ECONOMIC ANALYSIS OF WORKPLACE INJURIES

Kenneth Noble Fortson

A DISSERTATION PRESENTED

TO THE FACULTY

OF PRINCETON UNIVERSITY

IN CANDIDACY FOR THE DEGREE

OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE

BY THE DEPARTMENT OF ECONOMICS

NOVEMBER 2005

ABSTRACT

This dissertation comprises three chapters, each of which explores a different topic related to work injuries and the social insurance designed for such injuries, workers' compensation.

The first chapter examines workers' compensation in conjunction with an alternative means of compensating injured workers: the tort system. This chapter uses data from Texas—the only state not mandating workers' compensation—to determine how benefits paid under workers' compensation compare to awards injured workers would receive in the tort system. Average payments are substantially higher in the tort system, but an examination of the entire distribution suggests that only workers who have abnormally low levels of risk aversion would prefer the distribution of awards paid out in the tort system to the distribution of workers' compensation benefits.

Chapter 2 considers the consumption smoothing benefits of workers' compensation. Using data from the Panel Study of Income Dynamics, I examine how benefit generosity affects the drop in household food consumption following a work injury. The results suggest that in the absence of workers' compensation, consumption would fall by a substantially larger amount. However, further analysis suggests that current benefit levels are higher than the theoretically optimal level.

The final chapter uses data on workers' compensation claims from Texas together with Current Population Survey worker schedule data to examine how the workplace injury rate varies throughout the day. The injury hazard is substantially higher at night than during regular daytime work hours. As explanations for this result, I consider several factors that are correlated with night work but not inherent features of it. These include compositional changes in industry and age, as well as shift duration changes that would create artificial

variation in the injury rate; however, both of these fail to fully explain the higher nighttime injury rate.

ACKNOWLEDGMENTS

Many people have been instrumental in advising and otherwise helping me with this dissertation. Proper thanks for all of the comments and guidance I have received along the way would require another entire chapter; this section grossly understates my gratitude to everyone mentioned (and surely many others who are not).

Alan Krueger has been a fantastic advisor to me for the past four years. In addition to many conversations about my work, I have learned at least as much by observing his creative and energetic approach to his own research. Jeff Kling has been endlessly generous with his time, and has written tomes of insightful comments that greatly improved both this work and my general understanding of statistical methods. I am a far more competent economist than I would have been without the patient guidance of Alan and Jeff.

Harvey Rosen has been a kind mentor to me on matters both professional and personal, and has made the graduate school experience both smoother and more enjoyable. Hank Farber's enthusiasm for empirical work has infectiously inspired me, as have his regular smiling interrogations and discussions concerning how my work is progressing.

Several other Princeton professors and visitors have offered insightful advice along the way, especially Orley Ashenfelter, Olivier Deschenes, Bo Honore, Adriana Lleras-Muney, Bentley MacLeod, Jesse Rothstein, Ceci Rouse, and Albert Yoon. Fellow students including Melissa Clark, Bill Congdon, Anders Frederiksen, Helena Holmlund, Radha Iyengar, Alex Mas, Giovanni Mastrobuoni, Nick Racculia, Yaron Raviv, Brad Strum, Matt Weinberg, and Wes Yin have provided useful comments and stimulating discussions, and have more generally made work fun. Linda Belfield, Joyce Howell, Charlotte Howard, and Meghan Weeks have been invaluable in sorting out budgets, travel arrangements, and technological snafus.

My family has been supportive since I arrived at Princeton. My mom and dad kept me grounded through homesickness and other tough patches, always with a concerned but prudent word of encouragement. My brother and sister don't always understand what I am up to, but listen patiently while I prattle on about it. My mother- and father-in-law have defied stereotype to provide constructive comments and support.

Lastly, my wife Jane has been a willing ear for questions, complaints, and speculation. She has provided frequent assistance compiling and interpreting data, and even more important emotional assistance. Her indispensable work editing and proofreading many, many drafts of each chapter has significantly refined the papers herein.

TABLE OF CONTENTS

CH. 1. WORKERS' COMPENSATION AS AN ALTERNATIVE TO THE TORT SYSTEM

1.1. Introduction

Under the tort system in place at the turn of the twentieth century, a worker who was injured on the job and was seeking compensation had to prove in court that her employer's negligence was the proximate cause of injury. Because many industrial accidents were caused by seemingly inherent dangers of the work, fault was difficult to assign under this system (Fishback and Kantor 2000). Employers often defended themselves in such suits (with full or partial success) by demonstrating the employee's "contributory negligence," however minor. When the worker did prevail, the resulting award was frequently ruinous for the firm. The system was considered a failure for both parties: workers had very little chance of remuneration, even for severe, career-ending injuries, and employers could be financially devastated if they lost in court (Thomason, Schmidle, and Burton 2001).

In the last century, states have turned to workers' compensation insurance—originally called "workmen's compensation"—as a lower-variance alternative to the tort system. Now, in nearly all states, firms are mandated to purchase workers' compensation insurance from a state fund or a private carrier, or to self-insure. Only the smallest firms and a handful of narrowly-defined occupations are excused from the mandates.[1] Workers' compensation is now one of the largest social insurance systems in the United States. In 2001, 48.7 billion dollars in workers' compensation benefits were paid to injured workers, 50 percent more than the total amount paid in unemployment benefits (U.S. Census Bureau).

Under workers' compensation, when a worker is injured in the course of employment, the firm's insurance carrier pays for the worker's medical care. If the injury

[1] The specific exclusions vary by state, with the most common exemptions being for very small employers or for agricultural workers (Department of Labor 2004).

causes the worker to miss work, she is also paid some fraction of her average weekly wage while she recovers. To receive these benefits, the worker must only demonstrate that the injury was work-related—benefits are paid regardless of the relative negligence of the employee and her employer.[2] In exchange for this guaranteed outcome, an employee (implicitly) waives her right to sue her employer for work-related injuries. If she is injured on the job, the worker's sole course of financial remedy is through workers' compensation (Little, Eaton, and Smith 1999).[3]

However, little attention has been paid to the issue of whether workers' compensation in its present incarnation is indeed a good alternative to the tort system.[4] The purpose of this chapter is to examine just that—in particular, to assess whether non-medical benefits for injured workers covered by workers' compensation insurance are higher than awards paid to those in the tort system for otherwise comparable injuries. The workers' compensation system has evolved substantially since its inception. Beginning in the early 1990s, benefits began to decline as business leaders fought back against increasing insurance premiums; this followed a period of increasing benefits, especially through the 1970s (Thomason, Schmidle, and Burton 2001). As labor and business interest groups continue to wrestle over the sufficiency of benefit levels, determining whether current workers' compensation benefits adequately replace the awards paid under the tort system is important to inform the debate. Furthermore, it has been proposed that other areas of tort law in the

[2] Excepting injuries caused by the worker's intoxication, or injuries intentionally inflicted by the worker upon herself (Hardberger 2000).
[3] The employer is still liable for intentional acts of malice.
[4] An exception is work by Fishback and Kantor (2000), which examines the introduction of workers' compensation in the early 1900s. Using data primarily on fatally-injured workers from the turn of the century, they find that the introduction of workers' compensation insurance significantly improved the benefits paid to the families of workers who died in the course of their employment.

United States—most notably, medical malpractice—be replaced with workers' compensation-style systems, as has already been enacted in other countries.

In studying this question, I focus my analysis on firms in Texas. Texas is the only state where employers can effectively elect whether or not to subscribe to workers' compensation insurance, and as such, it provides a unique opportunity to directly compare the tort system and workers' compensation.[5] (Firms that elect to carry workers' compensation insurance are called "subscribers," while firms that do not are called "non-subscribers.") By looking closely at firms in Texas, I will be able to look at how outcomes for injured workers differ under the two systems.

My data on subscribing firms come from the Texas Department of Insurance's (TDI) Detailed Claim Information. To cull data on court cases in the tort system, I administered a survey of attorneys who have handled non-subscriber cases. Using propensity score matching, I study outcomes for comparable injury cases, estimating how workers in the tort system would have fared if they were instead employed by firms covered by workers' compensation.

The analysis focuses on the comparison of the income benefits received by injured workers employed by subscribing firms, and the verdict and settlement amounts of employees (of non-subscribers) who litigate. In particular, I am interested in differences in the consistency and timeliness of the two systems. Using the survey data on tort cases and a matched sample of workers' compensation cases, I find that the awards paid for verdicts and settlements in the tort system (net of attorney fees) are on average higher than the lost wage benefits paid by workers' compensation. However, comparing the quantiles of the two

[5] Technically, workers' compensation insurance is not mandatory in New Jersey. However, the structure of New Jersey law has induced all New Jersey employers to opt into the system (Shields and Campbell 2002).

systems reveals that the higher average award for tort cases obscures an extremely bimodal distribution that pays a handful of injured workers substantial amounts while many more receive little or no compensation. In contrast, indemnity benefits paid under workers' compensation have substantially lower variance. Furthermore, whereas benefits are paid almost immediately under workers' compensation, most tort cases take two or more years to resolve. I demonstrate that for plausible levels of risk aversion, an injured worker would prefer the distribution of benefits provided by workers' compensation, even though benefits are lower on average than awards paid in the tort system. This finding is especially remarkable given that my results are likely to understate the disparity between benefits under workers' compensation and awards in the tort system, as my research design has built-in biases that favor the tort system.

The next section provides institutional background about the two systems. Section 1.3 introduces the theoretical framework used to compare payments for injured workers under the two systems. The data are discussed in section 1.4, and section 1.5 explains the propensity score matching procedures employed to construct the matched samples. Section 1.6 discusses the outcome measures used in the analysis. Section 1.7 presents the results, and is followed in section 1.8 by an examination of which system is preferable for injured workers. Section 1.9 addresses the timeliness of the tort and workers' compensation systems, and section 1.10 discusses the results. I conclude in the final section.

1.2. Background

1.2.1. Workers' Compensation in Texas

When a worker at a subscribing firm is injured during the course of her employment, she must contact her employer, and must file a notice of injury with the Texas Workers'

Compensation Commission within one year of the injury. The injured worker may seek evaluation and treatment from the doctor of her choice; the insurer pays the costs associated with such medical treatment directly to the worker's health care provider. In addition, if for more than one week the injury either prevents her from returning to work or necessitates a temporary job reassignment, the worker becomes eligible for income replacement benefits. If the injury persists for 30 or more days, she becomes eligible for (retroactive) payment of income replacement benefits for the first seven days.

These lost wage benefits are calculated as 70 percent of the worker's pre-injury average weekly wage, or 75 percent if the worker was earning less than $8.50 per hour. The benefits are subject to a statutory minimum of 15 percent of the state's average weekly wage (about $81) and a statutory maximum of 100 percent of the state's average weekly wage ($562). In comparison, the tort system allows the worker to potentially be compensated for the full amount of her lost wages. Lost wage benefits are mandated under workers' compensation to replace only a fraction of foregone earnings in order to curb moral hazard on the part of the worker. At the same time, the purpose of the system is to prevent an injury from driving the worker into destitution, and indemnity benefits must be sufficient to accomplish this; determining and maintaining the appropriate balance between these two competing goals are delicate tasks that have troubled policymakers, especially in recent years.

Though workers' compensation insurance protects employers against many potentially devastating lawsuits, there are situations in which subscribing firms may be sued. In particular, the family of an injured worker is allowed to sue a subscribing employer for gross negligence, but only if the worker is killed (Hardberger 2000). Because in such cases the workers' families could potentially seek remuneration through both workers' compensation and the tort system, and because I do not have information on gross negligence lawsuits for workers' compensation cases, my analysis does not include cases involving fatal injuries.

1.2.2. The Tort System in Texas

Under the tort system, a worker who is injured in the course of her employment may sue her employer to seek benefits; in doing so, she must demonstrate that the employer's negligence caused the injury. Employers can be held liable for medical and lost wage damages, as well as non-economic damages, which include payments for pain and suffering. On January 1, 1991, the tort system was modified to favor the worker more than it previously had. Non-subscribing employers sued for negligence were stripped of the "assumption of risk," "comparative" or "contributory negligence," and "fellow servant" common law defenses (Sorrels and Ostrom 2000; Hardberger 2000). Put plainly, the "assumption of risk" defense allowed employers to argue that the worker understood the inherent dangers of the work, and by signing the employment contract, assumed the risks of the job. The second defense allowed an employer to concede its own negligence, but avoid paying the full amount of damages claimed by demonstrating that the employee was also negligent. Because employers were stripped of this defense, it is in theory sufficient to win the verdict for the worker to demonstrate that the firm was as little as one percent negligently responsible for the injury.[6] Finally, barring the "fellow servant" defense prohibits the employer from defending itself by arguing that it was the negligent actions (or inactions) of a fellow worker that caused the injury (Little, Eaton, and Smith 1999).

[6] The attorneys I have spoken with have suggested that, in practice, this one percent threshold is only loosely adhered to, as juries seemingly account for the comparative negligence of the two parties in calculating the award.

However, the tort system is not a strict liability system, unlike workers' compensation. Outcomes in this system are usually predicated on whether the employer had taken sufficient precautions and provided sufficient training (Sorrels and Ostrom 2000). Employers also sometimes question the circumstances of the injury, or argue that the worker has exaggerated the severity.[7] Nonetheless, while employers are far from defenseless, this tort system is more advantageous for the worker than a pure tort system with the full range of common law defenses.

1.3. Conceptual Framework

Identifying which of these two systems is better from the vantage point of the injured worker is complicated by the fact that, even independent of the *level* of remuneration, the two systems have very different *distributions* of remuneration. Whereas workers' compensation is designed to provide nearly all injured workers with fractions of their lost wages, the tort system will, by design, award some injured workers very high payments—with full replacement of lost wages plus payments for pain and suffering—but award many others nothing at all. Intuitively, the two systems may exhibit tradeoffs in terms of expected payment and variance, with workers' compensation having a more certain but lower outcome on average, and the tort system having a higher expected award but with greater variance.

To assess which system benefits injured workers more, I consider three alternative measures. Conceptually, the most straightforward of the three is the comparison of expected utility. While comparing expected payments would identify which system a risk-

[7] As with workers' compensation, employers are not liable for self-inflicted injuries or injuries caused by the worker's intoxication.

neutral injured worker would prefer, it is more plausible that injured workers exhibit some degree of risk aversion, in which case the spread of the distribution is also an important factor. Determining which system is preferable for an injured worker depends on which system yields greater expected utility (EU):

$$(1.1) \qquad EU_T = \int U(A)\, dF_T(A) \quad \substack{> \\ <} \quad \int U(A)\, dF_{WC}(A) \quad = \quad EU_{WC},$$

where $U(A)$ is the utility derived from an award of amount A, T and WC index the tort and workers' compensation systems, respectively, and F_T and F_{WC} are the distribution functions for the two systems. The relationship depicted in equation (1.1) will depend on the functional form of U, and in particular on the associated level of risk aversion. Consequently, in the empirical analysis that follows, I examine (1.1) over a range of assumptions about the level of risk aversion.

For policymakers, there may be different objectives than utility maximization, however. For example, we might only be concerned that injured workers are not impoverished. This criterion is more in line with a Rawlsian social welfare function that would rank the tort and workers' compensation systems based solely on the remuneration given to injured workers in the low end of each of the two systems' distributions of awards. Still another criterion would be to examine under which system a greater number of injured workers are better off. While the analysis and discussion in this chapter will focus on the expected utility comparison in (1.1), I consider each of these other two criteria as well.

1.4. Data Sources

1.4.1. Data on Workers' Compensation Injuries

Data on workers' compensation injuries come from the TDI Detailed Claim Information. These data report all relatively severe injuries between 1991 and 2003 at employers in Texas with workers' compensation insurance. Only injuries that had combined medical and indemnity benefits or expected benefits of at least $5000 are included. The data set includes a total of 148,807 closed, non-fatal cases that I use in the analysis. The data are also limited to single-injury cases to ensure high-quality matches with the tort cases, which I discuss in more detail in the next subsection.[8]

1.4.2. Data on Tort Cases

In order to collect data on tort cases, I conducted a survey of two groups of attorneys. The first "Small Group" consists of 494 attorneys who have each represented either a plaintiff or a defendant in a non-subscriber case that was reported in Westlaw's *Blue Sheets* for Texas. Because each lawyer in this group has dealt with a non-subscriber case in the past, I expect a relatively high response rate to the survey. I mailed each of the 494 attorneys in this group a paper version of the survey. In addition, I also provided a user name, password, and URL that the attorney could use to take an online version of the survey that mirrored the paper version. The survey was designed to take 15 to 20 minutes to complete. Attorneys who completed the survey—either the paper version or the online version—had the option of receiving a $20 check as a token of appreciation.[9]

---

[8] Approximately two-thirds of the TDI workers' compensation claims data are single-injury cases.

[9] The check was not intended to compensate the attorney for her time, as this amount is very low compared to the hourly rate of most attorneys. However, the cover letter encouraged

A concern with the Westlaw cases, however, is that they contain a disproportionate share of verdicts and very few settlements.[10] Mechanically then, the cases that the attorneys in the Small Group report may also contain a disproportionate number of verdicts compared to settlements.[11] Consequently, I supplemented the cases reported by the Small Group by surveying a much larger but less precisely targeted group of attorneys. The "Large Group" consists of all attorneys who are board-certified by the Texas State Bar in either personal injury law or employment and labor law, or are members of the Employment and Labor Law Section of the Texas State Bar.[12] I mailed each of the 4341 attorneys who composed the Large Group a letter that provided a user name and password and invited them to complete the survey online.[13] Like those in the Small Group, attorneys in the Large Group were each offered $20 for completing the survey; the survey content was identical for both groups.

Following two screening questions asking whether the attorney has recently represented a worker or an employer in a non-subscriber tort case, the survey asks the attorney for such background information as the size of the attorney's firm, how long she has been practicing law, and the number of non-subscriber cases she has handled in the past five years. The survey then proceeds to ask for the details of the attorney's most recently completed non-subscriber case, followed by the details of her second most recent.

---

attorneys to give the survey to paralegals to complete. Many attorneys appear to have done this, and others requested that I donate their $20 to charity. Still others politely declined the check.

[10] Aside from concerns about how representative they are, the Westlaw cases also do not contain sufficient information to allow me to include them in the analysis.

[11] Indeed, a much higher fraction of the cases in my sample that come from the Small Group were tried to a verdict compared to cases from the Large Group.

[12] Roughly one-third of the Small Group belongs to at least one of these three groups. Attorneys in the intersection of the Small and Large Groups were treated as part of the Small Group and are counted as such in the analysis.

[13] Due to the size of the Large Group and the anticipated low yield, I did not include paper versions of the survey in mailings to the larger group except upon request.

Attorneys are asked to describe the nature and circumstances of the injury and to report the outcome of the lawsuit, including whether the case resulted in a settlement or verdict as well as a breakdown of any monetary award. The survey also included questions about the worker's industry of employment, pre-injury annual earnings, amount spent on medical care, and attorney fees.

Responses were collected over a six-week period. Two and a half weeks into the data collection, I sent reminder postcards to both groups.

Table 1.1 reports the number of responses from each of the two groups. Eighteen percent (86 out of 482) of the attorneys in the Small Group responded, excluding 12 attorneys for whom I had incorrect mailing information. The three subsections of the State Bar that I surveyed were those most likely to contain potential non-subscriber case attorneys, but I anticipated that only a relatively small fraction of these attorneys would have handled non-subscriber tort cases. Indeed, ten percent of the Large Group responded, with 198 of the 431 reporting that they had handled non-subscriber cases in the past, and 233 reporting that they had not.

It is likely that the lower response rate of the Large Group is indicative of the lower proportion of the group that has handled non-subscriber tort cases.[14] To examine this, I randomly selected a set of 45 non-respondents in the Large Group and attempted to conduct a short phone interview with each of them. Of the 39 that I was able to contact, only 9 had handled non-subscriber cases in the last five years. Assuming that these 39 attorneys are representative of the larger group of non-respondents, this suggests that the

[14] Of course, there are other factors that probably contribute as well, most notably the fact that the Large Group's survey was administered almost exclusively online, while the Small Group could take either the paper or the online version. Over 70 percent of Small Group respondents used the paper version, suggesting it may have elicited a higher response rate.

response rate among Large Group members who handle non-subscriber cases is about 18 percent.

