ESTIMATING EMPIRICAL BLACKSTONE RATIOS IN TWO SETTINGS: MURDER CASES AND HIRING

*Shawn D. Bushway**

ABSTRACT

There is a growing awareness in the legal literature of the need to estimate the prevalence of errors that exist within the criminal justice system. A majority (but not all) of the time, the focus is on the false positive, or wrongful conviction rates. Yet, a complete picture of the decision process requires estimates of both false positives and false negatives. In this paper, I generate an estimate of the false negative rate for a representative sample of murders in Chicago. I also estimate the cost ratio of false negatives to false positives that would be needed to justify using records of incarceration to identify people at risk in the Chicago metropolitan area. Both estimates should shed meaningful light on the growing debate about what rules should be set to achieve more socially optimal decisions in both the criminal justice system and the labor market. Future work should focus on replicating and extending these preliminary estimates.

I. INTRODUCTION

Some decisions involve a choice between two options. In the case of the trial, the jury is trying to decide if a person is guilty or innocent, starting from the null hypothesis that the person is innocent. In the case of employment, an employer is concerned about hiring a risky employee who will harm fellow employees or clients. In this simplest kind of decision framework, there are two kinds of errors. False positives are innocent or not risky people who are identified as guilty/risky.[1] False negatives are guilty/risky

* Associate Professor, School of Criminal Justice, University at Albany.

[1] Henry L. Chambers, Jr., *Getting it Right: Uncertainty and Error in the New Disparate Treatment Paradigm*, 60 ALB. L. REV. 1, 5–6 (1996).

people who were not identified as guilty/risky.[2]

The Blackstone ratio on which this special issue is based makes it clear that policymakers can specify the nature of the tradeoff between these error rates. Specifically, in the context of conviction, Blackstone hypothesizes that it would be ideal to have a justice system that generates ten false negatives for every false positive.[3] There has been considerable subsequent debate about the relative desirability of a conviction of an innocent man versus allowing a guilty man to go free, or whether such a tradeoff is even morally acceptable.[4] A detailed review of the literature by Alexander Volokh found that the most commonly accepted standard in the U.S. is that ten guilty men should go free before one innocent man should be found guilty.[5] Volokh also found states that advocate for a one-to-one standard, as well as one state, Oklahoma, with a standard of one false positive for every one hundred false negatives.[6]

I am not aware of a systematic effort to describe the false negative rate in the U.S. criminal justice system using archival data, but Shawn Bushway and Brian Forst provide a back-of-the-envelope estimate based on aggregate data of 1500 to 3000 false negatives to each false positive in the U.S. criminal justice system.[7] In this paper, I use existing data to determine a rudimentary estimate of the number of false negatives in murder investigations in Chicago in 1979. I find an empirical Blackstone ratio of sixty-one false negatives for every false positive, with a lower bound of thirty-five-to-one.

Employers and others who use criminal history records are more concerned about false negatives than false positives; that is, they want to avoid hiring risky people who will harm someone while working. As a result, employers are plausibly willing to tolerate a certain number of false positives for every false negative. In the extreme case, employers would be worried only about false negatives. Most of the costs of false positives, like increased crime due to frustration or lack of legitimate income, are born by agents other than the employer. However, society can make employers feel

---

[2] *Id.*

[3] 4 WILLIAM BLACKSTONE, COMMENTARIES *358.

[4] *See, e.g.*, D. Michael Risinger, *Innocents Convicted: An Empirically Justified Factual Wrongful Conviction Rate*, 97 J. CRIM. L. & CRIMINOLOGY 761, 788–89 (2007).

[5] Alexander Volokh, *N Guilty Men*, 146 U. PA. L. REV. 173, 174 (1997).

[6] *See id.* at 204.

[7] Shawn Bushway & Brian Forst, *Studying Discretion in Sentencing* 7 (working paper, 2010) (on file with author).

some of those costs (for example, through the threat of Title VII litigation). And employers can also have direct costs if they cannot find enough qualified employees. Indeed, there is evidence that employers do willingly hire ex-offenders.[8]

I am aware of no literature that tries to quantify the acceptable tradeoff between false positives and false negatives by employers. In the second half of the paper, I conduct an exercise to back out the implied "acceptable" ratio of false negatives to false positives that is implied by an employer which uses a prison record as a predictor for homicide. I find that employers who decide not to hire people with prison records would have to argue, at minimum, that 930 false positives have the same cost as one false negative.

II. Estimating the Implied "Blackstone Ratio" in Murder Cases

There are two good outcomes for any criminal case—the guilty are convicted, and the innocent are not convicted. But there are two errors that can and will be committed in any system of justice—false negatives and false positives. Part of the value of the Blackstone ratio, in my view, is the explicit recognition that wrongful convictions are not the only errors that are committed in any system. Moreover, in my view, the only way to assure no false positives would be to fail to convict anyone (and vice versa).

