Michael D. White. 2009. A Study of Career-Ending Misconduct among New York City Police Officers. *Criminology and Public Policy* 8:737–69.

———. 2013. *Jammed Up: Bad Cops, Police Misconduct, and the New York City Police Department*. New York: New York University Press.

Kaplan, Jacob. 2017. Uniform Crime Reporting Program Data: Offenses Known and Clearances by Arrest, 1960–2016 (computer file). Ann Arbor, MI: Interuniversity Consortium for Political and Social Research.

Kappeler, Victor E. 2006. *Critical Issues in Police Civil Liability*. 4th ed. Long Grove, IL: Waveland Press.

King, Shaun. 2015. Police Brutality Is Getting Worse and Shows No Signs of Slowing Down. *Daily Kos,* September 22. https://www.dailykos.com/stories /2015/9/22/1423847/-Police-brutality-is-getting-worse-and-shows-no-signs-of -slowing-down.

———. 2018. Data Shows Police Brutality in America Is Getting Worse—2018 Could Be the Most Deadly in Years. *Appeal,* April 17. https://theappeal.org /data-shows-police-brutality-in-america-is-getting-worse-2018-could-be-the -most-deadly-in-years-90c9fa503580/.

Kotkin, Joel, and Alan M. Berger. 2017. The Future of America's Suburbs Looks Infinite. *Orange County Register,* November 19.

Lawson, Edward, Jr. 2018. Police Militarization and the Use of Lethal Force. *Police Research Quarterly* 72:177–89.

Legewie, Joscha, and Jeffrey Fagan. 2016. Group Threat, Police Officer Diversity, and the Deadly Use of Police Force. Public Law Research Paper No. 14-512. Columbia Law School, New York.

Leong, Nancy. 2012. Making Rights. *Boston University Law Review* 92:405–81.

Lersch, Kim Michelle, and Tom Mieczkowski. 2000. An Examination of the Convergence and Divergence of Internal and External Allegations of Misconduct Filed against Police Officers. *Policing: An International Journal of Police Strategies and Management* 23:54–68.

Levitt, Steven D., and Thomas J. Miles. 2006. Economic Contributions to the Understanding of Crime. *Annual Review of Law and Social Science* 2:147–64.

Liederbach, John. 2005. Addressing the "Elephant in the Living Room": An Observational Study of the Work of Suburban Police. *Policing: An International Journal of Police Strategies and Management* 28:415–34.

———. 2007. Controlling Suburban and Small-Town Hoods: An Examination of Police Encounters with Juveniles. *Youth Violence and Juvenile Justice* 5:107–24.

Lytle, Daniel J., and Ryan Randa. 2015. The Effects of Police Satisfaction on

Fear of Crime in a Semirural Setting. *International Criminal Justice Review* 25:301–17.

Mapping Police Violence. 2019. Mapping Police Violence. https://mapping policeviolence.org/nationaltrends.

Mas, Alexandre. 2006. Pay, Reference Points, and Police Performance. *Quarterly Journal of Economics* 121:783–821.

Mastrofski, Stephen D., Jeffrey B. Snipes, and Anne E. Supina. 1996. Compliance on Demand: The Public's Response to Specific Police Requests. *Journal of Research in Crime and Delinquency* 33:269–305.

Mazerolle, Lorraine, Emma Antrobus, Sarah Bennett, and Tom R. Tyler. 2013. Shaping Citizen Perceptions of Police Legitimacy: A Randomized Field Trial of Procedural Justice. *Criminology* 51:33–63.

McCoy, Candace. 2010. How Civil Rights Lawsuits Improve American Policing. Pp. 111–60 in *Holding Police Accountable*, edited by Candace McCoy. Washington, DC: Urban Institute Press.

McLaughlin, Eliott C. 2015. We're Not Seeing More Police Shootings, Just More News Coverage. *CNN.com*, April 21.

Mesic, Aldina, Lydia Franklin, Alev Cansever, Fiona Potter, Anika Sharma, Anita Knopov, and Michael Siegel. 2018. The Relationship between Structural Racism and Black-White Disparities in Fatal Police Shootings at the State Level. *Journal of the National Medical Association* 110:106–16.

Metzloff, Thomas B. 1991. Resolving Malpractice Disputes: Imaging the Jury's Shadow. *Law and Contemporary Problems* 54:43–129.

Morgan, Stephen L., and Joel A. Pally. 2016. Ferguson, Gray, and Davis: An Analysis of Recorded Crime Incidents and Arrests in Baltimore City, March 2010 through December 2015. Unpublished manuscript. Johns Hopkins University, Department of Sociology, Baltimore. https://socweb.soc.jhu.edu/faculty /morgan/papers/MorganPally2016.pdf.

Morin, Rich, Kim Parker, Renee Stepler, and Andrew Mercer. 2017. *Behind the Badge.* January 11. Washington, DC: Pew Research Center. http://www.pew socialtrends.org/2017/01/11/behind-the-badge.

New York City Civilian Complaint Review Board. 2017. *Annual Report 2017.* New York: New York City Civilian Complaint Review Board. https://www1 .nyc.gov/assets/ccrb/downloads/pdf/policy_pdf/annual_bi-annual/2017_annual .pdf.

Nicholson-Crotty, Sean, Jill Nicholson-Crotty, and Sergio Fernandez. 2017. Will More Black Cops Matter? Officer Race and Police-Involved Homicides of Black Citizens. *Public Administration Review* 77:206–16.