The response rate in my survey is well within the range of typical response rates to establishment surveys. The response rates for the surveys employed in eight studies reviewed by Kling (1995) averaged 26 percent, with a high rate of 66 percent and two surveys with response rates as low as two percent. It is also comparable to the response rate for the survey of employers in Texas by Shields and Campbell (2002): excluding the 17 percent of their original sample for whom the contact information was poor, their response rate was 20 percent. In appendix 1.C, I examine the geographic distribution of respondents and non-respondents; the limited evidence I have suggests that the respondents in my sample do not systematically differ from non-respondents.

The survey yielded a total of 251 cases that were sufficiently complete to consider.[15] I excluded 14 cases that involved a worker fatality, because, as discussed in section 1.2, these cannot be precisely compared to cases in the workers' compensation data, since the worker's family is not prohibited from filing a gross negligence lawsuit in the event of the worker's death, even if the employer had workers' compensation.

Forty-one of the 251 tort cases involve multiple injuries. Accurately matching these cases to workers' compensation cases is difficult, however, because the TDI data do not specify the exact nature of injuries when there is more than one. An additional 19 cases involved little or no time off from work and low medical expenditures, and these injuries were unlikely to be severe enough to exceed the $5000 threshold of the workers'

[15] The instructions asked for information about the attorney's two most recently completed cases; nonetheless, a handful of attorneys reported cases that were still pending. These cases are excluded. There were also two situations in which the same case was reported by two different attorneys. For each case, one of the duplicates was discarded.

compensation cases. To facilitate high-quality matches between the tort cases and workers' compensation cases, the 19 "small potatoes" cases and 41 multiple-injury cases are excluded from the analysis.

Of the remaining 191 tort cases, 16 are settlements with insufficient award information. While I cannot include them in the analysis, simply dropping them could bias the results. Hence, I construct weights that account for these cases by giving settlements that are in the analysis extra weight, such that the weighted share of cases that settled is the same as if the 16 settlements with insufficient information were included.[16]

All told, the survey yielded 175 cases that were sufficiently complete with single, non-fatal injuries, in addition to the 16 settlements with missing award amounts. Table 1.2 breaks down the tort cases by outcome. Nearly 60 percent of tort cases settled out of court, and factoring in the cases that were dropped, only 32 percent of tort cases in the sample were tried to a verdict. No evidence that I am aware of exists about the fraction of non-subscriber cases in Texas that settle. At least compared to other tort systems, though, the trial rate in my sample is relatively high. Eisenberg and Farber (1997), for example, study federal civil litigations, and in their data, less than five percent of cases are tried to a verdict. I discuss the implications of this finding in section 1.10.

Table 1.2 also shows that the worker prevailed in 44 percent of the cases in court (27 out of 62), while the verdict was in favor of the employer in 56 percent of cases. This deviates notably from the 50-50 breakdown predicted by Priest and Klein (1984), although this difference could possibly be sampling error.[17] Eisenberg and Farber (1997) develop an

[16] This weighting procedure is discussed in more detail in appendix 1.A.
[17] If the cases are independent, then the number of verdicts won by the plaintiff is a binomial random variable. Under the null hypothesis that the true chances of the employer and the

alternative theoretical model in which the plaintiff win rate is inversely related to the fraction of cases that go to trial. Following their model, the low fraction of plaintiff verdicts observed in my sample may reflect the higher propensity to go to trial in non-subscriber tort cases.

1.5. Matching Procedure

1.5.1. Propensity Score Matching

Thirty-five percent of Texas employers do not subscribe to workers' compensation insurance. However, because many of these firms are small, non-subscribing firms employ only 16 percent of Texas workers. High quoted premiums and a lack of control over medical providers are two primary factors in the decision not to subscribe (Shields and Campbell 2002). Not surprisingly, there is substantial variation across industries in who chooses to subscribe. Table 1.3 breaks down the non-subscription rates by industry. Mining has the lowest non-subscription rate, possibly reflecting the dangers of mining occupations, whereas retail trade has a very high non-subscription rate. However, manufacturing and agriculture also have considerable non-subscription rates, despite being relatively risky, and finance, presumably fairly unrisky, has a very low percentage of workers who are not covered by workers' compensation.

As is clear, employers do not randomly opt out of workers' compensation insurance. Instead, industry and the size of the firm play important roles in the firm's decision to subscribe or not to workers' compensation. Even conditional on a worker actually getting injured, there is still likely to be selection into the tort system based on the severity of the

---

worker winning the verdict are equal, the probability that the worker will win 27 or fewer verdicts out of 62 is 19 percent.

injury. Workers with very minor injuries are unlikely to be willing to spend the time and effort to pursue litigation, especially when doing so could strain or jeopardize their employment. They might also have difficulty finding legal representation willing to take their cases. I employ propensity score matching to help control for case circumstances and characteristics, and to help alleviate these selection concerns. (For discussions of propensity score techniques, see Rosenbaum and Rubin 1983 and, more recently, Heckman, Ichimura, and Todd 1998 and Dehejia and Wahba 2002.) For each tort case, I seek workers' compensation cases that provide an appropriate counterfactual, i.e., what would have happened if the injured worker in the tort case had instead been in the workers' compensation system.

The basic approach is to estimate the probability of being a tort case rather than a workers' compensation case as a function of worker and employer observables. Using the propensity score—namely, this probability—I then find the three "nearest neighbors" in the workers' compensation data for each tort case. Because of the size of the workers' compensation sample, I am able to match without replacement, facilitating more precise estimates.

I use a probit to estimate the propensity score $(\hat{P}_i^{tort})$ as a function of the body part injured, nature of the injury, and cause of the injury (collectively represented by the vector $\mathbf{I}_i$); industry and categorical size of the employer $(\mathbf{E}_i)$; and a fourth-order polynomial for the worker's pre-injury annual earnings $(\mathbf{W}_i)$:[18,19]

$$\hat{P}_i^{tort} = \Phi(\mathbf{I}_i'\hat{\beta}^{injury} + \mathbf{E}_i'\hat{\beta}^{employer} + \mathbf{W}_i'\hat{\beta}^{earnings}). \tag{1.2}$$

Following Dehejia and Wahba (1999), I isolate the observations on the common support of the two propensity score distributions by discarding 80,867 workers' compensation cases that have propensity scores below that of the lowest tort case. The distribution of propensity scores for the workers' compensation cases and tort cases on the common support is provided in figure 1.1. The support of the distribution is narrower than the full unit interval possible; none of the tort cases or the workers' compensation cases has a propensity score greater than .60. This occurs because there are so few tort cases compared to workers' compensation cases. Indeed, in these circumstances, even an estimated propensity score as low as .05 is very high.

Not surprisingly, the tort cases are more concentrated in the higher bins than the workers' compensation cases; however, because the TDI data have such a large sample size, there are a sufficient number of workers' compensation cases to provide three high-quality matches for nearly all of the tort cases. Among the few observations with propensity scores greater than .25, there are fewer good matches, but still at least one for each tort case.

---

[18] Amemiya (1981) notes that probit and logit models can yield different results if the dependent variable has very few of one outcome compared to the other. Since there are so few tort cases and so many workers' compensation cases, this is a legitimate concern with my data. However, re-estimating the model using a logit distribution yielded comparable results.

[19] To distinguish abdominal hernias from herniated discs, I also include an interaction term that identifies herniated discs.

Because I conduct the matching without replacement, the order in which I match could be important (Rosenbaum 1995). I choose to randomly determine the order in which tort cases are matched to workers' compensation cases. To evaluate the sensitivity of the results to this choice, I also consider an alternative procedure to determine matching order. In analysis not presented here, I rank the tort cases by their observational weight, using the composite weights described in appendix 1.A. Such a procedure would ensure that the tort cases with the most influence on the parameter estimates would also have the highest-quality matches. This procedure had little impact on the matches, and so I employ the standard, simpler procedure of randomly ordering the data before matching.

The sample that I use in my analysis includes 175 tort cases matched to 525 workers' compensation cases. Table 1.4.a breaks down the type of injury for the full workers' compensation sample (column 1) and the matched workers' compensation sample (column 3), and compares them to the tort cases (column 2). Table 1.4.b presents comparable results for body part of injury. In table 1.4.a, the differences between the full sample and the tort cases are most pronounced for more severe injuries: there are relatively more amputations, fractures, and hernias in the tort sample than there are in the full workers' compensation sample.[20] The reverse is true for less severe injuries, especially contusions and strains. In contrast, the tort and matched workers' compensation samples are quite similar for both the injury type distribution in table 1.4.a and the body part distribution in table 1.4.b. The

[20] The disparity is especially large for hernias. On the whole, the injury types are well-defined and specific; however, an exception is the herniated spinal disc category. It appears that such injuries are sometimes classified as hernias, as dislocations, or as sprains by the insurance companies who file the Detailed Claim Information. To ensure that tort cases involving herniated discs are properly matched, I classify these injuries as hernias, since this is the only category for which I can be certain that all of the workers' compensation cases are truly herniated discs.

propensity score method also very successfully balanced the distribution of injury causes, industrial mix, and annual pre-injury earnings, which for brevity I do not report here.

1.5.2. Other Selection Concerns

The key identifying assumption of a propensity score approach is that, conditional on the propensity score, observations are as good as randomly assigned to the treatment group (in this case, the tort cases) and the control group (workers' compensation cases). However, there are several avenues not encompassed in my estimation of the propensity score through which non-random selection is possible.

First, workers could conceivably select into jobs with (or without) workers' compensation within the same industry. The evidence available suggests that this is not a substantial concern, however. Legally, workers employed by subscribers to workers' compensation can themselves opt out of the employer's coverage. In practice, virtually no one opts out, suggesting that there is little initiative on the part of the worker to actively select into subscribing or non-subscribing status (Robertson 1993). Madrian and Shea (2001) find that workers adhere strongly to the employer's default in participating in 401(k) plans, and it seems likely that this behavioral bias toward the status quo would apply to workers' compensation as well.

Previous research has also considered the question of whether workers' compensation insurance leads to moral hazard, i.e., whether firms covered by this insurance take fewer safety precautions. Evidence for such a phenomenon is weak at best: Krueger and Burton (1990), for example, find that workers' compensation costs have increased almost dollar for dollar with benefits, and that if there are any moral hazard effects on employers or workers, they offset each other. Shields and Campbell (2002) find that there do not appear to be appreciable differences in safety precautions between subscribing and non-subscribing employers of comparable size. Butler and Worrall (1991) find evidence that higher benefit levels may cause higher claim rates, but do not lead to diminished safety. In Texas, it is not clear if safer firms would be more likely to opt into or out of workers' compensation. Employers who are exceptionally safe may be more likely to decide that insurance is unnecessary. At the same time, exceptionally unsafe firms may have a harder time securing affordable coverage, inducing them to "go bare." Even in the presence of differing levels of safety precautions, the fact that my analysis focuses on ex post worker injuries means that the differing safety measures can only account for observed differences in the awards and benefits under the two systems if injury severity varies systematically between the two systems *within a given injury category*. Given the precision of the injury categories, this is unlikely to be a problem.

A final source of bias that should be considered comes from the probable event that injured workers employed by non-subscribers will elect to sue only if it increases their expected utility to do so. Employees with weak cases may decide that it is not worth the effort—or risk of antagonizing their employers—to sue. A separate issue is that small employers might have few assets, and even a successful case against them will yield little or no recovery of damages for the worker. In both cases, the worker may have difficulty finding an attorney to take the case. Since these weak cases do not show up in my sample of court cases, my analysis will be biased in favor of the tort system. Hence, the disparity between the payments under workers' compensation and the tort system that I report in section 1.7 understates the full extent to which the distribution of payments is preferable for workers under workers' compensation.

1.6. Outcome Measures

Because workers' compensation pays for all of the employee's medical expenses, medical benefits under workers' compensation are widely considered at least adequate; any debate about medical benefits centers on whether workers' compensation medical benefits are *too* generous.[21] In contrast, the adequacy of indemnity benefits is much more heavily disputed (see, for example, Thomason, Schmidle, and Burton 2001). Hence, I focus on non-medical remuneration in my analysis, comparing indemnity benefits to non-medical tort awards.[22]

For those in the tort system, remuneration is measured in two ways. The first measure is the total non-medical award from the verdict or settlement, net of legal fees and expenses:

(1.3) $$\textit{Net Non-Medical Award}_i = \textit{Total Award}_i - \textit{Medical Expenditures}_i - \textit{Fees}_i - \textit{Expenses}_i$$

where $\textit{Medical Expenditures}_i$ is the amount that worker *i* spent on medical care, not including any medical benefits paid by the employer. ($\textit{Net Non-Medical Award}_i$ and all subsequent monetary outcomes are measured in 2003 dollars.)

Several defense attorneys who responded did not know the exact amount of legal fees or expenses paid by the worker. In such cases, I impute attorney fees as one-third of the total award and expenses as four percent of the award. In my data, and according to the lawyers in Texas with whom I have spoken, a contingency fee equal to one-third of the total

---

[21] This is especially true in Texas, where covered workers not only receive full medical benefits, but also have complete choice of treating physician. Most efforts to reform the medical benefits of workers' compensation concern the administration of the plan, such as whether fee schedules should be implemented, or whether the insurance company should be able to restrict the set of treating physicians an injured worker can visit.
[22] In separate analysis not reported here, I included medical awards (for tort cases) and benefits (for workers' compensation cases). As expected, this inclusion strengthens my findings.

award is by far the most common rate. Nearly all of the rest charge a contingency fee of 40 percent; thus, my imputations for missing attorney fees may understate true attorney fees and hence may overstate the true net non-medical award. Legal expenses are on average slightly less than four percent of the total award in my sample without missing data. More importantly, though, legal expenses are more non-linear than my imputation implies, as there is an apparent fixed cost component of legal expenses. Hence, my imputations may understate legal expenses (and overstate net awards) more for lower awards than higher awards. It should be emphasized that understating legal fees and expenses for tort cases where the amounts were not known will bias my results in favor of the tort system.

Firms who do not subscribe to workers' compensation still often choose to provide alternative benefits to injured workers, however. This can impact the damages claimed by workers in lawsuits. For example, if a non-subscribing employer pays an injured worker for all of her lost wages as part of an alternative benefits package, the worker would probably not claim damages for lost earnings as part of the suit. Instead, the worker would claim only non-economic damages, such as pain and suffering. In this case, the measure of $\textit{Net Non-Medical Award}_i$ defined above would fail to capture the total amount of non-medical compensation the worker received for her injury. Indeed, in a survey of employers in Texas, Shields and Campbell (2002) find that 39 percent of non-subscribing employers in Texas pay income benefits of some kind. Because larger employers are more likely to provide such benefits, 65 percent of the workers in the non-subscribing workforce could potentially receive such benefits.[23]

[23] Of the non-subscribing employers providing alternative benefits, however, only 35 percent have formal, written plans. This suggests substantial variability in the coverage and administration of such alternative benefits programs.

Table 1.5 reports the fraction of workers in my sample of tort cases who received either lost wage benefits or medical benefits. A small percentage of respondents did not know if the worker received benefits or not. Of the remainder, well over half received medical benefits and 40 percent received lost wage benefits of some kind. Interestingly, while the latter percentage is high, it is well below the 65 percent of the non-subscribing workforce that is employed by firms providing such benefits (Shields and Campbell 2002). There are several possible explanations. The simplest is that some employers pay benefits only for certain injuries or in limited circumstances. Frequently, alternative benefit plans will cover temporary disabilities but not permanent or chronic injuries. It is likely that there are case selection issues at play as well. For example, workers who receive benefits are probably less inclined to file lawsuits, both because they may have less financial need and because they are less likely to perceive that injustice was done. Although this chapter is unable to address this phenomenon, it certainly merits future investigation.

As discussed previously, the share of cases in which the worker received income benefits is high and should be accounted for in the analysis. To this end, I construct a second measure of remuneration for tort cases by tacking onto $\textit{Net Non-Medical Award}_i$ any lost wage benefits paid by the employer. This is not necessarily a superior measure, however. While it captures additional payments that the injured worker would be entitled to if successful in court absent employer benefits, it may also overstate the award in cases where the worker was paid income benefits but did not have a strong lawsuit. Monetary outcomes under a tort system without employer-provided benefits would lie somewhere between these two measures of remuneration. While neither measure on its own can completely replicate outcomes under such a system, analyzing them together allows us to pin down informative lower and upper bounds on the generosity of a hypothetical system in which litigation is the only resort for an injured worker.

For workers' compensation, I have only one measure of remuneration: indemnity benefits. As with my measure of remuneration for workers in the tort system, I calculate benefits for workers covered by workers' compensation net of attorney fees. Although most workers' compensation claims are filed without the help of an attorney, the system can be decidedly complex when a severe injury is involved, and it is not uncommon to seek legal aid. Workers hired attorneys in 10 percent of the cases in my workers' compensation sample. Unfortunately, TDI does not collect information about the amount of legal expenses paid by the worker. However, state law restricts attorney fees to be no more than 25 percent of the total income benefits that the worker receives, and so I measure net benefits for injured workers under workers' compensation insurance as follows:

(1.4) $$Net\ WC\ Benefits_i = (1 - 0.25 \times Attorney_i) \times Indemnity_i$$

where $Attorney_i$ is a dummy variable equal to one if worker *i* hired an attorney and equal to zero otherwise, and $Indemnity_i$ is the amount of lost wage benefits paid to the worker. It is worth noting that, in contrast to my imputation for attorney fees in the tort system, my imputation of attorney fees for those covered by workers' compensation as 25 percent of lost wage benefits may *overstate* attorney fees and, consequently, may understate the net workers' compensation benefits for workers who hire attorneys. Once again, this bias works in favor of the tort system and against workers' compensation.

This measure of net workers' compensation benefits is used to compare remuneration under workers' compensation insurance to each of the two measures of the net award in the tort system. While workers' compensation provides nearly certain payment, benefits do not fully replace lost wages. The tort system, however, can award injured workers full replacement of lost wages. Perhaps more importantly, injured workers in the tort system can also sue for non-economic damages, such as pain and suffering. It is this trade-off between certainty and size of the award that I address in the econometric analysis presented in the next section.

1.7. Results

1.7.1. Non-Medical Awards vs. Indemnity Benefits

I employ a number of weighting specifications in the results that follow. Here, I briefly discuss their purposes; the specific construction of the weights is discussed in appendix 1.A. First, as mentioned in section 1.4, in all specifications the weights account for tort cases that settled but did not have the amount of the settlement reported. I also construct weights to make the composition of industries represented in my sample comparable to the distribution of industries in the population of injured workers at non-subscribing firms in Texas. Appendix 1.D reports additional results that are weighted to account for differences in attorney caseload. Including such weights strengthens the main results; hence, I focus on the specifications that do not weight for attorney caseload.

Table 1.6 reports the difference between $Net\ Non\text{-}Medical\ Award_i$ in the tort cases and net benefits in the matched workers' compensation cases. The first column shows the difference between the means of the two groups. Considering the simplest weighting scheme used in the first panel, which only accounts for the settlements with missing award amounts, the average tort case receives an award of $24,805 more than the average workers' compensation case benefit level, $19,869. This suggests that, on average, an injured worker in the tort system receives remuneration that is well over twice the size of payment received by an observationally-similar injured worker in the workers' compensation system.