Blackstone ratio also implies that the ratio in any given system of justice is open to manipulation. Conceivably, policymakers could set certain rules and standards that would change the "justice" outcome to achieve the desired "Blackstone Ratio," holding constant the kinds of technological innovations that could conceivably reduce both forms of error simultaneously.

However, this idea is hypothetical at this point because there are few good estimates of the amount of errors made in the system. Through the advent of DNA testing, there has been a recent burst of identifying (and correcting) wrongful convictions. Although some are skeptical that good estimates are even possible, others believe that good strategies exist for generating reliable estimates of wrongful convictions in the system. A recent detailed effort found that 3.3% of those convicted in capital death cases were factually

[8] Michael A. Stoll & Shawn D. Bushway, *The Effect of Criminal Background Checks on Hiring Ex-Offenders* 3 (Nat'l Poverty Ctr., Working Paper No. 07-08, 2007), *available at* http://npc.umich.edu/publications/u/working_paper07-08.pdf.

innocent,[9] and there is some consensus from recent reviews that one to five percent of all felony convictions are of factually innocent people.[10]

Within this literature, there is a consistent effort to differentiate the "factually innocent" from those who are acquitted and might actually be guilty. However, I found no explicit attempt to estimate the number or rate of false negatives.[11] But these two errors are linked systematically and unavoidably. Identifying one without the other captures only half of the error story, and could lead to changes in policy that dramatically increase the number of guilty people who are not convicted, for instance.

In the following, I make use of a little known study of police investigative techniques, based on crimes committed in 1979, to conduct a simple exploration of false negative rates for murder in Chicago.[12] The authors of the study randomly sampled seventy-two murders in Chicago that were reported to the police in 1979.[13] There were 856 murder victims reported to the police in 1979,[14] so this represents 8.4% of the population. The authors collected information from the police about the investigation, as well as court data on the final outcome for each suspect.[15]

The goal of this section is to fill in the decision grid for the individuals who committed these murders. The decision grid is provided in Table 1. The columns represent the decision of the criminal justice system, and the rows involve the factual innocence

[9] Risinger, *supra* note 4, at 780.

[10] For an excellent review of research on wrongful convictions, see Jon B. Gould & Richard A. Leo, *One Hundred Years Later: Wrongful Convictions After a Century of Research*, 100 J. CRIM. L. & CRIMINOLOGY 825 (2010). *See also* Samuel R. Gross & Barbara O'Brien, *Frequency and Predictors of False Convictions: Why We Know So Little and New Data on Capital Cases*, 5 J. EMPIRICAL LEGAL STUD. 927 (2008); Marvin Zalman et al., *Officials' Estimates of the Incidence of "Actual Innocence" Convictions*, 25 JUST. Q. 72 (2008).

[11] I did find several papers, and there are no doubt others, that look at false negative rates for eyewitness identification. The logic that is used in these papers is applicable for the following archival analysis of false positives. *See generally* Gary L. Wells et al., *Eyewitness Evidence: Improving Its Probative Value*, 7 PSYCHOL. SCI. PUB. INT. 45 (2006); Gary L. Wells & Elizabeth A. Olson, *Eyewitness Identification: Information Gain from Incriminating and Exonerating Behaviors*, 8 J. EXPERIMENTAL PSYCHOL.: APPLIED 155 (2002); Steven Penrod, *Eyewitness Identification Evidence: How Well Are Witnesses and Police Performing?*, 54 CRIM. J. MAG. 36 (2003), *available at* web.jjay.cuny.edu/~spenrod/papers/PenrodCJ03.doc.

[12] Joseph L. Peterson et al., *Description and Citation—Study No. 8186: Forensic Evidence and the Police, 1976–1980*, ICPSR (2002), http://www.icpsr.umich.edu/icpsrweb/ICPSR/studies/8186/detail.

[13] *Id.*

[14] Carolyn Rebecca Block & Richard L. Block, *Description and Citation—Study No. 6399: Homicides in Chicago, 1965–1995*, ICPSR (2005), http://www.icpsr.umich.edu/icpsrweb/ICPSR/studies/6399/detail.

[15] Note that not every case had a suspect.

or guilt of the people.

TABLE 1: THEORETICAL DECISION TABLE

	Not Convicted	Convicted	TOTAL
People Who Committed the Murders	A	B	C
People Who Did Not Commit the Murders	D	E	F
TOTAL	G	H	I

Although a more nuanced decision table is possible, this table represents the simplest possible description of the problem. The most important thing to note about this table is that it involves all people who murdered the victims in these cases, not just those suspects who go to trial or are arrested.