Parker, Karen F., John M. MacDonald, Wesley G. Jennings, and Geoffrey P. Alpert. 2005. Racial Threat, Urban Conditions, and Police Use of Force: Assessing the Direct and Indirect Linkages across Multiple Urban Areas. *Justice Research and Policy* 7:53–79.

Payne, Brian K., Bruce L. Berg, and Ivan Y. Sun. 2005. Policing in Small Town America: Dogs, Drunks, Disorder, and Disfunction. *Journal of Criminal Justice* 33:31–41.

Pinker, Steven. 2011. *The Better Angels of Our Nature: Why Violence Has Declined*. New York: Viking Penguin.

Powell, Zachary A., Michele Bisaccia Meitl, and John L. Worrall. 2017. Police Consent Decrees and Section 1983 Civil Rights Litigation. *Criminology and Public Policy* 16:575–605.

Priest, George L., and Benjamin Klein. 1984. The Selection of Disputes for Litigation. *Journal of Legal Studies* 13:1–55.

Rappaport, John. 2017. How Private Insurers Regulate Public Police. *Harvard Law Review* 130:1539–1614.

Ratcliffe, Michael, Charlynn Burd, Kelly Holder, and Alison Fields. 2016. *Defining Rural at the U.S. Census Bureau*. American Community Survey and Geography Brief No. 1. Washington, DC: Census Bureau. https://www.census.gov/content/dam/Census/library/publications/2016/acs/acsgeo-1.pdf.

Reaves, Brian A. 2015. *Local Police Departments, 2013: Personnel, Policies, and Practices*. Bulletin No. NCJ 248677. Washington, DC: US Department of Justice, Office of Justice Programs, Bureau of Justice Statistics.

Ross, Cody T. 2015. A Multi-level Bayesian Analysis of Racial Bias in Police Shootings at the County-Level in the United States, 2011–2014. *PLoS ONE* 10:1–34.

Ross, Darrell L. 2000. Emerging Trends in Police Failure to Train Liability. *Policing: An International Journal of Police Strategies and Management* 23:169–93.

Ross, Darrell L., and Madhava R. Bodapati. 2006. A Risk Management Analysis of the Claims, Litigation, and Losses of Michigan Law Enforcement Agencies: 1985–1999. *Policing: An International Journal of Police Strategies and Management* 29:38–57.

Ross, H. Laurence. 1970. *Settled Out of Court: The Social Process of Insurance Claims Adjustment*. Chicago: Aldine.

Rozema, Kyle, and Max Schanzenbach. 2019. Good Cop, Bad Cop: Using Civilian Allegations to Predict Police Misconduct. *American Economic Journal: Economic Policy* 11:225–68.

Schwartz, Joanna C. 2014. Police Indemnification. *New York University Law Review* 89:885–1005.

———. 2016. How Governments Pay: Lawsuits, Budgets, and Police Reform. *UCLA Law Review* 63:1144–1298.

———. Forthcoming. Civil Rights Ecosystems. *Michigan Law Review*.

Shjarback, John A., David C. Pyrooz, Scott E. Wolfe, and Scott H. Decker. 2017. De-policing and Crime in the Wake of Ferguson: Racialized Changes in the Quantity and Quality of Policing among Missouri Police Departments. *Journal of Criminal Justice* 50:42–52.

Smith, Brad W. 2004. Structural and Organizational Predictors of Homicide by Police. *Policing: An International Journal* 27:539–57.

Stephens-Davidowitz, Seth, and Hal Varian. 2015. A Hands-On Guide to Google Data. Working paper. Google, Mountain View, CA. http://people.ischool .berkeley.edu/~hal/Papers/2015/primer.pdf.

Terrill, William, Fredrik H. Leinfelt, and Dae-Hoon Kwak. 2008. Examining Police Use of Force: A Smaller Agency Perspective. *Policing: An International Journal of Police Strategies and Management* 31:57–76.

Terrill, William, and Michael D. Reisig. 2003. Neighborhood Context and Police Use of Force. *Journal of Research in Crime and Delinquency* 40:291–321.

Toner, Casey, and Jared Rutecki. 2018. 113 Suburban Cop Shootings, Zero Discipline. *WBEZ Chicago,* January 8. http://interactive.wbez.org/taking-cover/ zero-discipline.

Vidmar, Neil, Paul Lee, Kara MacKillop, Kieran McCarthy, and Gerald McGwin. 2005. Uncovering the "Invisible" Profile of Medical Malpractice Litigation: Insights from Florida. *DePaul Law Review* 54:315–56.

Walker, Samuel, and Charles M. Katz. 2008. *The Police in America: An Introduction*. 6th ed. Boston: McGraw-Hill.

Weisheit, Ralph A., L. Edward Wells, and David N. Falcone. 1995. *Crime and Policing in Rural and Small-Town America: An Overview of the Issues*. Washington, DC: National Institute of Justice.

Wines, Michael, and Sarah Cohen. 2015. Police Killings Rising Slightly, Data Suggests. *New York Times,* April 30, pp. A1, A18.

Winton, Richard. 2017. Deputy Misconduct Claims Cost Millions. *Los Angeles Times,* April 9, pp. A1, A4.

Worrall, John L., and Ricky S. Gutierrez. 1999. Potential Consequences of Community-Oriented Policing for Civil Liability: Is There a Dark Side to Employee Empowerment? *Review of Public Personnel Administration* 19:61–70.