However, the apparent advantage of the tort system appears to be concentrated at the upper end of the distribution. The remaining columns in table 1.6 report estimates from quantile regressions of the $\textit{Net Non-Medical Award}_i$ on an intercept and a tort case dummy with no other controls. Conceivably, including additional controls could improve the precision of the estimates. Appendix 1.D discusses estimates that include controls for other covariates; the point estimates are similar to the results reported here, and including additional covariates does not appreciably improve the standard errors. The coefficient on the tort case dummy variable from the quantile regression for the p^{th} percentile is simply the difference between the p^{th} percentile of the distribution of tort cases and the p^{th} percentile of the distribution of workers' compensation cases.[24] In the first panel, the median (50th percentile) tort case receives \$9626 less than the median workers' compensation case (\$11,626)—implying that the median award for a tort case is only 17 percent of the median for an otherwise comparable workers' compensation case. The quantile differences are large and heavily favor the workers' compensation cases through the 70th percentile, and the standard errors (in parentheses) suggest that the differences are significant not only in economic terms, but in statistical terms as well. Standard errors are computed using bootstrapping to account for both the heteroskedasticity that is inevitable in a sample of court cases, as well as the fact that estimation error in the propensity score can lead to variation in the matches (Dehejia and Wahba 2002).[25] The standard errors are also clustered to allow for correlation between two cases reported by the same attorney. Although the 75th percentile for the tort cases is more than \$9000 lower than for the workers' compensation

[24] For an introductory discussion of quantile regression, see Koenker and Hallock (2001).
[25] The bootstrap procedure uses 200 replications. Bootstrap samples are drawn from each of the original data sets, and each replication re-estimates the propensity score before matching the samples.

cases, it is only statistically significant at the 10 percent level. At the 85th percentile, the tort system has a higher award, and the gap grows substantially after this point. The $24,805 gap between the mean award for the tort system and the mean benefit level for workers' compensation is entirely concentrated in the upper quintile of the distribution, with workers' compensation dominating for the lower four quintiles of the distribution.

The estimates in the second panel use the industry weights discussed earlier. The quantile regression results are almost identical to those in the first panel, suggesting that my results are not driven by sampling differences in the industrial mix. However, the mean award for the tort system is somewhat lower in the second panel compared to the first.

1.7.2. Non-Medical Awards Plus Alternative Wage Benefits vs. Indemnity Benefits

Table 1.7 reports estimates analogous to those in table 1.6, with the difference being that the measure of non-medical awards for the tort cases is augmented with any additional wage benefits paid by the injured worker's employer. In several of the tort cases, the attorney knew that the worker had received lost wage benefits but did not know the amount. I construct weights similar to those constructed for missing settlement amounts to account for the cases with missing benefits. It should also be noted that, while the measure of monetary remuneration for workers' compensation cases is the same in table 1.7 as it is in table 1.6, the point estimates differ because excluding tort cases with missing benefit amounts also changes the sample of matched workers' compensation cases.

The general pattern of results is similar when the tort case awards include alternative wage benefits to when the awards do not, but not surprisingly, the results do not favor workers' compensation quite as strongly. In both panels, the estimated quantile difference between the tort system and workers' compensation is negative and statistically significant until just before the 60th percentile, and the tort system has a higher point estimate beginning at about the 75th percentile.

Table 1.8 breaks down the distribution of awards in the two systems. The bimodal nature of the tort distribution is clear: whereas only 5 percent of workers' compensation cases do not receive any lost wage payments, 38 percent of tort cases receive nothing. At the same time, one in eight tort cases receives an award in excess of $70,000, compared to just over three percent of workers' compensation cases. This same phenomenon is depicted in figure 1.2, which displays kernel density estimates of the distribution of net non-medical awards including wage benefits (for tort cases) and net indemnity benefits (for workers' compensation cases), weighted to account for missing benefits and missing settlement amounts as in the first panel of table 1.7.[26] The graph includes only the 62 percent of tort cases and 95 percent of workers' compensation cases with awards greater than zero (i.e., rather than integrating to one, the density function integrates to .62 for tort cases and .95 for workers' compensation cases). The distribution is also top-coded at $70,000 to highlight the bimodality of the tort system. While the tort system has far more cases with very high payouts, it also has far more injured workers who receive no remuneration. In contrast, workers' compensation payments are heavily concentrated between $4000 and $32,000. From that point up until $68,000, the two distributions look relatively similar.

Figure 1.3 presents the cumulative distribution function (CDF) associated with the density function plotted in figure 1.2. In figure 1.3, the zeros are included. The horizontal distances between the two distributions correspond to the quantile differences estimated in table 1.7. Consistent with the results in that table, the two CDFs cross just past $28,000, the point which nearly 75 percent of each distribution is at or below. For each award value

[26] These estimates use a Gaussian kernel with a halfwidth of 1000.

below \$28,000, the tort system has a larger percentage of cases at or below that award, while for values above \$28,000 the CDF instead favors the tort system.

1.8. Are Injured Workers Better Off?

The results discussed in the previous section have demonstrated that until the upper percentiles, benefits under workers' compensation are higher than awards in the tort system (with or without alternative benefits included). This suggests that if our criterion for choosing which system is "better" for injured workers is based on which system provides higher payments for the most workers, as discussed in section 1.3, workers' compensation is better. Similarly, the fact that 38 percent of injured workers involved in tort cases receive no non-medical award versus only 5 percent of injured workers in the matched workers' compensation sample suggests that if we adopt the criterion that the better system is the one that pays injured workers at the bottom of the distribution more, workers' compensation is also preferable.

Alternatively, we can evaluate the two systems based on the expected utility of injured workers. The average award is higher in the tort system in all of the specifications. Hence, a risk-neutral worker who is injured would unambiguously prefer to be in the tort system. However, injured workers are unlikely to be risk neutral. For a risk-averse injured worker, we cannot say which system is preferred without knowing how risk averse she is. I assume injured workers have exponential utility functions defined as follows:

$$U(A) = -\exp(-\lambda^a \times A), \tag{1.5}$$

where A is the amount of the award (as in section 1.3) and $\lambda^a \in (0,1)$ is the Arrow-Pratt coefficient of absolute risk aversion $(= -U''/U')$, with risk aversion increasing with λ^a.[27] I then use the utility function represented in (1.5) coupled with the kernel density estimates (recomputed without censoring the higher awards at $70,000) of the awards under the two systems—with alternative employer benefits included for the tort cases—and factor in the zeros to construct estimates of the distribution functions F_T and F_{WC} in equation (1.1) necessary to compute expected utility for each system. For values of the coefficient of absolute risk aversion λ^a greater (less) than .000016, the distribution of awards provided by workers' compensation yields higher (lower) expected utility than the distribution of awards in the tort system. The associated coefficient of *relative* risk aversion, $\lambda^r = -A \times U''/U'$, evaluated at A = $12,000 is .19.[28]

The requisite levels of risk aversion computed above suggest that, for the tort system to have a higher expected utility than workers' compensation, workers would have to be more risk neutral than is consistent with most estimates. Arrow (1976) originally argued that the coefficient of relative risk aversion would be close to 1, and most current studies claim plausible levels of relative risk aversion are closer to 4, with absolute risk aversion between

---

[27] It may be more intuitive to think of utility as a function of total wealth, i.e., U as a function of pre-existing wealth W_0 plus the award A, instead of A alone. However, because wealth enters the function exponentially, replacing A in equation (1.5) with $W_0 + A$ amounts to the monotonic transformation of (1.5) that multiplies $U(A)$ by the constant term $\exp(-\lambda^a \times W_0)$. Hence, the conclusions drawn from considering the case where initial wealth is zero carry over to the more general case where initial wealth is positive.

[28] With an exponential utility function, relative risk aversion increases linearly with A and depends on whether initial wealth is considered or not. $12,000 is approximately equivalent to the median of the workers' compensation distribution. Alternatively, we could evaluate λ^r at A = $32,000, which is approximately equal to the median and mean pre-injury annual earnings for the two samples ($20,000) plus $12,000. This yields an estimated coefficient of relative risk aversion of .51.

.0001 and .0003.[29] Using levels of risk aversion that earlier research finds are consistent with behavior, the distribution of benefits in the matched workers' compensation cases provides injured workers with greater expected utility than the distribution of awards observed in the tort system. Furthermore, this analysis has used the measure of tort awards that includes alternative employer benefits; if we instead exclude such benefits and only consider the award received through the tort process, the distribution of workers' compensation benefits is even more favorable for injured workers.

1.9. Expedience of Awards

The previous two sections discussed differences in the distributions of awards under workers' compensation and the tort system. Although the amount of remuneration is measured in constant (2003) dollars in all cases, the awards have not been adjusted to account for discount rates.

My data on worker's compensation claims have information only on the year that the worker was injured and the year that she first received indemnity payments. Of the workers' compensation cases that received payment, 93 percent received payment within the same calendar year as the injury, and another 6 percent received payment in the next calendar year. In contrast, for tort cases that received any award, the calendar year of the verdict or settlement was the same as the calendar year of the injury in only seven percent of cases. Twenty-three percent, 37 percent, and 25 percent of the awards were received in the next year, two years later, and three years later, respectively. If I assume that injured workers in the tort system who received alternative wage benefits were paid immediately, then the

[29] See, for example, Bliss and Panigirtzoglou (2004), Aït-Sahalia and Lo (2000), Gomez-Limon, Arriaza, and Riesgo (2003), and Parker (2001).

fraction of workers in the tort system who first received payment in the same calendar year as their injuries rises to 50 percent, with 15 percent in the following year. However, this is still much less timely than workers' compensation payments.

In table 1.9 I report the distribution of time elapsed between the injury and the verdict or settlement for workers in the tort system. A pattern emerges similar to that discussed above for calendar years: only 4 percent of the awards are issued within 6 months of the injury, with an additional 11 percent coming within the first 12 months.

Once again, assuming that those who were paid wage benefits received them immediately, the percent who received payment within six months of injury rises to 46, with another 8 percent receiving payment in the next six months. This is still well below the standard set by workers' compensation. All told, the evidence presented here strongly indicates that accounting for time preferences in comparing remuneration in the two systems would tip the scales even more in favor of workers' compensation.

1.10. Discussion

1.10.1. Internal Validity

There are a number of potential sources of bias in my analysis. Most notably, as I have discussed earlier, the imputation methods that I employ are biased in favor of the tort system by design. An important systemic source of bias mentioned before is the election of injured workers who are not covered by workers' compensation to litigate or not. The fact that workers with weaker cases are presumably less likely to litigate implies that there may be many injured workers who are employed by non-subscribing firms and receive little or no remuneration; my survey method will fail to pick up these injured workers. Many workers employed by small firms may also opt not to litigate if the firm has insufficient assets to pay for any award won in court.

A related concern in the workers' compensation data is that only successful claims are represented in the data. To the extent that some tort cases that lose in court would have also been denied under workers' compensation, and because denied claims do not appear in the TDI data, this would bias the results in favor of workers' compensation. Work by Biddle (2001) suggests a ballpark workers' compensation claim denial rate of 15 percent. Assuming that an additional 15 percent of workers' compensation cases would have received no payment and employing the analysis similar to that in section 1.8, the requisite level of absolute risk aversion for the distribution of workers' compensation payments to be preferred to the tort distribution approximately doubles to .000032. However, this is still well below the range of estimated values of absolute risk aversion, suggesting that accounting for denied claims does not change the qualitative results.

The analysis presented in the previous sections has also included responses from attorneys in both the Small Group and the Large Group. Section 1.4 discussed concerns that the cases reported by the Small Group may have a disproportionate number of verdicts and too few settlements. However, over most of the distribution, the quantiles for tort case awards are actually *greater* for the Small Group than the Large Group; any bias induced by including the Small Group would favor the tort system. Hence, to the extent that there is such bias, the results understate the favorability of workers' compensation. Settlements and dropped cases may also be underrepresented in my sample for other reasons. Cases that are taken to trial are fewer and farther between, and as such, attorneys may recall them more easily, whereas settlements (and especially dropped cases) are shorter affairs. Respondents also may not have initially realized that dropped cases qualified as "completed." To address this, I analyzed the settlements and dropped cases (and their matched workers' compensation cases) in isolation. Quantile regressions analogous to those in sections 1.7.1 and 1.7.2 yielded qualitatively similar results to those in tables 1.6 and 1.7, with settlements and dropped cases performing marginally better than the full tort sample in the lower quantiles, but not as well in the upper quantiles.

1.10.2. External Validity

The analysis in sections 1.7 and 1.8 has focused on the compensation of workers *conditional* on being injured. However, another arguably important feature of the tort system is its potential to deter negligent behavior. While experience rating to some extent adjusts the premiums paid by employers who subscribe to workers' compensation in accordance with the employers' injury histories, only the largest employers are completely experience rated. Consequently, employers who are insured might not have as much of an incentive to maintain a safe work environment as employers who do not have workers' compensation insurance. However, as noted in section 1.5, Shields and Campbell (2002) find no appreciable differences between the safety measures at subscribing and non-subscribing firms that are otherwise similar—if anything, subscribing firms appear to be safer. This suggests that if workers' compensation did reduce incentives for employers to take appropriate safety precautions, insurance companies may have effectively neutralized this distortion by providing safety consulting and by monitoring the behavior of subscribing employers.

It is worthwhile to consider the implications of this chapter for other areas of tort reform as well. With medical malpractice, for example, one proposal has been to adopt a no-fault compensation scheme that is similar in spirit to workers' compensation, as is done in a handful of other countries, most notably Sweden and New Zealand (Studdert, Mello, and Brennan 2004). The results from this chapter suggest that a no-fault system for medical malpractice could yield a preferable distribution of benefits. At the same time, Localio et al. (1991) find that only two percent of potential medical malpractice claims actually result in legal action; a no-fault system patterned after workers' compensation would almost certainly induce a higher claim rate. As a result, the improvement in the distribution of payments could be accompanied by greater overall expense.

1.11. Conclusion

The evidence presented in this chapter strongly suggests that the benefits paid by workers' compensation insurance compare favorably to the distribution of awards injured workers receive in the tort system. Although the mean award is higher in the tort system than in workers' compensation, the difference is heavily concentrated in the upper tail of the distribution. An injured worker would need to be nearly risk neutral to prefer the distribution of awards in the tort system to that for a comparable set of injuries under workers' compensation. In terms of accomplishing its original goal of substituting for the tort system, current benefit levels under workers' compensation appear to be adequately compensating injured employees.

The finding that injured workers are better off in a workers' compensation system is in spite of the fact that more is actually paid on average to injured workers in the tort system. Furthermore, the available evidence suggests that workers' compensation more efficiently compensates victims than the tort system (i.e., the cost-benefit ratio is lower under workers' compensation). Tillinghast-Towers Perrin (2003) reports that for every dollar spent in the tort system—factoring in administrative expenses and legal fees for plaintiff and defendant attorneys—only 46 cents is actually paid to victims. In contrast, the Congressional Budget Office (2003) finds that 80 percent of total workers' compensation costs go to benefits for injured workers. However, compared to the tort system, pursuing a workers' compensation claim is quick and relatively easy, and a higher claim rate could conceivably mean that the workers' compensation system is more costly than the tort system, even if it is more efficient.

1.12. References

Aït-Sahalia, Yacine and Andrew W. Lo. 2000. "Nonparametric Risk Management and Implied Risk Aversion," *Journal of Econometrics*, 94, pp. 9-51.

Amemiya, Takeshi. 1981. "Qualitative Response Models: A Survey," *Journal of Economic Literature*, 19(4), pp. 1483-1536.

Arrow, Kenneth J. 1976. *Essays in the Theory of Risk Bearing* (North-Holland Publishing Co.: Amsterdam).

Biddle, Jeff. 2001. "Do High Claim-Denial Rates Discourage Claiming? Evidence From Workers' Compensation Insurance," *Journal of Risk and Insurance*, 68(4), pp. 631-658.

Bliss, Robert R. and Nikolaos Panigirtzoglou. 2004. "Option-Implied Risk Aversion Estimates," *Journal of Finance*, 59(1), pp. 407-446.

Butler, Richard J. and John D. Worrall. 1991. "Claims Reporting and Risk Bearing Moral Hazard in Workers' Compensation," *Journal of Risk and Insurance*, 58(2), pp. 191-204.

Congressional Budget Office. 2003. *The Economics of U.S. Tort Liability: A Primer* (The Congress of the United States: Washington, DC).

Dehejia, Rajeev H. and Sadek Wahba. 1999. "Causal Effects in Nonexperimental Studies: Reevaluating the Evaluation of Training Programs," *Journal of the American Statistical Association*, 94(448), pp. 1053-1062.

Dehejia, Rajeev H. and Sadek Wahba. 2002. "Propensity Score Matching Methods for Non-Experimental Causal Studies," *Review of Economics and Statistics*, 84(1), pp. 151-161.

Department of Labor. 2004. *State Workers' Compensation Laws.*

Eisenberg, Theodore and Henry S. Farber. 1997. "The Litigious Plaintiff Hypothesis: Case Selection and Resolution," *RAND Journal of Economics*, 28, pp. S92-S112.

Fishback, Price V. and Shawn E. Kantor. 2000. *A Prelude to the Welfare State* (University of Chicago Press: Chicago, IL).

Gomez-Limon, Jose A., Manuel Arriaza, and Laura Riesgo. 2003. "An MCDM Analysis of Agricultural Risk Aversion," *European Journal of Operational Research*, 151, pp. 569-585.

Hardberger, Phillip D. 2000. "Texas Workers' Compensation: A Ten-Year Survey—Strengths, Weaknesses, and Recommendations," *St. Mary's Law Journal*, 32.

Heckman, James J., Hidehiko Ichimura, and Petra Todd. 1998. "Matching as an Econometric Evaluation Estimator," *Review of Economic Studies*, 65, pp. 261-294.

Hirano, Keisuke and Guido W. Imbens. 2004. "Estimation of Causal Effects using Propensity Score Weighting: An Application to Data on Right Heart Catheterization," *Health Services and Outcome Research Methodology*, forthcoming.

Kling, Jeffrey R. 1995. "High Performance Work Systems and Firm Performance," *Monthly Labor Review*, 118(5), pp. 29-36.

Koenker, Roger and Kevin F. Hallock. 2001. "Quantile Regression," *Journal of Economic Perspectives*, 15(4), pp. 143-156.

Krueger, Alan B. and John F. Burton, Jr. 1990. "The Employers' Costs of Workers' Compensation Insurance: Magnitudes, Determinants, and Public Policy," *Review of Economics and Statistics*, 72(2), pp. 228-240.

Little, Joseph, Thomas A. Eaton, and Gary R. Smith. 1999. *Workers' Compensation: Cases and Materials* (West Group: St. Paul, MN).

Localio, A. Russell, Ann G. Lawthers, Troyen A. Brennan, Nan M. Laird, Liesi E. Herbert, Lynn M. Peterson, Joseph P. Newhouse, Paul C. Weiler, and Howard H. Hiatt. 1991. "Relation between Malpractice Claims and Adverse Events due to Negligence," *New England Journal of Medicine*, 325(4), pp. 245-251.

Madrian, Brigitte C. and Dennis F. Shea. 2001. "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior," *Quarterly Journal of Economics*, 116(4), pp. 1149-1187.

Parker, Jonathan A. 2001. "The Consumption Risk of the Stock Market," *Brookings Papers on Economic Activity*, 2, pp. 279-348.

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

Robertson, David W. 1993. "The Texas Employer's Liability in Tort for Injuries to an Employee Occurring in the Course of Employment," *St. Mary's Law Journal*, 24.

Rosenbaum, Paul. 1995. *Observational Studies* (Springer Verlag: New York, NY).

Rosenbaum, Paul R. and Donald B. Rubin. 1983. "The Central Role of the Propensity Score in Observational Studies for Causal Effects," *Biometrika*, 70(1), pp. 41-55.

Shields, Joseph and D.C. Campbell. 2002. *A Study of Nonsubscription to the Texas Workers' Compensation System: 2001 Estimates* (Research and Oversight Council on Workers' Compensation: Austin, TX).

Sorrels, Randall O. and Jason B. Ostrom. 2000. "Should an Employee's Negligence be Submitted in a Non-Subscriber Case?" *Texas Bar Journal*, 63.