A key fact of the criminal justice system is the high degree of selection from arrest to conviction. This selection is not random, but driven in large part by the probability of conviction.[16] The theory of "bargaining in the shadow of [the] trial" is just one of the most obvious examples of this logic, where legal theorists argue that pleas are driven by what happens at the trial.[17] The implication here is that upstream actors in the criminal justice system know the standards of conviction and will not bring forward cases, or even make arrests, that will not meet this standard. However, because this is true, the universe of cases that have an arrest, or go to trial, are a highly selected subsample of cases. Starting at arrest, rules that exist to prevent false positives will have the potential to create false negatives by eliminating guilty people from the criminal justice system against whom there is simply not enough evidence under the current standard of proof. Ignoring these people necessarily undercounts the number of false negatives. Thus, a reasonable alternative is to start with the number of people who commit the crime, and focus on the conviction decision.

[16] *See* Richard A. Berk, *An Introduction to Sample Selection Bias in Sociological Data*, 48 AM. SOC. REV. 386, 390 (1983); Steven Klepper et al., *Discrimination in the Criminal Justice System: A Critical Appraisal of the Literature*, *in* RESEARCH ON SENTENCING: THE SEARCH FOR REFORM 55 (Alfred Blumstein et al., eds., vol. II 1983).

[17] Stephanos Bibas, *Plea Bargaining Outside the Shadow of Trial*, 117 HARV. L. REV. 2463, 2486–87 (2004).

A priori, there are some elements which will not be easily knowable. For example, *F*, the total number of people who did not commit the crime, is a difficult conceptual number. Clearly, any number of people in Chicago did not commit the crime, so this number could be arbitrarily large. Conceptually, it might be cleaner to think of this as a list of "persons of interest" who could have plausibly committed the crime, but did not. It might be possible to make a list of "persons of interest" for the crime, but for the purpose of this discussion, I will not try to estimate *F*. For the same reason, I will also not try to estimate *D*, *G*, or *I*.

The dataset from Chicago has a unique feature in that it recorded the number of suspects that police believe, on the basis of the evidence, were linked to the murder. For the purposes of this exercise, I assume that this number is the total number of guilty people. Of course, it is possible that the witnesses or the evidence are mistaken, and in a later section I loosen this assumption.

There are seventy-two murders in the data. Nine of them had no information about the number of suspects and are dropped from the data.[18] Of the remaining sixty-three, forty-six had one suspect, ten had two suspects, five had three suspects, and two had four suspects, for a total of eighty-nine "guilty" suspects (C=89).[19]

The data also provide a detailed record of the final charge and the disposition of the case for each defendant. For example, we know if the person was convicted of murder by a jury, a court, or through a plea bargain. According to this data, a total of twenty-nine people were convicted of murder, attempted murder, voluntary manslaughter, or involuntary manslaughter for these sixty-three murders.[20] Two were convicted of other charges,[21] but I have decided to consider them not convicted in this analysis since they were not convicted of killing someone, the crime in question. Fourteen of the defendants pled guilty, and only three were convicted by a jury. Just as interesting, twenty-six of the sixty-three cases had no one even charged. Based on this data, I can conclude that H=29 for Table 1.

What remains unknown is the number of people who were falsely convicted, *E*. For the purposes of this discussion, I will assume one person was falsely convicted for a 3.4% rate of false conviction, which is consistent with the estimate by Risinger in capital death

[18] Block & Block, *supra* note 14.
[19] *Id.*
[20] *Id.*
[21] *Id.*

murder cases.[22] This assumption means that *B*=28 and *E*=1.

In Table 2, I fill in the estimates that I was able to generate from the police data. Because I was willing to make the strong assumption that the police estimate of the number of suspects based on witness testimony and evidence is a good estimate of the number of people who were actually guilty of murder, I now have an estimate of *A*, the number of false negatives. *A*, the number of guilty people who are not convicted of at least involuntary manslaughter, is just the total number of convicted guilty persons subtracted from the total number of guilty persons (*C* – *B*, or 89 – 28, therefore *A*=61). In other words, a reasonable estimate of the number of false negatives in this data is sixty-one. In rate terms, that means there is a false negative rate of 68.5%.[23] Given my estimate of one falsely convicted person, I estimate a Blackstone ratio in this data of sixty-one-to-one (sixty-one false negatives for every one false positive). This is clearly greater than the ten-to-one ratio, but not nearly as large as the ratio of 2500-to-1 estimated by Bushway and Forst for all felony crimes.[24]

An interesting question is what the false positive rate would have to be to produce a Blackstone ratio of ten-to-one. Simple algebra can be used to show that there would need to be six false positives (wrongful convictions) in order to create a ten-to-one ratio of false positives to false negatives. Six false positives would produce a false positive rate of 21%, which is quite high relative to current estimates.

