
RESPONSE
EMPIRICALLY VALIDATING
THE POLICE LIABILITY INSURANCE CLAIM[†]

Andrea Cann Chandrasekher^{*}

INTRODUCTION

Professor John Rappaport's paper is an innovative study of police misconduct and the possible effect that liability insurance could have in reducing it. His main argument is that liability insurers, through "differentiated premiums and the threat of coverage denial,"¹ can incentivize police departments to engage in loss-prevention activities, such as revisions of on-the-job trainings² and use-of-force policies.³ Assuming that these loss-prevention measures do in fact have a causal impact on police misconduct, liability insurance (or, more specifically, a liability insurance mandate) could be yet another strategy that policy-makers draw upon to reduce police misconduct.⁴

In a time when police misconduct is at the forefront of the nation's collective mind, Rappaport's liability insurance idea is a new and fresh spin on how to address what has become a national problem. As Rappaport readily and humbly points out, however, his is a very well-researched *hypothesis* about the possible deterrence effect of liability insurance on police misconduct. I, like Rappaport, believe that what is needed to take his idea to the next level is empirical validation. In Part I, I lay out why empirical validation is a necessary next step. In Part II, I outline what it would take to empirically validate Rappaport's claim that liability insurance reduces police misconduct.

I. WHY IS EMPIRICAL VALIDATION NECESSARY?

Before jumping into the details of what it would take to empirically validate the causal claim that liability insurance decreases police misconduct in Part II, it is important to first establish and acknowledge

[†] Responding to John Rappaport, *How Private Insurers Regulate Public Police*, 130 HARV. L. REV. 1539 (2017).

^{*} Acting Professor of Law and Martin Luther King Jr. Scholar, University of California, Davis, School of Law. J.D., Stanford Law School; Ph.D. (Economics), University of California, Berkeley; M.A. (Statistics), University of California, Berkeley. The author thanks Caroline Cox, Chris Elmendorf, David Horton, and Tom Joo for extremely helpful feedback on drafts.

¹ John Rappaport, *How Private Insurers Regulate Public Police*, 130 HARV. L. REV. 1539, 1547 (2017).

² *Id.* at 1548.

³ *Id.* at 1574.

⁴ *Id.* at 1612-14.

the evidence that Rappaport has already brought to bear on the claim. Without question, Rappaport's Article presents the most comprehensive data ever assembled on the relationship between liability insurance and policing. For his project, Rappaport interviewed a group of thirty-three subjects including police chiefs, commercial insurers, consultants, and insurance risk pool members.⁵ He also collected information from various resources like trade journals and insurance documents.⁶ Based on the data collected, Rappaport argues that, *at least for his sample of interviewees*, there is evidence that insurers "can and do shape police behavior — or, at the least, . . . influence policies, practices, and personnel decisions that are themselves proven or presumed to impact behavior."⁷ Rappaport says that "it matters little that [his] sources . . . are not necessarily a nationally representative sample" because his chief task is simply to establish that insurance-based deterrence is *possible*.⁸ Having successfully established that insurance incentivizes police to adopt risk-reducing strategies *in at least some cases*, and relying on social science research that these risk-reducing strategies reduce police misconduct, Rappaport carefully concludes that police liability insurance may be a significant, yet "overlooked," "policy lever."⁹ He points to mandatory police liability insurance and laws that could be used to "encourage insurers to regulate more closely" as policies that may "improve police behavior."¹⁰

In order to responsibly take his idea of insurance-driven police reform from the research realm to the policy realm, Rappaport suggests that the causal impact of insurance on police misconduct can and should be validated.¹¹ There are at least two reasons why empirical validation is necessary. The first reason is one that Rappaport himself shines a bright light on. Because insurers are in the business of reducing municipal liability and not public harm, it is not clear *ex ante* that police liability insurance would necessarily reduce police misconduct.¹² Only a well-designed empirical study can answer this question.

The second motivation for carrying out an empirical study, one that is more pressing in my mind, is that empirical validation is the

⁵ *Id.* at 1551.

⁶ *Id.*

⁷ *Id.* at 1549; *see also id.* at 1596 ("[I]nsurers typically require police agencies to maintain adequate policies on vehicle pursuits, the use of force, and other high-risk conduct. Studies suggest that these policies do, on balance, reduce the covered harms — at least when officers are properly trained on policy content. Insurers also facilitate the use of body-worn cameras and training simulators, both of which, research suggests, can reduce the inappropriate use of force.").

⁸ *Id.* at 1549.

⁹ *Id.* at 1540.

¹⁰ *Id.* at 1549.

¹¹ *See id.* at 1595 (acknowledging that "whether insurers elevate or lower the level of police misconduct" is a "pressing" empirical question).