Studdert, David M., Michelle M. Mello, and Troyen A. Brennan. 2004. "Medical Malpractice," *New England Journal of Medicine*, 350(3), pp. 283-292.

Thomason, Terry, Timothy P. Schmidle, and John F. Burton, Jr. 2001. *Workers' Compensation: Benefits, Costs, and Safety under Alternative Arrangements* (Upjohn Institute: Kalamazoo, MI).

Tillinghast-Towers Perrin. 2003. *U.S. Tort Costs: 2003 Update.*

U.S. Census Bureau. 2003. *Statistical Abstract of the United States* (Government Printing Office: Washington, DC).

Appendix 1.A. Construction of Weights

1.A.1. Missing Settlements

Sixteen cases (all settlements) did not have sufficient information about the amount of the award, but were otherwise generally complete. Simply dropping these settlements would understate the true number of settlements and, consequently, the number of tort cases with positive awards. To account for these settlements, I weight the remaining settlements to yield a weighted distribution of outcomes (i.e., verdicts, settlements, and dropped cases) that is the same as if the 16 missing settlements were included.

Let N^{settle} denote the number of cases that were settled and have complete award information. Weights to account for missing settlement amounts are then constructed as follows:

$$\omega_i^{settle} = \begin{cases} 1, & \text{if the case was dropped or resulted in a verdict} \\ (16 + N^{settle}) / N^{settle}, & \text{if the case was settled} \end{cases}$$

1.A.2. Missing Benefits

For tort cases where the injured worker received lost wage benefits but the amount was unknown, the observation is excluded from the analysis when the outcome variable includes benefits. This creates a problem similar to that discussed above for missing settlements: simply dropping observations with missing benefit amounts will understate the share of tort cases in which workers received lost wage benefits from their non-subscribing employers. I compute weights that are structurally similar to those for missing settlements. Let $N^{benefits}$ denote the number of workers who received lost wage benefits (including those missing the benefit amount) and N^{miss} denote the number of those workers who had missing benefit amounts. Weights to account for missing benefit amounts are then constructed as follows:

$$\omega_i^{benefits} = \begin{cases} N^{benefits} / (N^{benefits} - N^{miss}), & \text{if the worker received benefits in addition to the award} \\ 1, & \text{otherwise} \end{cases}$$

1.A.3. Industrial Composition

The final set of weights I employ attempts to adjust for the industrial composition of the sample by weighting cases such that the distribution of industries in my sample matches a hypothesized distribution of industries in the population of non-subscriber injuries. First, I construct the hypothesized distribution of non-subscriber industries. Let p_I denote the non-subscriber rate for industry *I* as reported in the second column of table 1.3, and let $Inj_I^{S,Pop}$ and $Inj_I^{NS,Pop}$ denote the number of injuries in the subscriber and non-subscriber populations, respectively. Assuming that the injury rate for firms that subscribe to workers' compensation is the same as that of non-subscribing firms within each industry, we have the following relationship:

$$\frac{Inj_I^{S,Pop}}{(1-p_I)} = \frac{Inj_I^{NS,Pop}}{p_I}.$$

Next, I use the workers' compensation claims data from TDI to estimate $Inj_I^{S,Pop}$ for each *I,* leaving $Inj_I^{NS,Pop}$ as the only unknown, which can be easily solved for. I use $Inj_I^{NS,Pop}$ as a measure of the population of (relatively severe) injuries at non-subscribers in each industry *I* and weight the tort cases in my sample such that the weighted distribution of industries matches this hypothesized distribution:

$$\omega_i^{industry} = \frac{Inj_I^{NS,Pop}}{Inj_I^{NS,Sample}}$$

where $Inj_I^{NS,Sample}$ is the total number of cases in industry *I* that are observed in my sample. Note that while the above discussion assumes that the injury rates are similar for subscribers and non-subscribers within a given industry, the weighting procedure is valid even if they are merely proportional.

1.A.4. Attorney Caseload

Table 1.A.1 tabulates the caseloads of attorneys in my sample. These are taken from responses to the question, "How many completed non-subscriber cases have you been involved with in the last five years?" Well over half have handled five or fewer cases, but 27 have handled 11 or more. Cases reported by attorneys who have handled more non-subscriber cases may be more representative of the population of non-subscriber cases than the tort cases reported by attorneys who have handled only a few such cases. It may also be

the case that attorneys with more non-subscriber cases are handling shorter, simpler cases. I construct weights for attorney caseload that assign more weight to the cases reported by attorneys with heavy caseloads:

$$\omega_i^{cases} = \frac{Cases_i^{5\ years}}{Cases_i^{sample}}, \text{ where } Cases_i^{sample} \in \{1,2\} \text{ and } Cases_i^{5\ years} \in [0,25]$$

Intuitively, this scheme weights cases such that the one or two cases reported by a given attorney will have a combined weight equal to her five-year caseload. For example, if only one of the cases reported by an attorney with a caseload of 12 is included in the analysis—i.e., if the second case were excluded because it involved multiple injuries, etc.—then that single case would receive a weight of 12. If both of the cases reported by this attorney were included, then each case would be given a weight of 6, such that the total weight of the cases equals the caseload (12).

A drawback to this weighting scheme is that the wide distribution in attorney caseloads can give excessive weight to just a few cases. The seven attorneys with outlying caseloads of 40 or more skew the results severely. As a disproportionate number of these attorneys report cases with little or no award, they generally skew the estimated award for tort cases downward. The bootstrapped standard errors are also grossly inflated, as the estimates from each bootstrapped sample vary greatly depending on which of the heavy caseload attorneys is included. Hence, in the specifications that account for attorney caseload, I exclude the seven attorneys with reported caseloads of 40 or more.

1.A.5. General Weighting Procedure

In many specifications, multiple weights are employed. In such cases, the weight for each observation is computed as the product of the relevant weights. For example, the sample weights for specifications that account for both missing settlements and industrial composition equal $\omega_i^{settle} \times \omega_i^{industry}$, where ω_i^{settle} and $\omega_i^{industry}$ are defined in subsections 1.A.1 and 1.A.3, respectively. Workers' compensation cases are assigned the same weights as their matched tort cases.

Appendix 1.B. Conversion of Awards into Real (2003) Dollars

Awards were converted into constant 2003 dollars using the annual Consumer Price Index from the Bureau of Labor Statistics. The CPI is used for the year of the verdict or settlement in tort cases and for the first year of payment in workers' compensation cases. Five tort cases have positive award amounts but no reported year when the case was completed. One of these cases had the year of injury reported, and for this case I used the CPI for two years later to convert the award into 2003 dollars (as most tort cases were completed two or more years after the injury). In the four other cases, neither the year the case was completed nor the year of injury was reported. Because most of the cases in my sample were completed in 2003 or 2004, I assumed that these four cases were completed in 2003 when I converted the awards into real dollars. Using earlier years did not qualitatively affect the results, however.

Appendix 1.C. Geographic Distribution of Respondents and Non-Respondents

Table 1.C.1 breaks down the respondents and non-respondents in each survey group by city. For the Large Group, the distribution of responding and non-responding attorneys is remarkably similar across cities. In the Small Group, there are several disparities in the distributions of respondents and non-respondents. In particular, a higher share of respondents are located in Austin, East Texas, and Fort Worth, whereas a lower share of

respondents come from Galveston and most notably, Houston. It is not clear what effect (if any) this differing distribution of cities has on my estimates, but given the relatively small differences, it is likely to be minimal. It is also encouraging that if I were to isolate the Large Group, which does not have observable differences in response rates across regions, this would strengthen my results (as I discuss in section 1.10).

Appendix 1.D. Alternative Specifications

1.D.1. Accounting for Attorney Caseload

The cases reported by attorneys who have handled many non-subscriber lawsuits may be more representative of the population of non-subscriber cases than cases reported by attorneys who have dealt with only one or two such cases. With this in mind, I construct caseload weights that are based on a survey question that asks the attorney how many non-subscriber cases she has handled in the past five years. Roughly speaking, five times as much weight is given to a case reported by an attorney who has handled five cases in the past five years compared to a case reported by an attorney who has only handled one case in the past five years. (These weights are discussed more in appendix 1.A.) Attorneys whose most recent non-subscriber cases were more than five years ago are not included in the estimates. I also excluded seven attorneys who reported abnormally heavy caseloads (40 or more). In addition to concerns about their veracity, including them severely skewed the estimates and substantially inflated the standard errors.[30] In tables 1.D.1 and 1.D.2, I replicate the procedures used for tables 1.6 and 1.7, but give greater weight to cases reported by attorneys with heavier caseloads. Accounting for attorney caseload does not qualitatively change the

[30] Nearly all of the cases reported by these seven attorneys had net awards of zero, and if excluding them induces any bias, it is in favor of the tort system.

results. The point estimates in tables 1.D.1 and 1.D.2 reveal an even larger disparity between tort awards and workers' compensation benefits across most of the distribution. However, the bootstrapped standard errors are notably larger, especially in table 1.D.2, where only the difference in medians is statistically significant at the 5 percent level in both panels.

1.D.2. Including Additional Covariates

The results reported in section 1.7 focus on OLS and quantile regressions with no covariates beyond the indicator variable for whether the case is from the tort system or workers' compensation. Although propensity score matching should yield an unbiased estimate of the difference between the awards in each system at each quantile, controlling for additional covariates can potentially increase the precision of the estimates. Hirano and Imbens (2004) propose selecting the covariates for inclusion in the regressions by iteratively regressing the outcome variable on the treatment indicator and each covariate, one at a time. Covariates with *t*-statistics from such regressions that are greater than a specified threshold are then included in the set of controls when the treatment effect is estimated. The intuition behind this procedure is to include the covariates with substantial correlation with the outcome measure, conditional on whether the observation was treated. I chose a threshold *t*-statistic of two, which yielded nine covariates—three types of injuries, three causes of injury, and three body parts injured—that had sufficient influence on the awards, conditional on whether the case was in the tort system. However, this procedure did not provide more precise estimates of the difference between payments under the two systems. The point estimates were qualitatively similar, with the most notable difference being that for the specifications in which the measure of tort awards does not include alternative employer benefits, the estimated differences were somewhat dampened.

Table 1.1. Survey Responses

	Small Group	*Large Group*
Online Response	24	265
Paper Response	61	3
Wrote/Phoned: No Such Case	1	163
Total Responding	86	431
No Response	396	3887
Response Rate (All)	18%	10%
Response Rate (Those with Cases)	18	18

Notes: In addition, 12 letters from the Small Group and 23 from the Large Group were returned for having incorrect or outdated address information. Of the 265 online respondents in the Large Group, 195 reported that they had handled non-subscriber cases.

Table 1.2. Tort Case Outcomes

	Count	*Share*
Settlement	113	59.2%
Verdict-Plaintiff (Worker)	27	14.1
Verdict-Defendant (Employer)	35	18.3
Dropped	16	8.4
Total	191	100.0

Notes: Settlement total includes 16 with missing award amounts.

Table 1.3. Non-Subscription Rates by Industry

	Percent of Firms	*Percent of Workers*
Agriculture	35%	14%
Mining	12	2
Construction	29	10
Manufacturing	36	14
Transportation, Utilities	29	20
Wholesale Trade	25	15
Retail Trade	48	29
Finance, Real Estate	28	8
Services	38	21
Overall	35	16

Notes: From Shields and Campbell (2002).

Table 1.4.a. Distribution of Injury Types

	WC (All) *N*=148,807	*Tort* *N*=175	*WC (Matched)* *N*=525
Amputation	1.3%	5.1%	5.5%
Burn	0.8	1.1	1.0
Contusion	9.4	0.6	0.0
Crushing	1.5	2.9	2.5
Dislocation	1.2	0.6	0.8
Electric Shock	0.1	1.1	1.5
Fracture	10.7	18.3	16.0
Hernia	3.3	29.1	28.4
Inflammation	1.3	0.6	1.5
Laceration	5.6	3.4	4.2
Puncture	0.9	0.6	1.7
Rupture	3.1	4.0	3.0
Severance	0.3	2.3	2.9
Sprain	11.6	19.4	17.9
Strain	44.6	8.6	10.3
Vision Loss	0.0	0.6	0.8
Loss of Hearing	0.0	0.6	0.8
Carpal Tunnel	1.3	1.1	1.3
Other Injury Types	3.0	0.0	0.0

Table 1.4.b. Distribution of Body Parts Injured

	WC (All) N=148,807	*Tort* N=175	*WC (Matched)* N=525
Skull	0.6%	1.1%	1.0%
Ear	0.1	0.6	0.4
Eye	0.3	1.1	1.1
Head-Soft Tissue	0.4	0.6	0.6
Neck-Vertebrae	0.3	0.6	1.0
Neck-Disc	0.5	4.6	4.6
Lower Arm	2.2	4.0	5.1
Wrist	5.4	1.7	1.3
Hand	4.5	5.1	5.1
Finger	6.5	5.7	6.9
Shoulder	1.9	4.0	4.0
Wrist and Hand	0.3	1.1	1.0
Lower Back	31.7	17.7	19.2
Back-Disc	1.0	33.1	33.7
Hip	1.0	0.6	0.4
Upper Leg	0.7	1.1	1.0
Knee	12.3	6.3	5.3
Lower Leg	2.3	4.0	2.9
Ankle	3.9	1.1	1.0
Foot	3.2	4.0	3.6
Abdomen	1.0	1.7	1.0
Other Body Parts	19.9	0.0	0.0

Table 1.5. Percent of Tort Cases Receiving Employer Benefits

	Lost Wage Benefits	*Medical Benefits*	*Either Kind of Benefits*	*Both Kinds*
Yes	38.9%	53.7%	54.3%	36.6%
No	57.7	44.6	41.7	59.4
Missing	3.4	1.7	4.0	4.0
Total	100.0	100.0	100.0	100.0

Notes: *N*=175.

Table 1.6. Difference Between Net Non-Medical Award (Tort) and Indemnity Benefits (Workers' Comp.)

			Quantiles								
		Mean	*50th*	*55th*	*60th*	*65th*	*70th*	*75th*	*80th*	*85th*	*90th*
Weights account for missing settlements. *N* =700	*Tort-WC Diff:*	24805	-9626	-10733	-11589	-12413	-11940	-9631	-1159	11916	47721
		(10729)	(2075)	(2337)	(2554)	(3346)	(4509)	(6644)	(9971)	(14799)	(45750)
	WC:	19869	11626	14072	17284	20595	24933	28631	32710	40032	50980
		(1574)	(1229)	(1479)	(1569)	(1782)	(2122)	(2398)	(2734)	(3254)	(4282)
Weights account for missing settlements and industry composition. *N* =700	*Tort-WC Diff:*	17266	-9706	-11408	-11584	-12668	-12847	-9606	1408	3628	18943
		(7995)	(2218)	(2516)	(2704)	(3653)	(5133)	(7514)	(8897)	(11329)	(30255)
	WC:	19876	12406	15408	17850	20682	24999	28606	32466	38872	50585
		(1818)	(1379)	(1683)	(1774)	(2082)	(2400)	(2684)	(3091)	(3687)	(4575)

Notes: Measured in 2003 dollars. Standard errors in parentheses, computed using bootstrapping, clustered by attorney.

Table 1.7. Difference Between Net Non-Medical Award w/Alternative Wage Benefits (Tort) and Indemnity Benefits (Workers' Comp.)

			Quantiles								
		Mean	*50th*	*55th*	*60th*	*65th*	*70th*	*75th*	*80th*	*85th*	*90th*
Weights account for missing benefits and missing settlements. N=648	*Tort-WC Diff:*	25671	-6238	-6655	-5130	-5765	-3558	3391	2972	16137	63786
		(8079)	(2348)	(2673)	(2770)	(3470)	(5832)	(6945)	(8360)	(17860)	(45914)
	WC:	19119	11588	14557	16403	19158	22558	27609	32028	38863	48867
		(1604)	(1194)	(1467)	(1491)	(1831)	(2175)	(2501)	(2679)	(3338)	(4336)
Weights account for missing benefits, missing settlements, and industry composition. N=648	*Tort-WC Diff:*	20543	-6086	-6655	-5112	-5401	-4233	3198	2016	14246	48029
		(7131)	(2303)	(2838)	(3068)	(3848)	(6204)	(7612)	(7907)	(14496)	(36057)
	WC:	19097	11586	14557	16385	18697	22644	27802	32650	38872	48867
		(1808)	(1312)	(1560)	(1626)	(2018)	(2340)	(2782)	(3038)	(3573)	(4748)

Notes: Measured in 2003 dollars. Standard errors in parentheses, computed using bootstrapping, clustered by attorney.

Table 1.8. Distribution of Awards, Including Employer Benefits

	Tort Cases	*WC Cases*
$0 = *Award*	38.0%	5.0%
$0 < *Award* < $10,001	19.9	39.0
$10,000 < *Award* < $20,001	13.1	21.7
$20,000 < *Award* < $30,001	3.4	11.3
$30,000 < *Award* < $40,001	6.6	9.0
$40,000 < *Award* < $50,001	1.9	4.4
$50,000 < *Award* < $60,001	3.7	3.2
$60,000 < *Award* < $70,001	1.1	3.4
$70,000 < *Award*	12.4	3.1

Notes: Awards for tort cases include employer benefits. Weighted to account for missing benefit amounts.

Table 1.9. Months Between Injury and Award

	Percent
6 months or fewer	4%
7 to 12 months	11
13 to 18 months	15
19 to 24 months	19
25 to 30 months	15
31 to 36 months	17
37 to 42 months	8
43 to 48 months	6
49 or more months	5
Total	100

Notes: Sample limited to tort cases with positive award. Cases with missing month or year of injury, or month or year of award, are excluded. Weighted to account for missing settlements. *N*=73.