[22] *See* Risinger, *supra* note 4, at 780.

[23] This number is generated by dividing sixty-one by eighty-nine. In the language of the wrongful conviction or false positive literature, this false negative rate can be thought of as the wrongful non-conviction rate for those who are factually guilty.

[24] Bushway & Forst, *supra* note 7, at 19. This is not necessarily a contradiction. Murders have higher clearance rates than most crimes, and therefore should have fewer false negatives than crimes such as burglary.

TABLE 2: EMPIRICAL DECISION TABLE

	Not Convicted	Convicted	TOTAL
People Who Committed the Crime	61	28	89
People Who Did Not Commit the Crime	D	1	F
TOTAL	G	29	I

My estimate of sixty-one-to-one represents a first estimate for the ratio of false negatives to false positives in archival data. As such, I expect it to be, at best, a starting point for future discussions. Like all archival estimates, it has a number of obvious shortcomings.

First, it rests crucially on the assumption that the police estimate of the number of perpetrators is a good estimate of the number of guilty parties. If this assumption is incorrect, the exercise is flawed. However, the magnitude of the flaw is necessarily bounded by the nature of the problem. Since there are sixty-three murders, the smallest number of murderers possible is sixty-three.[25] Recalculating Table 2 with C=63 instead of eighty-nine provides an estimate of thirty-five false negatives, which provides a false negative rate of 55.5%, and an empirical Blackstone ratio of thirty-five-to-one. As a result, I feel confident that the lower bound or floor of the Blackstone ratio in this data is thirty-five-to-one.

Second, this is one estimate based on sixty-three murders in one city in 1979. Murders receive a lot of attention because of the severity of the punishment and the existence of DNA evidence to estimate false positives. But other crimes are more common and are also important to consider. This dataset could facilitate estimates for other cities/crimes, although the sampling makes estimates for other crimes more problematic. Murder is the only crime for which all murders were included in the sampling frame. Other crimes were included only if there was a forensic investigation. There may be other datasets that have the necessary data to generate similar estimates. Ideal datasets need to have the universe of crimes reported to the police, the number of perpetrators, and the final court dispositions.

[25] I treat each murder case as a separate case, so serial murderers will count more than once.

III. ESTIMATING THE IMPLIED "BLACKSTONE RATIO" IN HIRING

A. Hypothetical Example

In the preceding section, I made a direct attempt to estimate the number of false negatives in the criminal justice system using archival data from a set of murder cases in Chicago. In this section, I attempt a similar exercise for employers and others who use criminal history records as a measure of risk.

Employers are particularly fearful of the false negative (someone who is going to do something harmful, and is therefore "risky" but was not identified as "risky" in screening) because of negligent hiring litigation.[26] Employers can be held liable for the criminal actions of employees that were "reasonably foreseeable."[27] The existence of a criminal history record, especially for offenses that are related to the offense that caused harm on the job, can be viewed as reasonably foreseeable by the courts.[28] Excluding all individuals with criminal history records to avoid the rare catastrophic criminal act by an employee is guaranteed to generate more false positives than false negatives by definition. There are simply very few people who commit these types of crimes in the first place—so even if the employer identifies no one as "high risk," there will be very few false negatives. As a result of this basic fact, the relevant "Blackstone" ratio for employers goes in the opposite direction than the Blackstone ratio for the criminal justice system in that employers are willing to tolerate more false positives relative to false negatives.

Unlike the situation in Part II, where I have actual data from the criminal justice system, I do not have data from employers on who they do not hire, or who goes on to commit serious crimes while employed. As a result, I cannot generate an estimate for the number of false negatives or false positives. However, I can use a typical hiring screen for ex-offenders, and estimates for the criminal activity among individuals with that characteristic, to estimate what the ratio would have to be to justify such a screen. While not definitive, this estimate should provide readers with a "benchmark"

[26] Cindy M. Haerle, *Employer Liability for the Criminal Acts of Employees Under Negligent Hiring Theory:* Ponticas v. K.M.S. Investments, 68 MINN. L. REV. 1303, 1303 (1984).

[27] Irina Kashcheyeva, Comment, *Reaching a Compromise: How to Save Michigan Ex-Offenders from Unemployment and Michigan Employers from Negligent Hiring Liability*, 2007 MICH. ST. L. REV. 1051, 1060 (2007).

[28] *Id.*

for the tradeoffs involved for employment relative to the criminal justice system.