¹² *Id.* at 1549, 1599.

only way to determine which of the two “polar forces,”¹³ moral hazard or loss prevention, has the stronger effect on police misconduct. Rappaport explains that police liability insurance likely puts both upward pressure on police misconduct (since insured police departments have less of an incentive to reduce the risk of harm because of the moral hazard problem) and downward pressure on police misconduct (since insured police departments must adopt risk-reducing measures as a condition of their insurance policy).¹⁴ From a theoretical perspective, it is unclear which of these two effects would dominate, and Rappaport does not directly tackle this question of relative influence. Moreover, the ways in which moral hazard would manifest in the policing context are unclear since, as Rappaport describes, both individual officers and cities are often indemnified from civil liability.¹⁵ Whatever the theoretical mechanism, because of the opposing nature of these effects, it is possible that the net empirical effect of insurance on police misconduct could be negative (if the loss-prevention effect dominates), positive (if the moral hazard effect dominates), or zero (if the effects cancel each other out). Again, only a well-designed empirical study can answer this question.

II. WHAT WOULD EMPIRICAL VALIDATION TAKE?

The work of empirically validating the causal impact of liability insurance on police misconduct would not be trivial. Generally speaking, there are two types of research designs for causal inference in social science research: randomized controlled trials (RCTs) and quasi-experiments. RCTs are the “gold standard” for proving causal effects but are notoriously difficult and expensive to carry out.¹⁶ Quasi-experiments are also capable of supporting causal inference but only when the researcher exploits some type of policy intervention or random social event that credibly replicates experimental variation.¹⁷ In this Part, I briefly outline the basics of the RCT and quasi-experimental designs and how they might be used to estimate the causal effect of police liability insurance on police misconduct. As I discuss, both strategies have their strengths and weaknesses.

¹³ *Id.* at 1547.

¹⁴ *Id.* at 1543–44, 1553–54.

¹⁵ *See id.* at 1547.

¹⁶ For an explanation of the use of RCTs for causal inference, see generally WILLIAM R. SHADISH, THOMAS D. COOK & DONALD T. CAMPBELL, *EXPERIMENTAL AND QUASI-EXPERIMENTAL DESIGNS FOR GENERALIZED CAUSAL INFERENCE* (2002). *See also id.* at 13 (referring to RCTs as the “gold standard for treatment outcome research”).

¹⁷ *See* Bruce D. Meyer, *Natural and Quasi-Experiments in Economics*, 13 J. BUS. & ECON. STAT. 151, 151 (1995) (explaining quasi-experiments).

A. Randomized Controlled Trials

There is one central hypothesis that requires empirical validation: that liability insurance reduces police misconduct — even net of the increase in misconduct that is likely to occur because of the moral hazard problem — relative to tort liability without insurance.¹⁸ The most credible method of determining the effect of liability insurance on police misconduct would be to conduct a randomized controlled trial where police agencies¹⁹ would be randomly assigned to two groups — a treatment group that would be required to carry liability insurance and a control group that would be required to self-insure.²⁰ The random nature of the assignment makes it highly probable that the treatment and control groups would be, on average, similar along all observable and unobservable dimensions that could potentially have an effect on police misconduct.²¹ So, for example, the police officers in both groups would be likely to have, say, similar average levels of education, policing experience, training, and stress levels. They would also be likely to face similar average levels and types of criminal activity in their work. As a result, any later differences that develop between the two in terms of police misconduct are likely to be due solely to the treatment (that is, the adoption of liability insurance) and not to other potential confounders.²²

The causal effect of police liability insurance on police misconduct, then, would be the difference between the average police misconduct rate in the treatment group and the average police misconduct rate in the control group at the end of the experiment. If the average police misconduct rate in the treatment group is lower, then the causal effect of police liability insurance on police misconduct would be negative, meaning that insurance has a deterrence effect on misconduct, even accounting for moral hazard. If the treatment group's misconduct rate is higher, then the reverse would be true.

¹⁸ See Rappaport, *supra* note 1, at 1595 (pointing out that the relevant question is “whether liability insurance actually does reduce misconduct relative to tort liability”).

¹⁹ Technically speaking, the random assignment would happen at the municipality level and not at the police department level. However, since most municipalities have only one major police department, a random assignment of municipalities is essentially equivalent to a random assignment of police departments.

²⁰ Essentially, the treatment group would face a mandatory insurance regime. This treatment is particularly relevant since mandatory insurance is one of the primary ways in which Rappaport suggests that his insurance scheme could be implemented.

²¹ See SHADISH ET AL., *supra* note 16, at 13 (“If implemented correctly, random assignment creates two or more groups of units that are probabilistically similar to each other on the average.”). Of course, even random assignment fails sometimes. In these cases, there may not be covariate balance between the treatment and control groups and so the RCT would not yield credible estimates of the causal effect of liability insurance on police misconduct.