Figure 1.1. Estimated Propensity Score on Common Support
(Excludes 80,867 Workers' Comp. Cases)

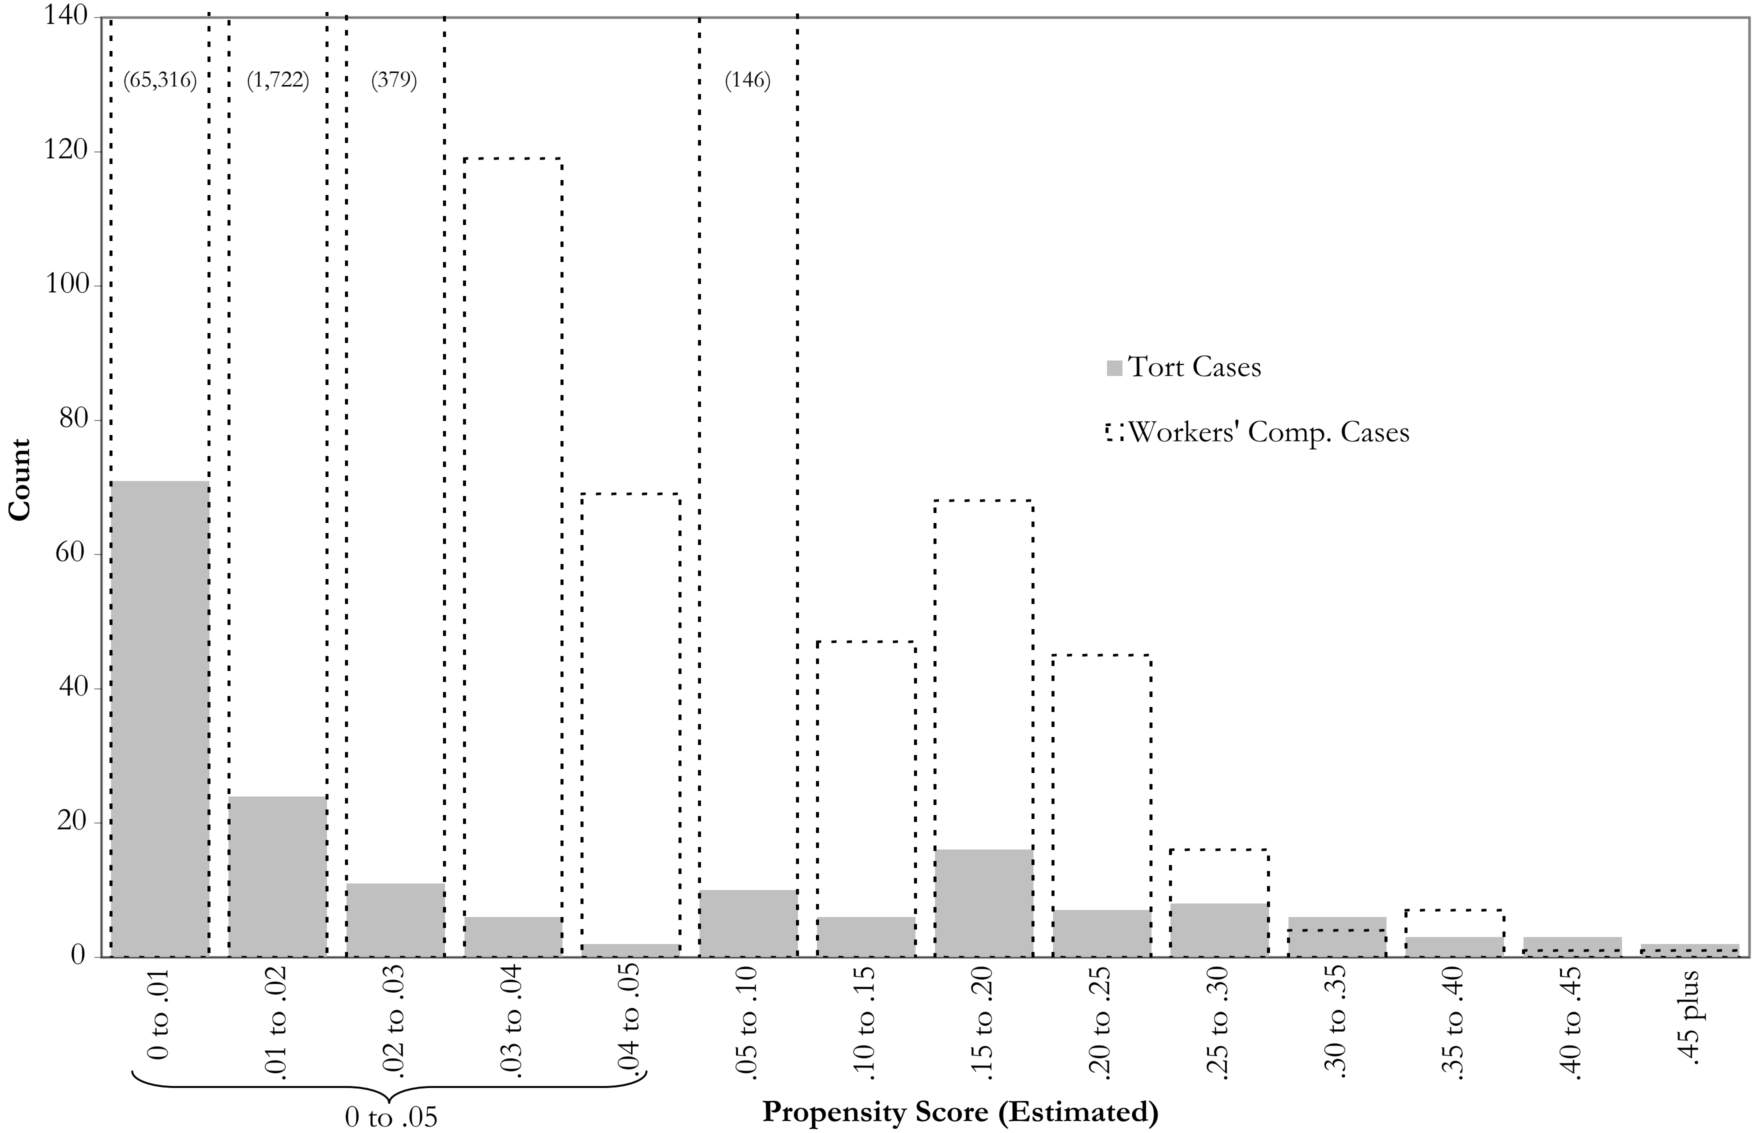

Figure 1.2. Density of Non-Medical Awards Including Wage Benefits, Excluding Zeros (Tort Cases and Matched Workers' Comp. Cases, Weighted to Account for Missing Benefits, Settlements)

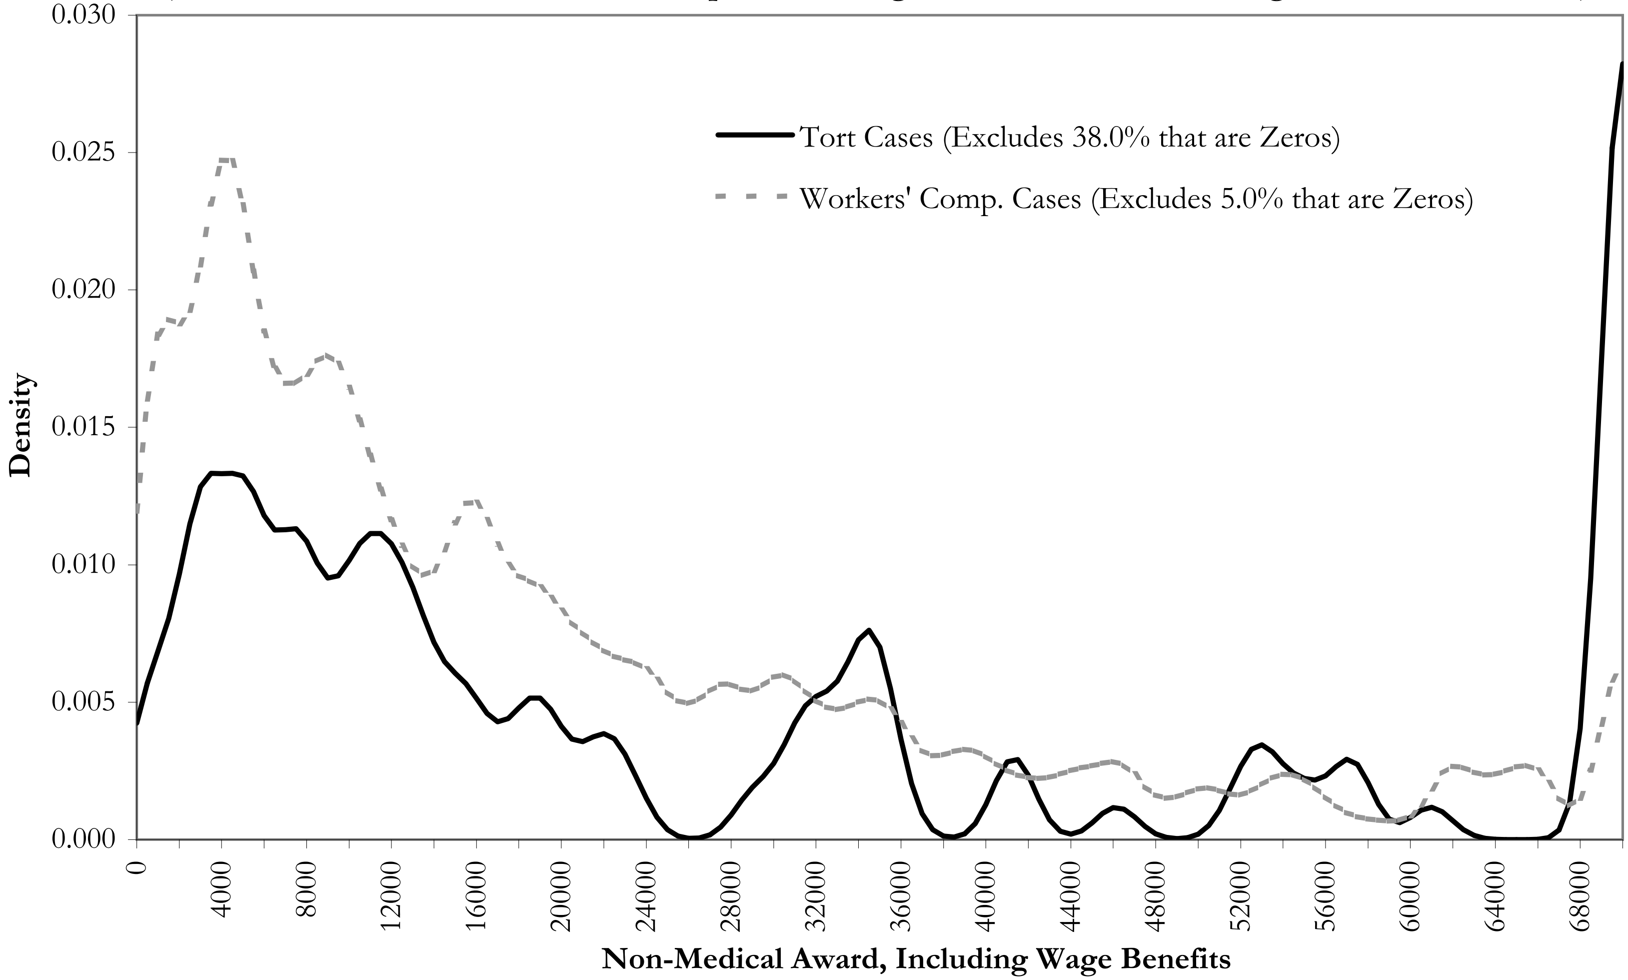

Figure 1.3. Cumulative Distribution of Non-Medical Awards Including Wage Benefits
(Tort Cases and Matched Workers' Comp. Cases, Weighted to Account for Missing Benefits, Settlements)

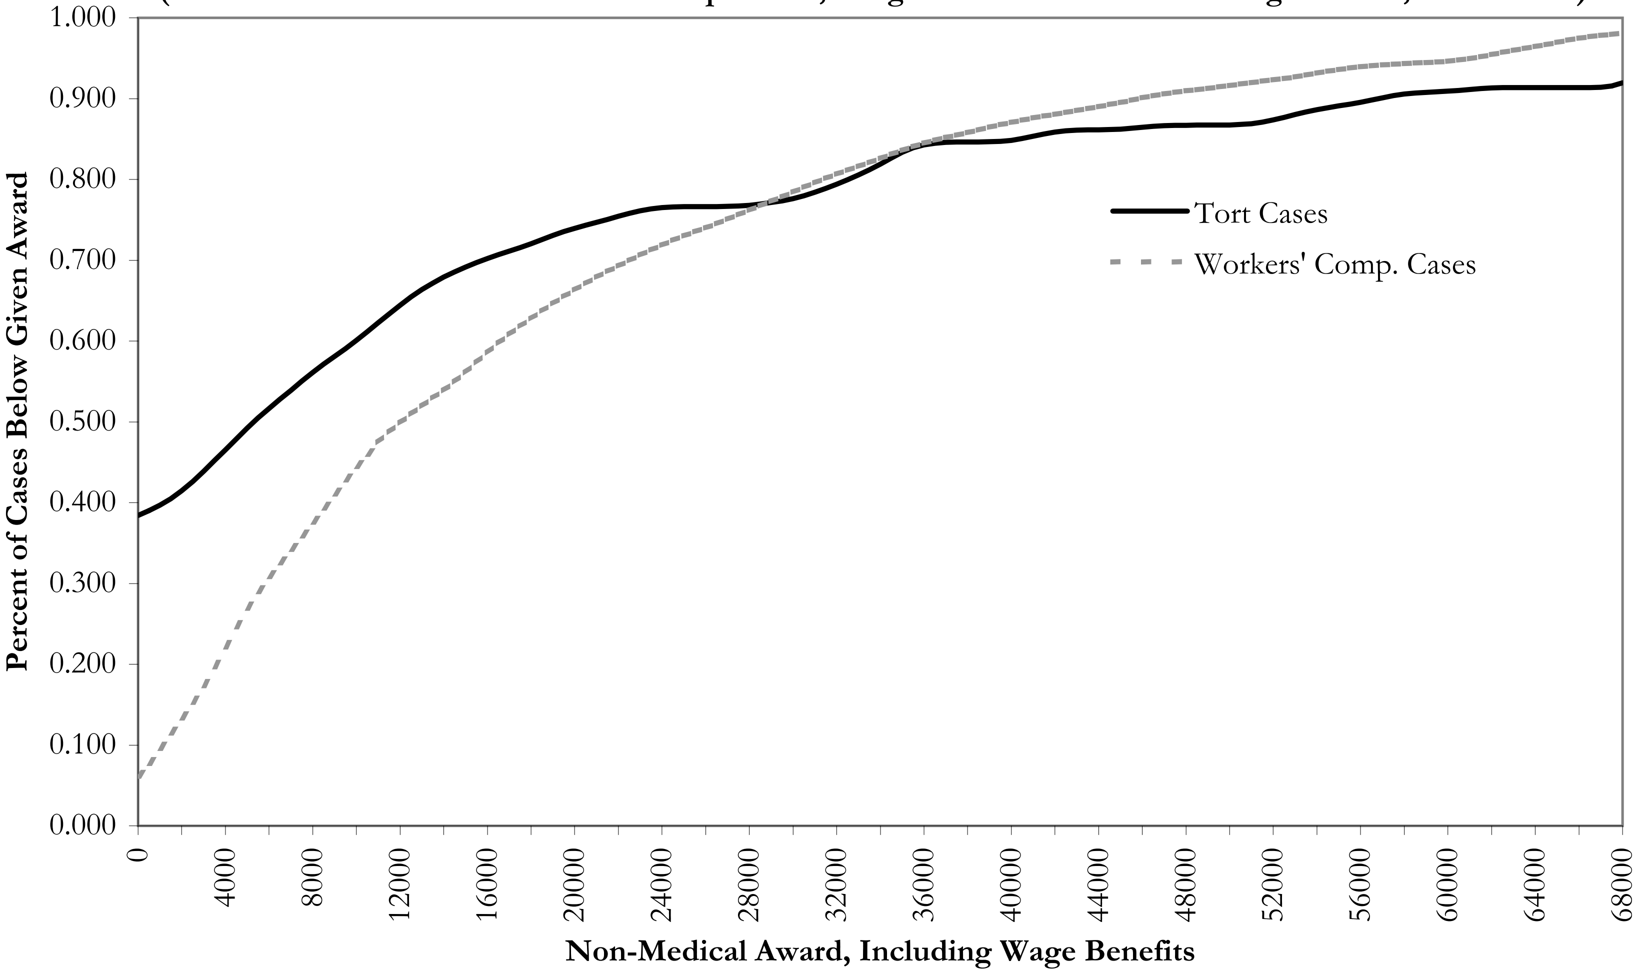

Table 1.A.1. Attorney Caseload

Number of Non-Subscriber Cases	*Number of Attorneys*
0 cases	4
1 to 5 cases	76
6 to 10 cases	23
11 to 20 cases	10
25 cases	10
40 cases	1
50 cases	2
130, 150, 200, 250 cases	1 each
Total	130

Table 1.C.1. Geographic Distribution of Respondents and Non-Respondents

	Small Group		***Large Group***	
	Percent of Non-Respondents	*Percent of Respondents*	*Percent of Non-Respondents*	*Percent of Respondents*
Amarillo/Lubbock	2%	3%	2%	3%
Austin	4	10	11	11
Corpus Christi	0	2	2	3
Dallas	14	13	23	21
East Texas	5	9	4	5
El Paso	4	3	3	6
Fort Worth	6	10	7	7
Galveston	7	2	4	5
Houston	49	38	28	25
Laredo	2	5	2	1
Midland/Odessa	2	0	0	0
San Antonio	3	3	8	8
Waco	1	0	2	1
Other States	1	0	3	3
Total	100	100	100	100

Table 1.D.1. Difference between Net Non-Medical Award (Tort Cases) and Indemnity Benefits (Workers' Comp. Cases), Weighted to Account for Attorney Caseload

			Quantiles								
		Mean	*50th*	*55th*	*60th*	*65th*	*70th*	*75th*	*80th*	*85th*	*90th*
Weights account for attorney caseload and missing settlements. *N* =632	*Tort-WC Diff:*	24283	-11196	-10888	-13060	-12304	-14375	-12741	-9913	-2535	19406
		(18955)	(2265)	(2587)	(2871)	(3510)	(4603)	(7173)	(10050)	(20564)	(80926)
	WC:	17983	11878	14161	17060	18570	21759	24893	30205	36900	41551
		(2446)	(1533)	(1813)	(2072)	(2319)	(2795)	(3318)	(3871)	(5281)	(7017)
Weights account for attorney caseload, industry, and missing settlements. *N* =632	*Tort-WC Diff:*	14131	-11953	-12853	-11950	-12764	-15481	-12822	-10119	-3012	11566
		(13009)	(2386)	(2820)	(3139)	(3645)	(5137)	(7333)	(9919)	(11771)	(26696)
	WC:	18354	11953	15930	17450	19031	22640	24974	30411	37377	41551
		(2831)	(1723)	(1991)	(2316)	(2721)	(3193)	(3815)	(4594)	(5923)	(7181)

Notes: Measured in 2003 dollars. Excludes cases for seven attorneys with abnormally heavy caseloads. Standard errors in parentheses, computed using bootstrapping, clustered by attorney.

Table 1.D.2. Difference between Net Non-Medical Award w/Benefits (Tort) and Indemnity Benefits (Workers' Comp.): Weighted to Account for Attorney Caseload

			Quantiles								
		Mean	*50th*	*55th*	*60th*	*65th*	*70th*	*75th*	*80th*	*85th*	*90th*
Weights account for missing benefits, missing settlements, and caseload. N=588	*Tort-WC Diff:*	31324	-7548	-6355	-5191	-5087	-7296	-6311	4138	8595	49092
		(25319)	(3635)	(4032)	(3884)	(4272)	(7330)	(10788)	(16124)	(59747)	(163899)
	WC:	17524	10821	13522	15027	17164	20192	25311	30227	32774	42959
		(3178)	(1806)	(2236)	(2546)	(2908)	(3397)	(3753)	(4564)	(5502)	(6975)
Weights account for missing benefits, missing settlements, caseload, and industry. N=588	*Tort-WC Diff:*	19135	-7833	-6649	-6658	-6045	-11040	-11859	1606	-2125	15704
		(17594)	(3550)	(4011)	(3995)	(4087)	(6373)	(9223)	(12267)	(19306)	(86562)
	WC:	18603	11158	13575	14924	17926	23192	27352	31130	37125	45253
		(3527)	(1912)	(2370)	(2617)	(3072)	(3441)	(4066)	(4707)	(6285)	(7916)

Notes: Measured in 2003 dollars. Excludes cases for seven attorneys with abnormally heavy caseloads. Standard errors in parentheses, computed using bootstrapping, clustered by attorney.

CH. 2. THE CONSUMPTION SMOOTHING BENEFITS OF WORKERS' COMPENSATION

2.1. Introduction

The primary role of social insurance is to smooth consumption in the event of an adverse occurrence, be it unemployment, injury, or poverty. In the case of an on-the-job injury, workers' compensation insurance pays the medical costs associated with the injury, impairment benefits (usually for permanently debilitating injuries), and lost wage benefits. Lost wage benefits are the component that is most responsible for smoothing consumption across temporary injury spells, which account for the high majority of workers' compensation claims. These benefits are calculated as a fixed percentage of the worker's pre-injury average weekly wage. The replacement rate varies by state (and over time), but is frequently equal to two-thirds of the pre-injury wage. However, the existence of the earnings replacement component of workers' compensation creates potential moral hazard, whereby workers may take unnecessarily long to recover from an injury or otherwise prolong their time out of work (Moore and Viscusi 1990). The optimal benefit rate should therefore trade off the behavioral distortions created by the insurance with the consumption smoothing benefits it provides.

There is a substantial literature on the costs and inefficiencies generated by a separate but similar social insurance program, unemployment insurance. (See Krueger and Meyer 2002 for a thorough survey.) Gruber (1997) attempts to estimate the positive effect unemployment insurance has on consumption, an important supplement to the aforementioned literature on the distortions of unemployment insurance. Further work by Chetty (2004) extends Gruber's earlier attempts to use the effect of unemployment insurance on consumption together with existing estimates of the inherent distortions created by unemployment insurance to determine the optimal level of unemployment benefits.