I start by establishing basic concepts before generating a reasonable estimate for the Chicago area.[29] Suppose that fifty percent of the potential employees have some characteristic that makes them risky. Furthermore, suppose there was no information about potential risk so, when faced with one hundred potential employees, the employer simply flipped a coin one hundred times and labeled those who received heads as risky, and those who received tails as "not risky". In this case, half the people who were labeled as risky would indeed be risky (twenty-five true positives), and half the people who were labeled as risky would be "not risky" (twenty-five false positives). In contrast, of those not labeled risky, half would be "not risky" (twenty-five true negatives), and half would be risky (twenty-five false negatives). This procedure produces an error rate of fifty percent. If false positives and false negatives are equally undesirable, this procedure would be equally accurate as saying no one was risky (fifty false negatives) or saying everyone was risky (fifty false positives).

The goal of a good risk predictor is to reduce the error rate. Suppose the employer knew that men are four times more likely to be risky than women, and there are fifty men and fifty women in the pool of applicants. This means that men have an eighty percent chance of being risky, whereas women have a twenty percent chance. The overall pool continues to have a fifty percent chance of being risky. Suppose the employer decided to use gender as a predictor of risk. Men would be called "risky" and women "not risky." Ten of the women are risky, so there will be ten false negatives, and ten of the men are not risky, so there will be ten false positives. This particular risk predictor would have reduced the error rate from fifty percent to twenty percent.

The problem, however, is that this ability to reduce errors depends critically on the base rate. A central truth from the prediction literature is that the difficulty of predicting events increases as the base rate moves away from 0.5.[30] Thus, the less frequent the underlying event, the greater the chance of an inaccurate prediction. In this context, an event that occurs with ten

---

[29] I chose Chicago simply because the data from the criminal justice exercise came from Chicago.

[30] For a definitive statement of this issue, see Stephen D. Gottfredson & Don M. Gottfredson, *Accuracy of Prediction Models*, *in* CRIMINAL CAREERS AND "CAREER CRIMINALS" 212 (Alfred Blumstein et al. eds., vol. II 1986).

percent probability is considered to be rare.[31]

Suppose, for example, a case where only ten percent of the population is risky, and men are still four times more likely to be risky than women. This would mean that men have a sixteen percent chance of being risky and women have a four percent chance. Suppose that the employer again decides to label men as "risky" and women as "not risky." Now, the error rate increases dramatically. There are two false negatives (since two of the women are risky), but there are forty-two false positives since forty-two of the men are "not risky." This makes a total error rate of forty-four percent.

Furthermore, suppose instead that the employer chose simply to not use a risk tool. This is equivalent to predicting that no one is risky. This decision rule would have an error rate of ten percent since there are only ten risky people, and the employer would have ten false negatives. From an accuracy perspective, it would have been reasonable for the employer not to use the available risk prediction tool based on gender because they were more accurate when they used no tool than when they used the available evidence.

However, acting as if both errors are equally undesirable (as is standard in the criminological literature) is not reasonable.[32] For the employer, hiring false negatives (risky people who are identified as "not risky") is substantially more undesirable than not employing false positives (a not risky person who was identified as "risky"). Or, to put it another way, employing someone who is at risk to harm a fellow employee or client is more costly than refusing to hire someone who would have been a relatively risk-free employee.

This relative desirability of not employing a false positive (someone who is not risky but labeled by the tool as "risky"), when compared with employing a false negative (someone who is risky but is labeled as "not risky"), does not imply that the employer does not care about false positives. Not allowing people who are risk-free to earn a living or supplement their income deprives these individuals of the ability to take care of their families, and might force people to commit crime to support themselves or their families. Moreover, the employer may be subject to, or suffer exposure to, anti-discrimination laws. Finally, the employer and its clients/customers suffer adverse consequences if the employer

[31] *See* Vernon L. Quinsey, *The Baserate Problem and the Prediction of Dangerousness: A Reappraisal*, 8 J. PSYCHIATRY & L. 329, 331 (1980).

[32] Richard Berk, *Asymmetric Loss Functions for Forecasting in Criminal Justice Settings*, 27 J. QUANTITATIVE CRIMINOLOGY 107, 107 (2010).

cannot identify eligible employees.

It is beyond my scope here to state the absolute consequences of a false negative or a false positive. However, it is possible to say something about the relative consequences implied by any given risk tool. In the above example, a person could justify the use of the risk tool if they were willing to argue that two false negatives and forty-two false positives impose fewer adverse consequences on the company and/or society than hiring five false negatives. In essence, this is an argument that the negative consequences of one false negative are at least as large as the consequences of fourteen false positives (i.e., that the three fewer false negatives offset the cost of not hiring the forty-two false positives).