²² *Id.*

This basic structure could of course be tweaked in various ways. For example, if the researcher thought that the *type* of insurance (that is, risk pools versus commercial insurance) affected the size or nature of the treatment effect in a way that she was interested in, she might randomly assign police departments to two different treatment groups — one group that insured within a pool and one that insured commercially.

Of course, such a study, while ideal in its design, would be exceedingly difficult to undertake in practice for several reasons. First, there would likely be a participation problem. For example, one way of proceeding would be to take all midsize to large cities, or at least a good-sized random sample, and then randomly assign them to treatment and control groups. However, because police departments can be hesitant to engage in social science research, it is likely that not all of the chosen police departments would participate. Now, all is not lost — as long as those police departments that choose to opt out do so before the random assignment, the RCT would still have *internal validity*; that is, the RCT would still credibly tell us the causal effect of police liability insurance on police misconduct for those police departments that opted in. However, the RCT would have limited *external validity*; that is, the study's results would apply only to the subgroup of police departments that opted in and not to police departments generally. Such an outcome would be problematic if the overall goal of the RCT is to make national policy prescriptions. Second, RCTs are extremely expensive and so the financial cost of undertaking such a study would be exceedingly high — probably prohibitively so without a generous grant. Third, it might be hard to maintain treatment consistency within the treatment group. In other words, in order to preserve the integrity of the design, it would be important to make sure that all police departments were receiving the same “treatment.” Accordingly, insurers working with the treatment group would all have to mandate similar loss-prevention strategies. The same consistency principle would have to be observed in the control group as well. Since there are different ways in which police departments self-insure, with some diligently setting aside money for future claims and others adopting a more “pay-as-you-go” strategy,²³ the RCT would have to be structured so that all police departments in the control group followed the same definition of self-insurance.

As is clear, an RCT of the magnitude described above would require a high degree of planning and expertise to execute successfully. An easier, more cost-effective and direct way of cracking the question of the causal effect of liability insurance on police misconduct may

²³ Rappaport, *supra* note 1, at 1561.

be to break the effect down into its component subparts and then prove each of those subparts separately using RCT. More clearly, in Rappaport's case, the underlying claims are that liability insurance will lead to the adoption of loss-prevention strategies (causal claim #1) and that those strategies ultimately will reduce misconduct (causal claim #2). Using interviews with commercial insurers and risk pool members, Rappaport demonstrates a *correlation* between insurance and loss-prevention strategy adoption for a specific group of insurers and police departments.²⁴ In other words, for his sample, he has shown that police departments that have insurance are more likely to adopt these strategies. However, in order to prove a causal connection between insurance and loss-prevention strategy adoption, police departments would have to be randomly assigned to insurance (treatment) and self-insurance (control) to determine whether the requirement to buy insurance, in and of itself, causes police departments to adopt loss-prevention strategies. This RCT would, of course, pose all the same implementation costs and challenges as the RCT study described above. One way around this would be to assume *rationality* on the part of municipalities. Rappaport has already established that insurers provide significant monetary incentives to cities to adopt loss-prevention strategies.²⁵ Assuming that municipalities are rational profit-maximizers that, all else equal, prefer low-cost insurance to high-cost insurance,²⁶ empirical validation of the causal claim that insurance adoption increases the adoption of loss-prevention strategies can, I think, be waived.

This leaves just causal claim #2 (that certain loss-prevention strategies reduce police misconduct) to prove. I agree with Rappaport that the best strategy here is to piggyback off the work of other policing studies that investigate the causal impact of loss-prevention strategies (like changes in departmental policies, adoption of body cameras, and the use of simulators) on police misconduct.²⁷ However, this causal inference shortcut is only as credible as the research underlying it, and my assessment of these external studies is that they do not always rely on strong research designs. While the research on body cameras relies on randomized controlled trials, the research on the other two loss-prevention strategies does not. For example, the research on the effect

²⁴ See *id.* at 1573–91.

²⁵ See *id.* at 1587–91 (discussing the role that underwriting and rating play in incentivizing police departments to adopt loss prevention strategies).

²⁶ But see Daryl J. Levinson, *Empire-Building Government in Constitutional Law*, 118 HARV. L. REV. 915, 965 (2004) (challenging the assumption that governments are rational profit-maximizers).