There is comparably little research on workers' compensation insurance, and none that I am aware of on the impact of workers' compensation on smoothing consumption in the event of an on-the-job injury. This paper attempts to help fill this void. Modeling the analysis after Gruber's, I use the Panel Study of Income Dynamics and exploit state and year variation in workers' compensation benefits to estimate how consumption changes are affected by the generosity of benefits. I estimate that a 20 percentage point increase in benefit generosity reduces the drop in consumption by 3.0 to 5.5 percent. Furthermore, my estimates suggest that in the absence of workers' compensation insurance, consumption would fall by at least 11.7 percent after an on-the-job injury, with estimates ranging as high as 21.3 percent.

The next section discusses the data as well as the construction of the replacement rate variable that is central to the subsequent analysis. The third section introduces the econometric model that is estimated in section 2.4. Section 2.5 summarizes the model for optimal unemployment benefit levels that was originally developed by Baily (1978) and applies this model to workers' compensation insurance. The sixth section concludes with a qualitative interpretation of the results and discusses some of the limitations of the analysis.

2.2. Data

2.2.1. Panel Study of Income Dynamics

The raw data come from the Panel Study of Income Dynamics (PSID). I focus on heads of households from 1983 through 1987, and 1991 through 1996. Prior to 1983, not enough workers' compensation information is provided, and in 1988 and 1989 data on consumption were not collected. (Since the data for any given year include lagged values of consumption, the lack of consumption data in 1989 also precludes using the 1990 data.)

Data on consumption for year *t* come from the survey for year *t*; however, data on workers' compensation receipt in year *t* come from the survey for year $t + 1$. Consequently, when the PSID moved to a biennial survey in 1997, the data were insufficient for this analysis.

Ideally, we would like to see how the average weekly benefits an injured worker receives affect her consumption. However, the data that are available in the PSID on the amount of workers' compensation received are unfortunately not consistently coded across years or detailed enough to use in the analysis. Instead, I focus on the benefits that an injured worker is eligible for, which I discuss further in the next subsection.

In contrast to Gruber's approach to unemployment insurance, I am not able to identify everyone who was *eligible* for workers' compensation, only those who actually *received* payments. We would also ideally like to observe consumption for months just prior to, during, and just after the injury spell. However, the PSID is an annual survey, and only asks for consumption as of the month of the interview.[31] Further complicating this issue, there are not enough observations to isolate only those receiving workers' compensation as of the date of the interview in year *t*; instead, I include all household heads who were employed and not receiving workers' compensation as of the interview date in year $t - 1$, but reported in the subsequent year *t* interview that they received workers' compensation in at least one month since the year $t - 1$ interview. I later discuss the possible impact of this inclusion in the empirical analysis.

The PSID data are useful because they contain information on both consumption and workers' compensation receipt. Unfortunately, the latter is limited to household food

---

[31] The PSID consumption questions take the form, "How much do you (and everyone else in your family) spend on food that you use at home in [time period]?" While the reference point is somewhat vague, Zeldes (1989) argues that the question refers to contemporaneous consumption.

consumption—which includes food prepared at home, eaten out, and delivered, as well as food stamps—and food consumption could conceivably respond differently to workers' compensation generosity than other consumption would. In particular, since food is a necessity, we might expect that food responses would be muted compared to responses for other consumption goods. Kingston, Burgess, and Walters (1978) and Burgess, Kingston, St. Louis, and Sloane (1981) find that in the case of unemployment, food consumption and other consumption respond very similarly. Browning and Crossley (2004) find much larger consumption drops for durable goods than nondurables, however, and it is important to bear this caveat in mind when interpreting the results.

2.2.2. Computation of Replacement Rate

I use the U.S. Department of Labor's *State Workers' Compensation Laws* supplemented with information from the *Monthly Labor Review* to create a detailed program that determines the benefit formula for a worker injured in a given state at a particular time. Because workers' compensation pays for all medical costs, I focus on lost wage payments—the primary source of interstate variation over time. I use states' temporary total disability schedules—which are closely related to permanent total disability in most states—to generate the benefits for which an injured worker is eligible.[32] There is some variation in the basic formula, which is usually two-thirds of the worker's pre-injury average weekly wage or eighty percent of her spendable earnings. However, substantial variation in maximums and minimums (especially the former) generates a wide range of replacement rates. While most states base benefits purely on the worker's pre-injury average weekly wage, in some cases

[32] About 70 percent of all claims are for temporary total disability (Krueger and Meyer 2002).

there are additional payments or different schedules based on marital status or dependents, generating further variation.

I compute benefits as of the month the worker most recently received workers' compensation prior to the date of the interview. Because workers' compensation benefits are not subject to federal income taxes, I calculate the replacement rate R_i by dividing the computed weekly benefit by the after-tax average weekly wage:

$$R_i = \frac{B_i}{(1 - m_i)W_i}, \tag{2.1}$$

where B_i is the calculated amount of weekly benefits, W_i is the average weekly wage of worker *i*, and m_i is the worker's marginal tax rate.

For the years 1982 to 1990, the marginal tax rate is included in the PSID. For 1991 and beyond, I use the taxable income of the head and spouse together with information on children and marital status to impute the marginal tax rate, based on the tax brackets, exemption amounts, and deductions for each year according to the Internal Revenue Service and the Tax Policy Center. The PSID also asks whether the individual took the standard deduction or itemized deductions on her tax return. Following the *Statistical Abstract of the United States*, I assume itemizers deduct 10 percent of their taxable income, and for others I calculate the standard deduction using marital status and number of dependents.

To check the accuracy of this imputation method, I replicated the above procedure for 1987 and compared the results to the reported marginal tax rates that year. The correlation of the actual and imputed marginal tax rates is .87. This suggests that, while imperfect, the imputation is at least a close approximation.

The dark bars in figure 2.1 represent the distribution of after-tax replacement rates. While over 40 percent of the sample have replacement rates between 80 and 89 percent, there is nonetheless substantial variation. Over one quarter have replacement rates below 70 percent, and 6 percent receive benefits that actually exceed their after-tax wages.

However, part of the replacement rate variation simply reflects the variation in average weekly wages in the sample. To isolate the variation that is due purely to differences in wages—i.e., the variation that is *not* due to differences in workers' compensation laws across states and time—I compute a generic replacement rate that subjects all individuals to the same benefits schedule. The schedule I use is artificial, but representative of the most common benefits structure: the basic benefit is two-thirds of the worker's pre-injury average weekly wage, subject to a minimum of 15 percent of the state's average weekly wage or the worker's average weekly wage (whichever is less) and a maximum that is 90 percent of the state's average weekly wage. I use the entire sample to compute the "state" average weekly wage that determines the minimum and maximum benefit.

The lighter bars in figure 2.1 show the distribution of replacement rates were everyone subject to this same generic benefits formula. The generic replacement rates are much more heavily concentrated in the 80 to 89 percent range than the distribution of actual replacement rates. This comparison demonstrates that there is substantial variation in effective replacement rates across time and states beyond that generated simply from the distribution of wages. While the generic benefits formula is comparable to the schedule used by many states, it is also clear that it is on average more generous.

2.2.3. Summary Statistics

Table 2.1 reports summary statistics for the main variables in the analysis. All results here and in the subsequent analysis are weighted using the provided family weights. The dependent variable of interest is the change in log consumption. The average is negative, consistent with the notion that consumption falls following an injury spell. It is curious that the consumption drop is very small, however. Still, for those at the twenty-fifth percentile of the distribution, the drop in consumption is substantial. The relatively small consumption drop overall may be caused by the inclusion of injured workers who have recovered fully by the time of the interview, a possibility I address in the empirical analysis. Another possible explanation is that the generosity of workers' compensation almost completely neutralizes any consumption dips following an injury: the average benefit is over three quarters of post-tax earnings. In contrast, unemployment insurance replaces less than half of a worker's after-tax earnings, on average (Krueger and Meyer 2002).

As a whole, injured workers are disproportionately young and male, and tend to have lower education levels. Consistent with this, the average worker in the sample has a high school education, and less than a quarter of the sample is female. The latter fact also partly reflects that the sample only includes household heads, and because of this, it is not surprising that at an average age of 39, the sample is not altogether young. The average weekly wage after taxes is $317 (in 1982-1984 dollars), and the average duration of an injury in the sample is 2.76 months.

2.3. Econometric Model

The primary econometric model focuses on the effect of the replacement rate on reducing the drop in consumption following an on-the-job injury. The base regression specification takes the form:

$$(2.2) \qquad \Delta \log C_i = \mathbf{x}_i'\boldsymbol{\beta}_1 + \mathbf{y}_i'\boldsymbol{\beta}_2 + \beta_3 R_i + \varepsilon_i ,$$

where the dependent variable is the difference in (the natural log of) consumption in the post-injury period and the pre-injury period, and $\mathbf{x}_i$ is a vector of individual characteristics that includes the real after-tax weekly wage, age, sex, marital status, education, number of children, and any changes in household food needs, as well as a constant term.[33] Variables are converted into real dollars whenever relevant using the corresponding CPI item for that month, and all estimates reported also include the year effects captured by $\mathbf{y}_i$. The variable of interest is the calculated after-tax replacement rate R_i described in the previous section. A positive estimate of the coefficient on R_i would indicate that higher benefit levels are associated with smoother consumption profiles.

2.4. Results

Table 2.2 reports the base results. Comparing the effect of the weekly wage when the replacement rate is excluded (column 1) and included (column 2), it is clear that the replacement rate and weekly wage are highly correlated. In contrast to Gruber (1997), however, the weekly wage has no impact on consumption smoothing when it is included alone, but the impact is much larger when the replacement rate is included. The point estimate of the effect of the replacement rate (.148) is reasonably large and positive in column 2, although it is not statistically significant. A 20 percentage point increase in the replacement rate is associated with an approximately 3 percent smaller drop in consumption. Were the replacement rate zero, these estimates suggest that consumption would fall by 11.7 percent, substantially greater than the almost zero percent drop that is observed.

The statistical insignificance of the results is driven at least in part by the small sample size. Even taking that into account, however, the effects I estimate here are smaller than the effect of unemployment benefits on consumption estimated by Gruber (1997).

[33] The household food needs variable is the Orshansky poverty threshold used by the USDA, and is included in the PSID. The variable is computed based on the size of the household.

We could imagine that an on-the-job injury causes an only temporary drop in consumption, but that by the time of the interview, consumption is not affected if the worker has returned to work. This would suggest that the effect of the replacement rate on consumption would be smaller for those who are further removed from their workers' compensation spells. However, including a term that interacts the replacement rate with the number of months since workers' compensation was last received does not affect the main results; the sign of the coefficient on this interaction term in column 3 is actually positive, but small and insignificant. However, given the large standard error, I cannot reject a negative estimate as large in magnitude as -.06, which would be substantial.

In table 2.3, I explore other potential explanations for the somewhat small estimates. In the first column, the regression includes state fixed effects. Including state effects increases the estimated effect of the replacement rate on consumption considerably, indicating that there are important unobserved, state-specific factors. The point estimate is .274—almost twice the effect estimated without state effects—and it is statistically significant. The implied consumption drop in the absence of workers' compensation is also much larger, about 21 percent. It is surprising that the estimates with and without state effects are so different, and this may be a symptom of the small sample size; at any rate, this certainly suggests caution and further investigation. Bearing this in mind, my analysis will report results both with and without state effects for the remainder of the paper.

Still another explanation for the relatively small estimates in table 2.2 is that very short injury spells may not affect consumption, while more persistent injuries may be more likely to lead to observable behavioral impacts. The injuries in my sample led to workers' compensation benefits lasting 2.76 consecutive months on average, with a median of 2.00 months. The high proportion of shorter-duration injuries could mask larger impacts for more severe injuries. Using the duration of the workers' compensation spell as a measure of the severity of the injury, column 2 of table 2.3 considers whether the impact of the replacement rate is affected by how long the worker is unable to work. Combining the main effect of $Duration_i$ with the interaction of $Duration_i$ and the replacement rate R_i (evaluated at the mean replacement rate of 76.8 percent) reveals that workers with longer spells experience greater falls in consumption, as expected. Each additional month of injury is associated with a drop in consumption of about 1.3 percent. However, the replacement rate actually has a much *smaller* effect on consumption for more prolonged injuries. This difference is also highly statistically significant. The next column reports similar results when state effects are included; once again, the effect of the replacement rate on consumption is much larger when I account for state effects.[34]

2.5. Optimal Benefit Level

Beginning with Baily (1978), several studies have examined the optimal level of unemployment benefits, including Gruber (1997) and Chetty (2004). Baily creates an additive, two-period model where the worker consumes and saves in the first period, and then faces an exogenous probability of being unemployed in the second period. In the second period, the worker decides the duration of her unemployment based on the level of unemployment benefits. Finally, aggregate benefits are financed by a tax on earnings.

It stands to reason that the Baily model can be easily applied to workers' compensation. With an injury, there is an exogenous time out of work while the worker

[34] Only counting consecutive months avoids mistaking two separate minor injuries for one major injury. However, it may also understate the injury severity for recurring or chronic injuries. Conducting similar analysis but with total months receiving workers' compensation instead of consecutive months yielded qualitatively similar results.

recuperates, but the worker may prolong this recuperation period, either by exaggerating the severity of the injury, or indirectly by not rehabilitating as rigorously as possible. In such cases, there will also be an endogenously-determined injury spell as more generous benefits lead injured workers to extend lost work time. Translating the findings from Baily's model to workers' compensation yields the following relationship:

(2.3) $$-\frac{\Delta C}{C_{t-1}}\gamma = \eta,$$

where $\Delta C = C_t - C_{t-1}$ is the consumption in the post-injury period t less consumption in the pre-injury period $t-1$; γ is the coefficient of relative risk aversion; and η is the elasticity of the duration of workers' compensation receipt with respect to benefits.

Equation 2.3 can then be re-written as:

(2.4) $$\Delta \log C = \log\left(\frac{C_t}{C_{t-1}}\right) = \log\left(-\frac{\eta}{\gamma}+1\right) \approx -\frac{\eta}{\gamma}.$$

The basic regression result from column 2 of table 2.2 implies that $\Delta \log C = -.117 + .148 \times R$ (all else equal), where R is the replacement rate. Combining these two expressions results in the optimal replacement rate as a function of the other parameters:

(2.5) $$R^* = .79 - 6.76 \times \frac{\eta}{\gamma}.$$

Krueger and Meyer (2002) survey studies that have attempted to estimate the duration elasticity η for workers' compensation. The estimates in most of these studies are generally in the range of .20 to .40 when all injuries are included. In the first column of table 2.4, I report the optimal benefit level for a range of assumptions about the coefficient of relative risk aversion γ when the duration elasticity is assumed to be .40. For "normal" levels of risk aversion in the neighborhood of 3.0, the optimal benefit level is at or near zero. Chetty (2004) argues that previous studies implicitly understate the income elasticity for the unemployed, and he demonstrates empirically that the unemployed may have higher levels of risk aversion, in the 4.5 to 5.0 range. Injured workers would arguably exhibit similar attitudes toward risk as the unemployed. Even at these higher levels of risk aversion, the optimal replacement rate is only 25 percent.

The second column of table 2.4 instead uses the results from the regression with state effects, such that the optimal replacement rate is given by:

$$(2.6) \qquad R^* = .78 - 3.65 \times \frac{\eta}{\gamma}.$$

The optimal benefits are much greater with this specification. Even with relative risk aversion of 5.0, however, the optimal replacement rate is less than 50 percent, and well below the average replacement rate in the sample (.768).

As evidenced by the relatively short duration of the injuries in my sample, the typical injury in my sample is relatively minor. Consequently, using the lower end of the range of elasticity estimates may be more appropriate. The third and fourth columns repeat the above exercise, except the duration elasticity with respect to benefits is halved to .20. When the behavioral distortions are lower, the optimal benefit level increases. Using the base results when relative risk aversion is assumed to be between 4.5 and 5.0, the optimal replacement rate is near one half, and using the state effects results yields optimal levels just above 60 percent.

2.6. Summary and Discussion

It is worth reiterating some of the caveats of the analysis reported here. The sample size is very small, and as is apparent from the differences between the estimates with and without state effects, the estimates are not comfortably robust to specification changes. Furthermore, even in the (unlikely) case that food consumption is measured accurately, the measures of consumption available in the PSID do not account for non-food consumption, which Browning and Crossley (2004) find is more responsive to benefit generosity than food consumption alone. A separate issue is the inclusion of only workers' compensation recipients, which ignores the fact that the decision to claim benefits conditional on an injury is likely related to benefit generosity. (See, for example, Krueger 1990.) Nonetheless, the results presented above are at least instructive, if crude.

The estimates of the effect of workers' compensation benefit generosity on consumption smoothing are consistently substantial and positive, although the magnitude depends greatly on whether or not state effects are included in the specification. However, comparing the most optimistic parameter estimates that account for state effects and assuming injured workers have very high levels of risk aversion but relatively low duration elasticities leads to an estimate of the optimal benefit level of 63.1 percent, 13.7 percentage points lower than the actual average replacement rate of 76.8. In contrast, Gruber (1997) and Chetty (2004) find that, for comparable levels of risk aversion, actual unemployment insurance benefits are very similar to optimal levels.

The fact that unemployment insurance has replacement rates very close to the optimal rate, while workers' compensation has much higher replacement rates than is optimal may partly reflect differences in the original purposes of the two systems. Workers' compensation evolved as an alternative to, and eventual replacement for, the tort system.

An injured worker who is covered by workers' compensation loses the opportunity to sue her employer for damages following an on-the-job injury, even if the injury was clearly the employer's fault. Fishback and Kantor (2000) argue that while workers clearly benefited from the introduction of workers' compensation insurance in the early twentieth century, employers also benefited by avoiding potentially devastating litigation.

On the other hand, unemployment insurance is social insurance in the purest sense; the sole beneficiary is the worker who is unemployed. Workers' compensation benefits may be higher than is socially optimal because injured workers give up some of their rights in exchange for these benefits.

2.7. References

Baily, Martin Neil. 1978. "Some Aspects of Optimal Unemployment Insurance," *Journal of Public Economics*, 10(3), pp. 379-402.

Browning, Martin and Thomas F. Crossley. 2004. "Shocks, Stocks, and Socks: Smoothing Consumption Over a Temporary Income Loss," McMaster University Working Paper 2003-07.

Bureau of Labor Statistics. 1983-1997. "Workers' Compensation: State Enactments," *Monthly Labor Review*, assorted issues.

Burgess, Paul L., Jerry L. Kingston, Robert D. St. Louis, and Joseph T. Sloane. 1981. "Changes in Spending Patterns Following Unemployment," Department of Labor, Employment and Training Administration Occasional Paper 81-3.

Census Bureau. 2005. *Statistical Abstract of the United States* (U.S. Government Printing Office: Washington, DC).

Chetty, Raj. 2004. "Optimal Unemployment Insurance When Income Effects are Large," University of California-Berkeley, mimeo.

Department of Labor. 1983-1996. *State Workers' Compensation Laws.*

Fishback, Price V. and Shawn E. Kantor. 2000. *A Prelude to the Welfare State* (University of Chicago Press: Chicago, IL).

Gruber, Jonathan. 1997. "The Consumption Smoothing Benefits of Unemployment Insurance," *American Economic Review*, 87(1), pp. 192-205.

Kingston, Jerry L., Paul L. Burgess, and Chris Walters. 1978. "The Adequacy of Unemployment Insurance Benefits: An Analysis of Adjustments Undertaken Through Thirteen and Twenty-five Weeks of Unemployment," Department of Labor, Employment and Training Administration Occasional Paper 78-2.

Krueger, Alan B. 1990. "Incentive Effects of Workers' Compensation Insurance," *Journal of Public Economics*, 41, pp. 73-99.