But what if the base rate is only one percent? Assume again that men have a risk level that is four times that of women, which means that women have a 0.4% chance of being risky, and men have a 1.6% chance of being risky. A policy that identifies men as risky would create 49.2 false positives and 0.2 false negatives, for a total error rate of 49.4%. Alternatively, suppose the employer simply decided not to predict risk. They would have a total of one false negative, for a total error rate of one percent. An argument that using this risk tool as an accurate tool would require that one false negative imposes substantially more adverse consequences than 61.5 false positives. Likewise, a base rate of 0.1% would require that false negatives are substantially more undesirable and risky than 624 false positives. It should be clear from this example that the accurate prediction of any characteristic or event depends critically on the base rate of the event.

B. Empirical Estimate for Chicago

In this section, I will try to calculate the relative risk ratio implied by the use of a criminal history record by an employer concerned about hiring someone who might murder someone in 2008 in the Chicago-Naperville-Joliet, Illinois Metropolitan Division ("Chicago MD"). A number of simplifying assumptions will be made to make this exercise tractable, but the estimate represents a first order approximation of the exercise conducted by employers interested in avoiding hiring someone who might murder someone while working for the company. Murder is not the only crime that an employer might be worried about, but it represents a serious harm that could be pursued in court, and has the best data available for this exercise. It is well-known that factors such as race, criminal history record, age, and gender are among the best predictors of future criminal activity.[33] In what follows, I will concentrate on criminal history record.[34]

I need two statistics for this analysis—the percentage of adults in the Chicago MD with criminal history records, and the percentage of fatal attacks that are committed by individuals with criminal history records. Given that employers will only know about convictions, and not incarceration spells, I would like to use convictions as the measure of criminal history, and find a statistically valid estimate of the percentage of adults with felony convictions. However, this statistic is not available from a federal statistical agency. Instead, the U.S. Bureau of Justice Statistics has published statistics on the percentage of adults who have been incarcerated. Therefore, the analysis that follows will use information on incarceration rather than conviction.

The Chicago MD represents the set of Illinois counties around the central city of Chicago that are substantially connected socially and economically. The Chicago MD accounts for "59.5 percent of the state's population and account[s] for 65.0 percent of the Illinois crime index offenses."[35] The Chicago MD can reasonably be viewed as the labor market pool for Chicago, and represents a reasonable area in which an employer in central city Chicago might wish to make inference. The example will proceed in three steps. First, I generate an estimate for the base rate of murderers in 2008 in the Chicago MD. Second, I generate an estimate for the base rate of murderers in 2008 among ex-prisoners in the Chicago MD. Finally, I calculate the ratio of false negatives to false positives that would

---

[33] Paul Gendreau et al., *A Meta-Analysis of the Predictors of Adult Offender Recidivism: What Works!*, 34 CRIMINOLOGY 575, 583 (1996).

[34] "Criminal history record" is a generic term used to refer to past involvement with the criminal justice system. It can refer to administrative records from the police, corrections, and/or the courts system. Employers typically have access to a hodgepodge of records available at the state or local level. From this hodgepodge of records, the best possible source of information is typically court records. *See generally* DEREK HINTON, CRIMINAL RECORDS BOOK: THE COMPLETE GUIDE TO THE LEGAL USE OF CRIMINAL RECORDS (Michael L. Sankey & Peter J. Weber eds., 2002). Court records have a number of limitations, including the fact that they are only available at the county level. They also often do not include the disposition of the case (including sentence received), and they contain no information on the time spent in prison. As a result, the standard criminal history record check conducted by a background check company will not contain information on prison terms. *See generally* LESTER S. ROSEN, THE SAFE HIRING MANUAL: THE COMPLETE GUIDE TO KEEPING CRIMINALS, IMPOSTERS AND TERRORISTS OUT OF THE WORKPLACE (Michael L. Sankey & Robert Peterson eds., 2004).

[35] ILL. UNIF. CRIME REPORTING PROGRAM, ILL. STATE POLICE, CRIME IN ILLINOIS 198 (2008), *available at* http://www.isp.state.il.us/docs/cii/cii08/cii08_Section_III_Pg195_to_200.pdf.

justify the use of this particular screen by employers.

C. Base Rate of Murderers in the Chicago MD Among the Labor Force

In the Chicago MD, there were 616 victims of murder and negligent manslaughter in 2008, according to the Illinois State Police.[36] Using numbers from the U.S. Bureau of Justice Statistics on the number of homicide events that have multiple victims, I calculate that there were approximately 584 murderers responsible for the 616 victims in this MD.[37] Although the Illinois crime report does not provide this number, a reasonable estimate from this state is that 85% of all arrests for murder were of individuals between the ages of fifteen and sixty-four, so a reasonable statistical estimate is that 584 individuals (85% of 616) who fatally attacked someone were between the ages of fifteen and sixty-four.[38] I will use the age range of fifteen to sixty-four to represent the bulk of the people in the workforce.