²⁷ See Rappaport, *supra* note 1, at 1596–97; *id.* at 1596 n.332, 1597 n.333 (listing social science papers that attempt to investigate the causal impact of these loss prevention strategies on police misconduct).

of the adoption of new departmental policies on police misconduct relies on nonexperimental research designs that cannot prove causation but instead can suggest only the existence of a correlation.²⁸ Furthermore, the body-camera research, clearly the most sophisticated research in the group, sometimes points in opposite directions. Specifically, the relevant research on body cameras includes three RCT papers, two of which show a strong deterrence effect and one that shows a perverse effect.²⁹

Overall, then, one way of empirically validating the claim that liability insurance deters police misconduct via RCT may be to conduct more and better deterrence studies that use RCT. Assuming that municipalities are rational, insurers should be able to import successful loss-prevention strategies into police departments through underwriting and differentiated premiums, as Rappaport suggests. What could help to complete the causal chain, then, is a strong series of RCT studies that determine which techniques reduce police misconduct and which do not.³⁰

B. Quasi-Experiments

Quasi-experiments might be a good option for exploring the causal relationship between liability insurance and police misconduct. In a quasi-experiment, the researcher finds a naturally occurring incident, like a random social event or policy change that, by sheer coincidence, creates experiment-like variation that can be used to measure causal effects.³¹ One special case of the quasi-experimental research design is the interrupted time series design.³² With interrupted time series, the researcher compares the level of some outcome variable (like police

²⁸ See *id.* at 1596 n.332 (listing research on the adoption of new departmental policies that proves only the existence of a correlation).

²⁹ See *id.* at 1597 n.333 (listing three body-camera studies that come to conflicting conclusions).

³⁰ There are two caveats that must be mentioned with regards to this two-part RCT strategy, though. First, the two-part strategy works by decoupling the effect of insurance from the effect of loss prevention strategies (like body cameras). We first ask if insurance causes the take-up of body cameras and then if body cameras cause a reduction in police misconduct. This is not a perfect solution, however, because even if insurance induces police departments to use body cameras and body cameras reduce police misconduct, it may be that the insurance plus body camera combination leads to higher misconduct than the counterfactual of no insurance and no body cameras *because of moral hazard*. In the end, only a full-blown RCT (where insurance and loss prevention are not decoupled) could resolve this question. The second caveat is that the two-part strategy assumes that the effect of loss prevention strategies on police misconduct is the same regardless of whether the police department adopts the loss prevention strategies on their own (as is the case in social science research) or is required by an insurer to do so. If police departments try to skirt around implementing insurer-imposed loss prevention strategies, or implement them differently when required to do so by an insurer, then this assumption will not hold.

³¹ *E.g.*, Meyer, *supra* note 17, at 151.

³² See Gene V. Glass, *Interrupted Time Series Quasi-Experiments*, in COMPLEMENTARY METHODS FOR RESEARCH IN EDUCATION 589 (Richard M. Jaeger ed., 2d ed. 1997).

misconduct) right before and right after the imposition of some policy that affects the availability (or level) of the treatment (in this case, police liability insurance). Assuming that other (possibly confounding) variables don't change sharply at the policy threshold, any change in police misconduct at the policy threshold can credibly be attributed to the availability of police liability insurance.

Of course, the difficulty comes in identifying a policy change or random social event from history that discontinuously impacts the availability (or cost) of police liability insurance but not any of the other determinants of police misconduct. One tempting option might be to use the municipal liability insurance crisis of the 1980s as a quasi-experiment. According to Rappaport, as a result of the crisis, "[p]olice liability insurance, for practical purposes, had vanished."³³ If so, then one possible path forward might be to compare police misconduct levels right before and right after the start of the insurance crisis for a sample of police departments that lost their insurance because of the crisis. (Let's call them the treatment group.) If insurance adoption does indeed have a deterrence effect on police misconduct, we might expect to see police misconduct increase in the period right after the insurance crisis began when the availability of insurance plummeted. To make the study richer, the rise (or fall) in police misconduct for those departments that lost their insurance because of the crisis could be compared to the change in police misconduct experienced by those departments that never had insurance. Those police departments that never had insurance would serve as the control group, and changes in their misconduct rate right after the start of the crisis (if any) could be treated as a benchmark for how much police misconduct levels would have changed at the threshold for the treatment group absent the insurance crisis. Any rise (or fall) in the treatment group's police misconduct rate, above and beyond this benchmark change, would be the causal effect of insurance (or more precisely, noninsurance) on police misconduct.

To be sure, such a quasi-experiment would face several challenges. In particular, collecting historical police misconduct data for multiple police departments might present a significant hurdle. However, the above example hopefully provides a helpful illustration of how quasi-experiments could work to uncover the causal effect of insurance on police misconduct.

³³ Rappaport, *supra* note 1, at 1556.

CONCLUSION

Rappaport argues that police liability insurance may have a deterrence effect on police misconduct. In order to take his idea to the next level, empirical validation through an RCT or quasi-experimental study is necessary. Even though RCT is the gold standard for causal inference, in most policy contexts, RCT-based proof is not feasible. In this context of police liability insurance, however, proving the causal link from insurance to police misconduct via RCT may be advanced by (1) assuming that municipalities are rational profit-maximizers, and (2) carrying out more and better RCTs where the deterrence effect of various loss-prevention strategies on police misconduct is studied.