Krueger, Alan B. and Bruce D. Meyer. 2002. "Labor Supply Effects of Social Insurance," *Journal of Public Economics, Vol. 4*, editors: Alan Auerbach and Martin Feldstein (North-Holland Publishing Co.: Amsterdam), pp. 2327-2392.

Moore, Michael J. and W. Kip Viscusi. 1990. *Compensation Mechanisms for Job Risks* (Princeton University Press: Princeton, NJ).

Zeldes, Stephen P. 1989. "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, 97(2), pp. 305-346.

Table 2.1. Summary Statistics

		Quantiles	
	Mean	*25th*	*75th*
Δ log(*consumption*)	-.003	-.224	.234
After-tax replacement rate	.768	.685	.846
After-tax weekly wage (real)	317	202	397
Duration of WC spell (months)	2.76	1.00	3.00
Age	39.3	--	--
Female	.237	--	--
Married	.556	--	--
Years of education	12.2	--	--
Number of children (under 18)	.839	--	--

Notes: Sample size is 621.

Table 2.2. Impact of Replacement Rate on the Change in Log Consumption

	Dependent variable: Δ log(*consumption*)		
	(1)	(2)	(3)
After-tax replacement rate	--	.148	.132
		(.123)	(.195)
Implied consumption fall at replacement rate of 0	--	-.117	-.124
After-tax weekly wage (x 1000)	.015	.149	.161
	(.133)	(.165)	(.163)
Months since WC receipt	--	--	.002
			(.031)
Months x *Replacement rate*	--	--	.006
			(.043)
Age	.003	.003	.003
	(.002)	(.002)	(.002)
Female	.039	.047	.047
	(.087)	(.087)	(.087)
Married	-.072	-.074	-.075
	(.070)	(.070)	(.069)
Years of education	-.006	-.007	-.006
	(.010)	(.010)	(.010)
Change in log food needs	.392	.391	.387
	(.141)	(.139)	(.137)
Number of children	.043	.045	.045
	(.018)	(.017)	(.017)

Notes: Huber-White standard errors in parentheses, clustered by household. All regressions include year effects. $N = 621$.

Table 2.3. Alternative Specifications

	Dependent variable: Δ log*(consumption)*		
	(1)	(2)	(3)
After-tax replacement rate	.274	.301	.425
	(.104)	(.076)	(.060)
Implied consumption fall at replacement rate of 0	-.213	-.062	-.156
Duration of WC spell	--	.052	.051
		(.013)	(.016)
Duration x *Replacement rate*	--	-.082	-.083
		(.016)	(.016)
State effects	Yes	No	Yes

Notes: Huber-White standard errors in parentheses, clustered by household. All regressions include year effects as well as all other covariates included in the second column of table 2.2. $N = 621$.

Table 2.4. Optimal Benefit Levels

	Duration Elasticity = .40		*Duration Elasticity* = .20	
CRRA	(1) Base Result	(2) w/State Effects	(3) Base Result	(4) w/State Effects
1.0	0	0	0	.047
1.5	0	0	0	.291
2.0	0	.047	.115	.412
2.5	0	.193	.250	.485
3.0	0	.291	.340	.534
3.5	.018	.360	.404	.569
4.0	.115	.412	.453	.595
4.5	.190	.453	.490	.615
5.0	.250	.485	.520	.631
10.0	.520	.631	.655	.704

Notes: Based on estimates reported in tables 2.2 and 2.3. CRRA is the coefficient of relative risk aversion.

Figure 2.1. Distribution of Replacement Rates

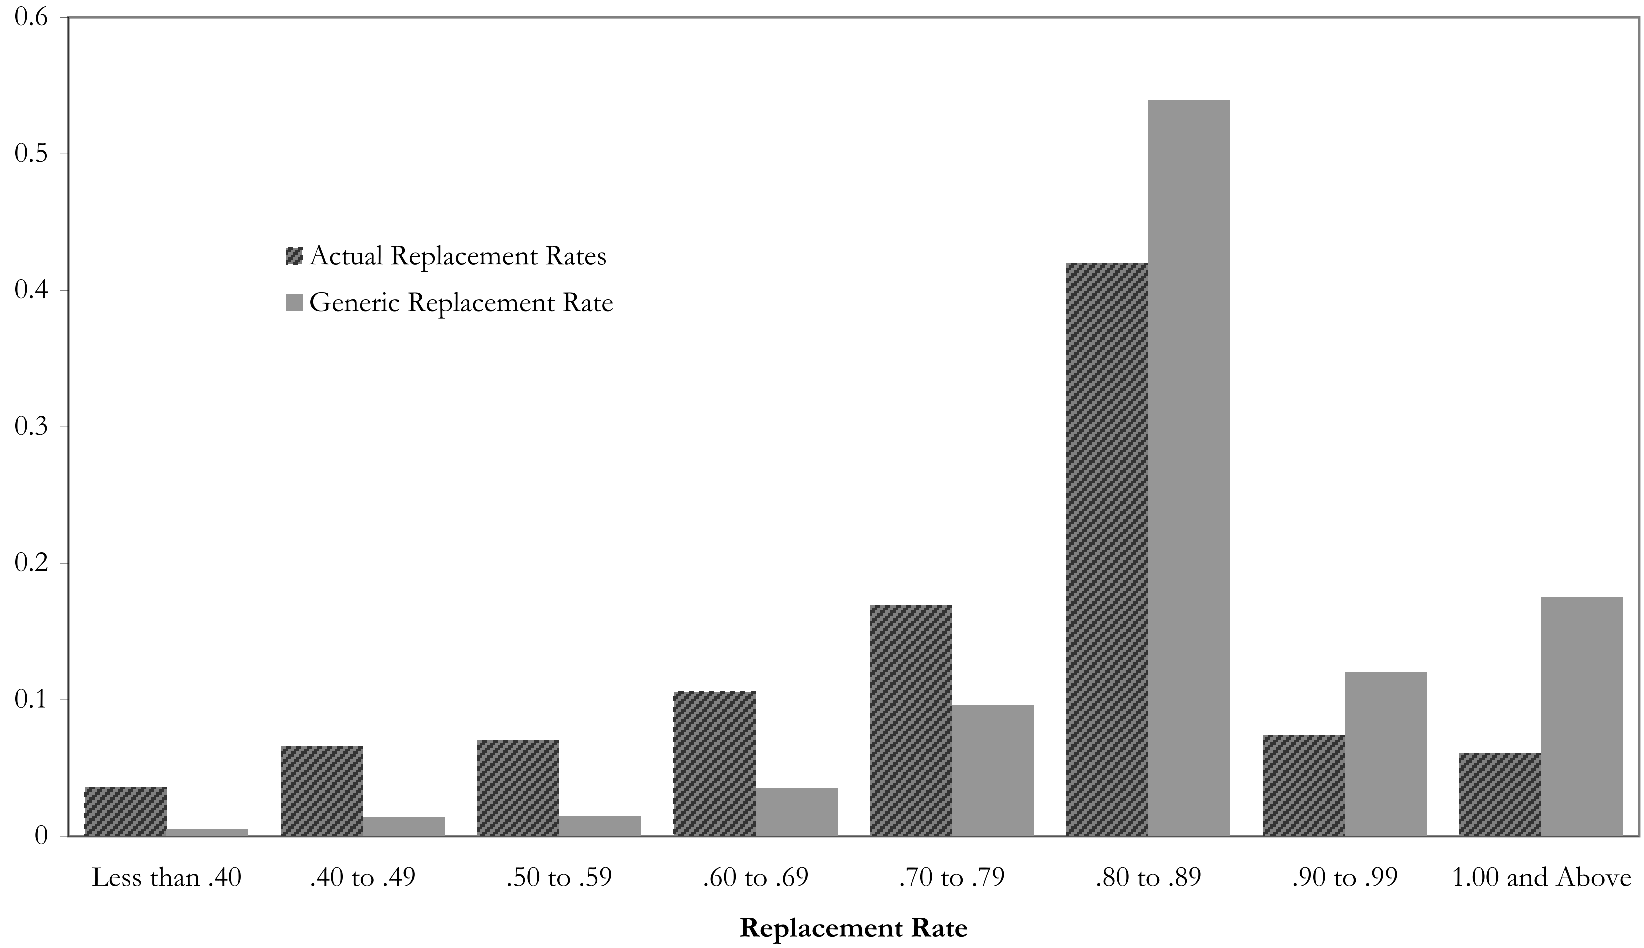

CH. 3. THE DIURNAL PATTERN OF ON-THE-JOB INJURIES

3.1. Introduction

Shortly after 4:00 AM on March 28, 1979, mechanical equipment at the nuclear power plant at Three Mile Island malfunctioned. In the course of responding, operators working the late-night shift made errors that grossly exacerbated the situation, resulting in the worst accident in the short history of U.S. commercial nuclear power (President's Commission on the Accident at Three Mile Island, 1979). Seven years later and halfway around the globe, at 1:23 AM on April 26, 1986, gross safety negligence by night shift workers at the nuclear reactor in Chernobyl led to an even more catastrophic nuclear disaster (OECD Nuclear Energy Agency, 1995).

In the popular press it is often asserted but rarely explained that workplace injuries are more common at night.[35] In the academic literature, economists have largely ignored the diurnal pattern of on-the-job injuries and, by extension, the economic ramifications. In this chapter I use data on workers' compensation claims from Texas to estimate the empirical distribution of injuries over the day. The results show that the injury rate is very high in the off-hours late at night and low during the regular nine-to-five shift.

In addition to providing empirical evidence of this phenomenon, an important contribution of this chapter is the decomposition of the factors causing the observed injury pattern. I explore the possibility that the empirical injury cycle is merely an artifact of compositional changes in the age or industry and occupation of workers throughout the day. Late night workers also work longer hours, and I examine whether fatigue can explain the

[35] See, for example, "In the Deep of the Night," *Time* (November 1, 1999) and "No More Nine-to-Five," *The Economist* (January 10, 1998).

injury pattern by computing the mean shift duration at each hour of the day. However, I reject each of these as the lone explanation of the injury pattern.

Instead, there may simply be inherent physiological implications of late-night work that make off-hours jobs more hazardous. This is an important distinction because it suggests that, when scheduling work hours, firms should consider shift *time* in addition to factors such as shift *length*, which is merely correlated with late night work and contributes to a higher injury rate but is not unique to night work.

Section 3.2 describes the data, and section 3.3 presents my estimates of the distribution of injuries throughout the day. Section 3.4 addresses alternative explanations for the empirical injury distribution, focusing in turn on compositional and physiological explanations. I conclude with section 3.5.

3.2. Data

3.2.1. Texas Workers' Compensation Commission

The data come from two sources. The Texas Workers' Compensation Commission (TWCC) provided data on age and time, date, and nature of injury for workers injured between 1998 and 2002. This is a complete count of all workers who were employed by a firm carrying workers' compensation insurance and were injured in this time period. Unlike nearly all other states, Texas does not mandate that firms provide workers' compensation insurance (United States Department of Labor, 2003), and my sample does not include those who work for a firm that does not carry workers' compensation insurance; however, I do not have a compelling reason to believe that the diurnal injury pattern would be different among firms opting out of the workers' compensation insurance system.

In addition to the time of day the worker was injured and the worker's age at time of injury, the TWCC provided information about the type, severity, and body part of injury. Over 400,000 injuries are recorded in the TWCC data; of these, I examine the 42,902 severe fractures or lacerations and the 29,074 severe falls. The primary advantage of focusing on these injuries is that they are acute and likely to be reported immediately, whereas back injuries, for example, are caused by cumulative conditions, and hence the time when they are reported is somewhat arbitrary.

3.2.2. Current Population Survey

Data on worker schedules come from the May 2001 Workers' Schedules Supplement (WSS) to the Current Population Survey (CPS). In addition to such common CPS data as age, education, industry, occupation, race, and gender, the 2001 WSS provides data on when each worker's shift usually began and ended. For respondents who reported that they work a regular work schedule, I use shift beginning and end times to determine whether a worker was at work during each hour of the day.

Workers in the WSS report their usual shift in two ways: by reporting the usual times they start and end work, and by a categorical description of the hours when they work. In the latter case, they indicate whether their shift is best described as a regular daytime schedule, an evening shift, a night shift, a rotating shift, a split shift, or an irregular schedule. Of the 47,047 observations I begin with, 9,636 lack data on either when the shift usually began or when the shift usually ended (or both).

In order to maximize my sample size, I have used the categorical description of the shift to impute the start and end times whenever possible. Among those who reported their shift start and end times, there is wide variation in actual schedules within the rotating shift category, as well as within the split shift and irregular shift categories. However, within the day, evening, and night shift categories, the typical start and end times are quite consistent. Thus, for these three schedule categories, when the individual either refused or did not know when the shift usually started, or if she reports that the start time varies, I coded her start time as the adjusted median start time of those who worked the same type of shift—day, evening, or night—but *did* report their start and end times. I used an adjusted median because shift times are reported as the time of day (on a 24-hour clock), and simply using the median would incorrectly estimate the usual schedule for each type of shift. For example, using the median would consider midnight as a very late time, and 1:00 AM as a very early time. Conceptually, however, it is usually more reasonable to consider 1:00 AM one hour *later* than midnight, rather than 23 hours earlier.

To get around this issue, for each of these three shifts I bisected the 24-hour day into two 12-hour segments, one for the shift start time and one for the shift end time. I limited the calculation of median shift times to those with a start time within the first segment and an end time within the second, and iteratively selected the 12-hour windows to maximize the number of observations used in the calculation for that shift.[36]

I rescaled each of the 12-hour windows such that when I calculate the median, the beginning of the window was treated as early and the end of the window as late. I then took the median of the rescaled windows. The resulting imputed shift times for those who reported working a day shift were 8:00 AM to 5:00 PM; for an evening shift, they were 3:00 PM to 11:00 PM; and for a night shift, they were 9:00 PM to 7:00 AM. For the day and

[36] The particular windows I chose using this procedure are as follows:
Day Shift: Begins between 1:00 AM and 1:00 PM; ends between 1:00 PM and 1:00 AM.
Evening Shift: Begins between 7:00 AM and 7:00 PM; ends between 7:00 PM and 7:00 AM.
Night Shift: Begins between 12:00 PM and 12:00 AM; ends between 12:00 AM and 12:00 AM.

evening shifts, the shift times I imputed using the adjusted median are identical to the modal shift times, and the modal shift times for night shifts are very close to the adjusted median, providing additional support for these imputations. 6,849 observations were imputed in this manner, and consequently, I am able to use 44,260 of the original 47,047 observations in the WSS.

3.3. Diurnal Injury Distribution

Table 3.1 reports the share of hours worked in each of the 24 time intervals, which I compute from the WSS for each of three age groups, presented in the first three columns, and for all workers between 21 and 69 years old (inclusive), presented in the final column. All calculations are weighted using the WSS sample weights.

Two important features stand out. First, hours of work are heavily concentrated during the day: a full 80 percent of the share of hours worked fall between 8:00 AM and 5:00 PM. Second, the distribution of hours throughout the day is remarkably similar for all three age groups. Only slightly more 21- to 39-year-olds work evening shifts than 40- to 49-year-olds and 50- to 69-year-olds, while the latter two groups work marginally more during normal business hours.

I use the TWCC data to compute the share of injuries incurred during each hour-long interval for two distinct categories of injuries: fractures and lacerations, and falls. These injury shares are then weighted by the share of hours in each interval *t* by taking the ratio of the share of injuries to the share of hours for each age group *a* from the WSS to obtain the injury ratio:[37]

$$(3.1) \qquad Injury \quad Ratio_{at} = \left(\frac{Injuries_{at}}{\sum_t Injuries_{at}} \right) \Bigg/ \left(\frac{Hours_{at}}{\sum_t Hours_{at}} \right).$$

If there is a constant hazard of being injured, the ratio in equation 3.1 should be constant at one across the day; that is, an increase in the share of injuries should be offset by a commensurate increase in the share of hours worked in that time interval. The results of this calculation for severe lacerations and fractures are compiled in figure 3.1 and for falls in figure 3.2. Both figures demonstrate that the injury rate is far from constant. Indeed, the injury ratio is almost three times higher very early in the morning than it is at mid-afternoon. The two categories of severe injuries display similar patterns, both peaking in the 1:01 AM to 2:00 AM hour, then steadily declining until 8:00 AM, from which point the injury rate stays low and flat until 5:00 PM before gradually rising again through the evening.

The profiles of the three age groups in figures 3.1 and 3.2 reveal little discernible difference between the older and younger workers in the injury rates throughout the day, a point to which I will return in the next section. There is wider dispersion between the age

[37] Note that this formula is equivalent to taking the injuries per hour and multiplying by a constant that depends only on the age group. Hence, the informational content of the injury *ratio* calculation used in equation 3.1 is the same as calculating an injury *rate*, which is the number of injuries divided by the number of hours. However, because the numerator and denominator in my analysis come from two separate sources, I have instead used the ratio of injury shares to hour shares to avoid confusion and to provide a statistic with a more apparent interpretation. Because of this relationship, for the remainder of the paper I use the terms "injury ratio" and "injury rate" interchangeably.

groups in the late night and early morning hours compared with the normal business hours of 8:00 AM to 5:00 PM, but the differences are not systematic, and given the much larger number of observations used in the calculations for daytime hours, it is not surprising that there is more noise in the wee hours of the night.

3.4. Possible Explanations

3.4.1. Compositional Differences

Hamermesh (1996) finds that age is negatively related to the probability of working late at night. Young workers may also be more prone to injuries due to inexperience, which could potentially drive the results. However, there is little difference in the distribution of hours across the day for the three age groups in my sample, as reported in table 3.1 and noted in section 3.3. Furthermore, as mentioned above, the injury ratio patterns reported in figures 3.1 and 3.2 are quite comparable for each age group. This is notable because it indicates that the overall injury pattern is not driven purely by compositional changes in the age of the workforce; that is, the injury pattern is not simply an artifact of a disproportionately young and inexperienced workforce late at night.

However, there are other compositional differences that we should be concerned about. The distribution of injuries throughout the day may be a corollary of the differential distribution of industries and occupations throughout the day. If more dangerous jobs are also more likely to have night shifts, for example, the composition of jobs could entirely explain the dramatic increase in injuries in the hours shortly before and after midnight. The first panel of table 3.2 reports the industry and occupation of employment for those working between 1:01 AM and 2:00 AM, while the second panel shows the same breakdown for those working between 1:01 PM and 2:00 PM. While the shares of blue-collar and white-collar workers are comparable in the early morning (35.7 percent vs. 40.5 percent, respectively), there are fewer than half as many blue-collar jobs as white-collar jobs in the afternoon (25.9 percent vs. 64.2 percent).[38,39]

To examine this more closely, I extracted additional data on workers' compensation from the CPS March 2001 Annual Demographic Survey. For each industry-occupation combination, I compute the percentage of people working in that industry-occupation category who received workers' compensation payments in the previous year, which I use as a measure of the injury rate in that industry-occupation cell. Not surprisingly, there is quite a bit of dispersion between industry-occupation cells: while less than half of a percent of white-collar workers in the commerce industry (including trade, finance, and insurance) received income from workers' compensation insurance, nearly two percent of blue-collar manufacturing workers reported that they received workers' compensation payments. In addition to being more likely to be injured, table 3.2 reveals that blue-collar manufacturing workers comprise 19.7 percent of workers between 1:01 AM and 2:00 AM but only 7.0 percent of workers between 1:01 PM and 2:00 PM. The reverse is true of white-collar commerce workers, who represent only 11.6 percent of the workers early in the morning but 18.1 percent by the afternoon.

---

[38] I label the following as white-collar occupations: executive, professional, administrative, managerial, technician, and sales. Blue-collar occupations are composed of precision production, machine operating, material moving, transportation handling and cleaning, farming, fishing, and forestry. Service occupations include private household services, protective services, and other services.