According to the U.S. Census Bureau, there were 5,173,926 individuals between the ages of fifteen and sixty-four in the Chicago MD in 2008.[39] Therefore, I calculate that the base rate of adults who fatally attacked someone was 0.000113.[40] In other words, one hundredth of one percent of the population between the age of fifteen and sixty-four in the MD fatally attacked at least one other person. By any standard, this is a very low base rate.

D. Base Rate of Murderers in the Chicago MD Among All Ex-

[36] Murder is defined here as a Part I index offense, as defined by the FBI. *See generally* U.S. DEP'T OF JUSTICE, FED. BUREAU OF INVESTIGATION, UNIFORM CRIME REPORTING HANDBOOK (2004), *available at* http://www2.fbi.gov/ucr/handbook/ucrhandbook04.pdf. For the purposes of this study, murder combines two categories: murder, defined as the "willful killing of a human being," and manslaughter, defined as the "killing of a human being by the negligence of another."

[37] JAMES ALAN FOX & MARIANNE W. ZAWITZ, BUREAU OF JUSTICE STATISTICS, HOMICIDE TRENDS IN THE UNITED STATES (2007), *available at* http://bjs.ojp.usdoj.gov/content/pub/pdf/htius.pdf. In 2005, the majority of all homicide events had one victim, but 4% involved 2 victims, 0.6% involved 3 victims, 0.1% involved 4 victims, and 0.05% involved 5 or more victims. I use these numbers to estimate that there was 0.946 of a murderer for every one murder victim.

[38] 58 = 80% of 72.

[39] *See generally* U.S. CENSUS BUREAU, STATE AND METROPOLITAN AREA DATA BOOK: 2010 119 (7th ed. 2010), *available at* http://www.census.gov/compendia/databooks/2010/tables/sma_B-03.pdf.

[40] I assumed that all of the people who fatally attacked someone in the Chicago MD were from this MD.

prisoners Aged Fifteen to Sixty-Four

Now, according to a report from the U.S. Bureau of Justice Statistics, 2.7% of the U.S. adult population has been in prison at least once.[41] Also, according to an analysis of State Courts Processing Statistics, 26% of those individuals twenty years of age or older indicted for murder in the largest seventy-five counties in the United States in May 2002 had prior prison sentences.[42]

Assuming that the national statistics hold in the Chicago MD, 2.7%, or 139,696, of the population in the Chicago MD have been in prison. And, if the national statistics hold, 26% of the 584 murderers (152 people) will have past prison records, which is a base rate of 0.00109 (one tenth of one percent) for ex-prisoners (15/14,027). The base rate of all adults was 0.000112 (one hundredth of one percent), about 1/10th the base rate for ex-prisoners. Or, to put it another way, ex-prisoners are ten times more at risk to commit a murder than the general population. A factor that can identify people with a ten times elevated risk is a strong predictor, which is the reason behind the use of criminal history records to predict future behavior. The question is how does this elevated level of risk translate into error rates in the effort to predict who in the Chicago MD will be a murderer in the next year.

E. Relative Risk of False Negatives and False Positives

If an employer labels all ex-prisoners as "risky," then they will have an error rate of 2.705%. All of the ex-prisoners (2.697%) who do not commit murder will be false positives, and the non ex-prisoners who commit murder (0.00835%) will be false negatives. In contrast, the error rate of *not* labeling anyone as "risky" will be 0.0113%, the rate of murderers in the target population. All of these will be false negatives. In general terms, using ex-prisoner status as a risk factor will increase the inaccuracy of risk prediction by *23,842%*. In order to justify this inaccuracy on the grounds that false negatives are more undesirable than false positives, the employer using the tool would need to assert that false negatives are more than 930 times more costly than false positives. Although

---

[41] THOMAS P. BONCZAR, PREVALENCE OF IMPRISONMENT IN THE U.S. POPULATION, 1974–2001 1 (2003), *available at* http://bjs.ojp.usdoj.gov/content/pub/pdf/piusp01.pdf.

[42] BUREAU OF JUSTICE STATISTICS, U.S. DEP'T OF JUSTICE, *Description and Citation—Study No. 2038: State Court Processing Statistics, 1990–2006: Felony Defendants in Large Urban Counties*, ICPSR (2010), http://icpsr.umich.edu/icpsrweb/ICPSR/studies/2038/detail.

I am not claiming that employers have gone through this exercise, I feel confident in stating that using ex-prisoner status as a background screen in a context where the employer is worried about murder implies that they put a cost on false negatives that is almost 1000 times larger than a false positive.

Although this seems like a large ratio, it is possible to make an argument that this ratio is at least reasonable. For example, suppose that the average negligent hiring lawsuit costs employers $1,000,000.[43] In that case, hiring a murderer (someone who will murder someone in the coming year) will cost an employer $1,000,000. This implies that an employer values a false positive (not hiring a perfectly good employee with a criminal history record) at $1075 (1,000,000/930). If anything, this seems high in terms of real costs to the employer for not hiring someone, and implies that the employer is considering societal costs.

In reality, only a small fraction of murders are committed on the job, and the employer is not liable for murders by employees while not working, although the negative publicity might hurt certain types of employers. A plausible alternative assumption would be that a false negative costs the employer $100,000. That would imply that the false positive costs $1.08, which seems low—it probably costs more than that just to have someone fill out an application.

It is beyond the scope of this article to provide a definitive answer to the question of the actual cost of a false negative. Attaching numbers to the estimated ratios should provide the reader with a sense of the implications of a "Blackstone" ratio of almost 1000 false positives to false negatives by employers. Furthermore, the larger goal of this section was to provide a reasonable estimate of the implied ratio of employers who use criminal history records, which are undeniably a "good" predictor of future crime, to predict rare behavior such as murder.

Pager and Quillian argue that employers are not screening with criminal history records to avoid negligent hiring lawsuits, but rather to provide information about simple, straightforward employee characteristics like reliability and trustworthiness.[44] If this is true, then the "risky" behavior will be less rare, and the implied ratio will be smaller. This makes sense because the cost of

[43] I have not been able to find definitive data that reports the average award.

[44] *See generally* Devah Pager & Lincoln Quillian, *Walking the Talk? What Employers Say Versus What They Do*, 70 AM. SOC. REV. 355 (2005).

a false negative will decrease if the targeted behavior is poor job performance rather than murder. The analysis in this section is not dependent on the targeted behavior. Although the ratio will change depending on what the employer is trying to predict/avoid, the logic of the exercise remains consistent.

IV. CONCLUSION

There is a growing awareness in the legal literature of the need to estimate the prevalence of errors that exist within the criminal justice system. Much (but not all) of the time, the focus is on the false positive, or wrongful conviction rates. Yet, a complete picture of the decision process requires estimates of both false positives and false negatives. While exact numbers cannot be estimated, it is possible to generate educated estimates. In this paper, I generate an estimate of the false negative rate for a representative sample of murders in Chicago from 1979. In this dataset, the majority of murderers were never convicted of killing someone. Although this error does not have the same moral weight as the conviction of a factually innocent person, it is not costless. Victims and their families are denied justice, and there is the risk that this person will continue to offend. The "true" cost of the tradeoff between the relative costs of false negatives and false positives is unknown—however, my estimates suggest that the system in Chicago in 1979 had sixty-six false negatives for every false positive, well above the mythical ten-to-one ratio suggested as ideal by Blackstone.

I also estimated that employers who use records of incarceration to identify people at risk for murder are acting as if they are willing to accept 930 false positives for every false negative. Both the nature of the decision by employers, and the consequences of the decision, are different than the decision by the criminal justice system. Employers are trying to predict the future, while the criminal justice system is trying to ascertain the truth about the past. Potential employees falsely identified as risky are simply denied employment, not incarcerated or even executed. Further, the two different decision makers weigh false negatives and false positives in very different ways. In the criminal justice system, false negatives are tolerated to avoid false positives, while in employment decisions, false positives are tolerated to avoid false negatives.

Despite these differences in answers, the basic approach to studying the question of error rates is the same in both contexts, and both contexts allow for the estimation of some reasonable estimates of the tradeoffs between false positives and false negatives. While at times it was cumbersome to consider both sets of estimates in the same paper, the effort should remind the reader that the tradeoff between the two types of errors is a universal constant in any decision process, and the basic conceptual framework transfers easily across contexts.

Future research should proceed in at least three directions. First, the estimates reported here should be replicated and then extended to other situations. While I am confident that my employment analysis will be very similar in other metropolitan areas, I am less sure that my criminal justice analysis will yield similar answers for other crimes, time periods, and geographical areas. Second, competent and comprehensive cost-benefit analyses should be done on the actual costs of the different kinds of errors in these two settings. The ratios of the costs of these two kinds of errors could then be compared with the actual ratios obtained from archival analysis like the one provided in this paper. Finally, research should study how different policy rules will affect both false positives and false negatives. This kind of research has been done in the area of eyewitness testimony,[45] and I see no reason why it could not also be done in other settings. Policymakers can, in theory, pick the desired tradeoff between false positives and false negatives. However, unless they understand how the different decision rules generate errors of both kinds, the idea of setting policy to achieve socially optimal levels of both errors will remain a theoretical possibility.

[45] *See generally* Wells et al., *supra* note 11.