[39] Industries are divided into five categories: (1) agriculture, mining, and construction, which also includes forestry and fishing; (2) manufacturing, including durable and non-durable goods; (3) infrastructure, comprising transportation, communications, utilities and sanitary services; (4) commerce, comprising retail and wholesale trade, finance, insurance, and real estate; and (5) services, including household, repair, personal, recreational, medical, social, educational, and other professional services.

I test the hypothesis that industry and occupation entirely explain the distribution of injuries throughout the day by creating a weighted average of injuries throughout the day using the percentage of workers in each cell i who received workers' compensation, denoted ω_i, from the CPS March Supplement and the percent of workers at each hour of the day t who were employed in cell i, denoted $S_{i,t}$, from the May 2001 WSS. I use ω_i and $S_{i,t}$ to compute the share of injuries that is explained purely by differential industry and occupation injury rates using the following formula:[40]

(3.2) $$Explained\ \ Share_t = \sum_i \omega_i S_{i,t}\ .$$

Figure 3.3 plots the explained injury share for each hour versus the actual share of fractures and lacerations in each hour, and the actual share of falls in each hour. The injury shares are normalized to sum to one throughout the day for comparability. The dotted horizontal line through the middle of the figure, labeled "Constant Injury Rate," represents the hypothetical flat line that we would observe if the ratio of injuries to hours worked were constant throughout the day. The shape of the explained share plot is similar to the actual injury rates, dipping below the constant injury rate during normal daytime hours and increasing above it during hours late at night and very early in the morning. Nonetheless, the magnitude of the difference between the constant injury rate and the explained share is less than half the difference between the constant injury rate and the actual shares. In other words, differences in industry and occupation compositions throughout the day account for less than half of the variation in injury rates throughout the day.

[40] This formula is similar to calculations that Shimer (1998) and Katz and Krueger (1999) use to examine how age and education composition differences in the United States workforce can account for unemployment patterns over the past three decades.

3.4.2. Fatigue

One explanation for the observed diurnal injury rate pattern could be fatigue. Most of the people working between 10:01 AM and 11:00 AM have only been at work for a couple of hours, while most of those at work between 3:01 AM and 4:00 AM have been at work for much longer. Further evidence of this comes from the adjusted median schedule calculations in section 3.2, which show that for night workers the typical schedule is 10 hours long, compared with 9 hours for day workers and 8 hours for evening workers. Such differences in the duration of work at each time of the day can have sizable effects on diurnal injury rates if workers are sensitive to the amount of time on the job.

I consider this possibility by first calculating the cumulative hours worked as of each hour of the day by each worker, and then taking the average cumulative hours in each hour of the day of employees working during that hour. Figure 3.4 displays the average cumulative hours worked for each hour. Although both the average shift duration and injury rate dip during the day and peak at night, the shape and extrema of the shift duration plot are remarkably dissimilar to the injury patterns in figures 3.1 and 3.2. The fractures, lacerations, and falls patterns have wide troughs bottoming out between 4:01 PM and 5:00 PM, while the average cumulative hours profile has a very narrow but deep trough that achieves its minimum between 8:01 AM and 9:00 AM. As cumulative hours on the job rise dramatically from 8:00 AM to 5:00 PM, injury rates remain low and actually decrease slightly. Given the dissimilarities between the diurnal injury patterns and the diurnal fatigue patterns, there is little evidence that fatigue is the lone factor contributing to the late night spike in the injury rate.

3.4.3. Other Physiological Factors

Up until this point I have discussed circumstances that are correlated with working a late shift but are not intrinsic features of late shifts. Working in a more dangerous industry or occupation and working long hours are relatively more prevalent among workers who work night shifts than day shifts, but neither of these can solely explain the high nighttime injury rate. However, Coburn and Moore-Ede (2001) argue that there are inherent characteristics of night activity that affect worker alertness. This may explain why the injury pattern has such large variation throughout the day.

A well-developed body of research in the physiology and neuroscience literature examines biological patterns known as circadian rhythms. Circadian rhythms are a biochemically-regulated process that generates intra-day variation in the body's level of alertness. Wyatt et al. (1999), for example, experimentally assess the influence of circadian rhythms on such behavioral functions as short-term memory, reaction time, and visual vigilance by scheduling episodes of sleep to "desynchronize" circadian rhythms from the duration of wakefulness, thus independently identifying the two processes. They find that functional impairment peaks just after the nadir of the circadian cycle, which is observed in the early morning hours. Although each subject was intensively evaluated throughout the course of 15-24 repetitions of a 20-hour cycle, one limitation of their study is that it is entirely based on only six subjects. However, several related studies conducted by many of the same researchers have found very similar effects of the circadian cycle on alertness. Although the evidence presented in this chapter is not conclusive, these experimental studies coupled with the empirical results in sections 3.3 and 3.4 provide strong evidence that workers are sub-optimally alert during night shifts, contributing to hazardous work conditions for themselves as well as their fellow employees.

3.5. Conclusion

There are both supply-side and demand-side reasons workers might work at night. As Hamermesh notes in his 1996 book *Workdays, Workhours, and Work Schedules*, working unusual times is a more usual event than we might expect. He finds that women with young children often choose to work late at night, arguably because of a lack of affordable childcare during the day. On the demand side, firms can potentially increase the productive capacity of a plant by sustaining a night shift to supplement the day shift.

However, there is a trade-off for firms employing night shift workers. As this chapter has demonstrated, injuries are much more prevalent late at night than during normal business hours. The evidence presented here suggests that this is not simply because of compositional changes in the age, industry, or occupation of late-night workers. I also reject the idea that the difference is attributable to late-night workers having been at work longer. My failure to find sufficient explanation from any of these factors leads me to the conclusion that inherent features of night work make it more hazardous than day work, and this conclusion is consistent with evidence in the physiology literature.

3.6. References

Coburn, Ed and Martin Moore-Ede. 2001. "Keeping the Night Shift Alert," *Journal of Workers' Compensation*, 10(2), pp. 22-35.

Hamermesh, Daniel S. 1996. *Workdays, Workhours, and Work Schedules* (W.E. Upjohn Institute for Employment Research: Kalamazoo, MI).

Katz, Lawrence F. and Alan B. Krueger. 1999 "The High-Pressure U.S. Labor Market of the 1990s," *Brookings Papers on Economic Activity*, pp. 1-87.

OECD Nuclear Energy Agency. *Chernobyl—Ten Years On*. 1995. Paris, France: Organisation for Economic Cooperation and Development. Available from www.nea.fr.

President's Commission on the Accident at Three Mile Island. 1979. *Report of the President's Commission on the Accident at Three Mile Island.* (US Government Printing Office: Washington, DC).

Shimer, Robert. 1998. "Why is the U.S. Unemployment Rate So Much Lower?" in *NBER Macroeconomics Annual 1998*, edited by Ben Bernanke and Julio Rotemberg (MIT Press: Cambridge, MA).

Wyatt, James K., Angela Ritz-De Cecco, Charles A. Czeisler, and Derk-Jan Dijk. 1999. "Circadian Temperature and Melatonin Rhythms, Sleep, and Neurobehavioral Function in Humans Living on a 20-h Day," *American Journal of Physiology*, 277, pp. R1152-R1163.

Table 3.1. Share of Hours Worked in Each Hour of the Day

	Age Group			
Hour of Day	*21 to 39*	*40 to 49*	*50 to 69*	*21 to 69*
24:01 to 1:00	0.007	0.006	0.005	0.006
1:01 to 2:00	0.007	0.007	0.006	0.007
2:01 to 3:00	0.007	0.006	0.006	0.007
3:01 to 4:00	0.007	0.006	0.006	0.007
4:01 to 5:00	0.007	0.007	0.006	0.007
5:01 to 6:00	0.009	0.009	0.008	0.009
6:01 to 7:00	0.017	0.018	0.018	0.018
7:01 to 8:00	0.039	0.044	0.043	0.041
8:01 to 9:00	0.079	0.084	0.085	0.082
9:01 to 10:00	0.090	0.093	0.094	0.092
10:01 to 11:00	0.092	0.095	0.096	0.094
11:01 to 12:00	0.093	0.095	0.097	0.095
12:01 to 13:00	0.092	0.095	0.096	0.094
13:01 to 14:00	0.092	0.094	0.095	0.093
14:01 to 15:00	0.091	0.093	0.094	0.092
15:01 to 16:00	0.087	0.087	0.087	0.087
16:01 to 17:00	0.073	0.070	0.070	0.071
17:01 to 18:00	0.032	0.029	0.027	0.030
18:01 to 19:00	0.019	0.016	0.016	0.017
19:01 to 20:00	0.014	0.012	0.012	0.013
20:01 to 21:00	0.013	0.010	0.010	0.011
21:01 to 22:00	0.012	0.009	0.009	0.010
22:01 to 23:00	0.011	0.008	0.008	0.009
23:01 to 24:00	0.008	0.007	0.007	0.008

Source: Author's calculations from Current Population Survey May 2001 Workers' Schedules Supplement. Total sample size is 44,260, including 20,125 ages 21 to 39; 12,761 ages 40 to 49; and 11,374 ages 50 to 69.

Table 3.2. Occupation and Industry Composition 1:01 AM to 2:00 AM vs. 1:01 PM to 2:00 PM

	1:01 AM to 2:00 AM				***1:01 PM to 2:00 PM***			
	Type of Occupation				***Type of Occupation***			
Industry	*White Collar*	*Blue Collar*	*Services*	*Industry Share*	*White Collar*	*Blue Collar*	*Services*	*Industry Share*
Agriculture, Mining, Construction	0.008	0.036	0.000	0.044	0.026	0.088	0.001	0.115
Manufacturing	0.030	0.197	0.003	0.230	0.068	0.070	0.002	0.140
Infrastructure	0.046	0.047	0.005	0.098	0.040	0.029	0.002	0.070
Commerce	0.116	0.050	0.059	0.225	0.181	0.035	0.020	0.236
Services	0.206	0.026	0.170	0.403	0.327	0.037	0.075	0.440
Total Occupation Share	0.405	0.357	0.238	1.000	0.642	0.259	0.099	1.000

Source: Author's calculations from Current Population Survey May 2001 Workers' Schedules Supplement.

Figure 3.1. Share of Lacerations and Fractures Over Share of Hours by Hour Worked
All Workers Under 70

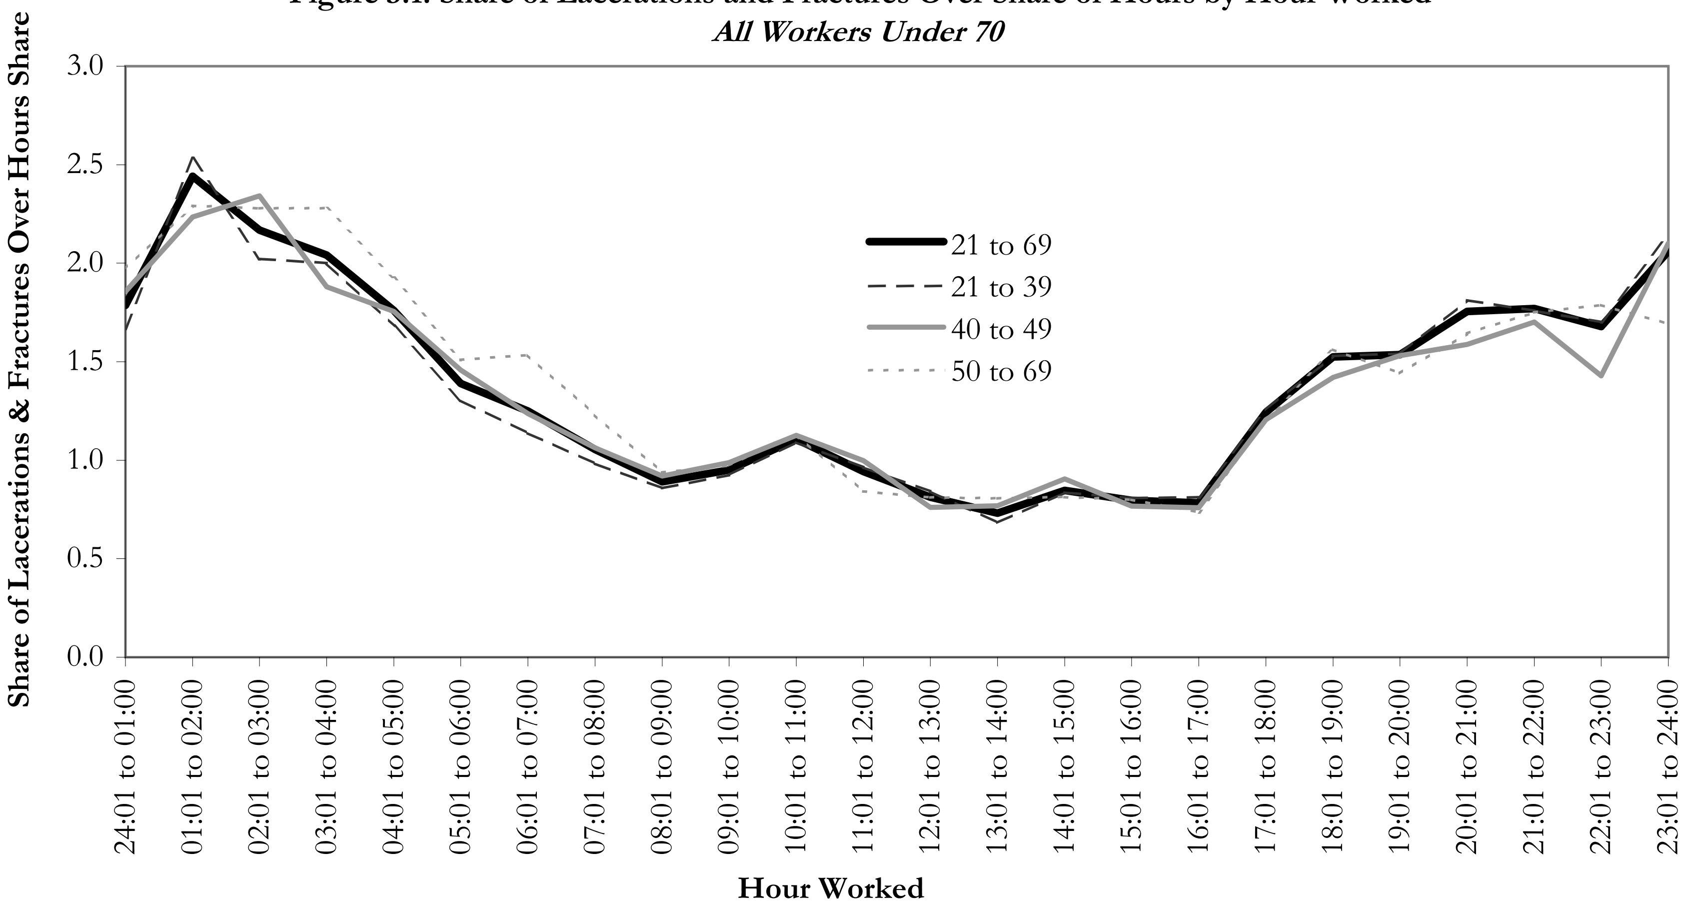

Source: Author's calculations from Current Population Survey May 2001 Workers' Schedules Supplement and Texas Workers' Compensation Commission unpublished data.

Figure 3.2. Share of Falls Over Share of Hours by Hour Worked
All Workers Under 70

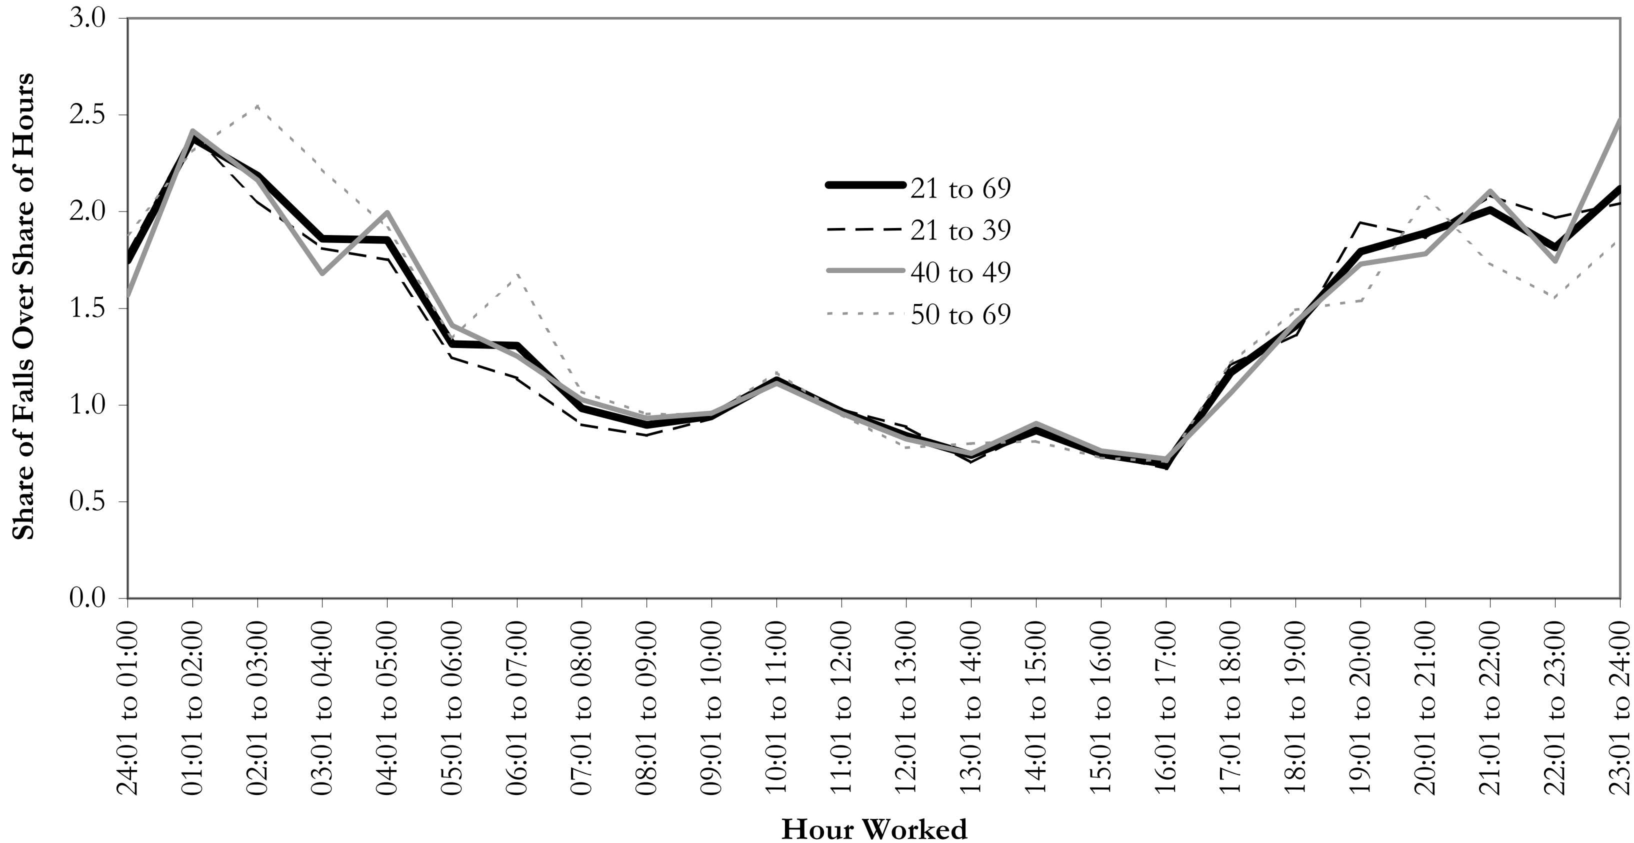

Source: Author's calculations from Current Population Survey May 2001 Workers' Schedules Supplement and Texas Workers' Compensation Commission unpublished data.

Figure 3.3. Actual Injury Share vs. Share of Injuries Explained by Industry and Occupation Composition

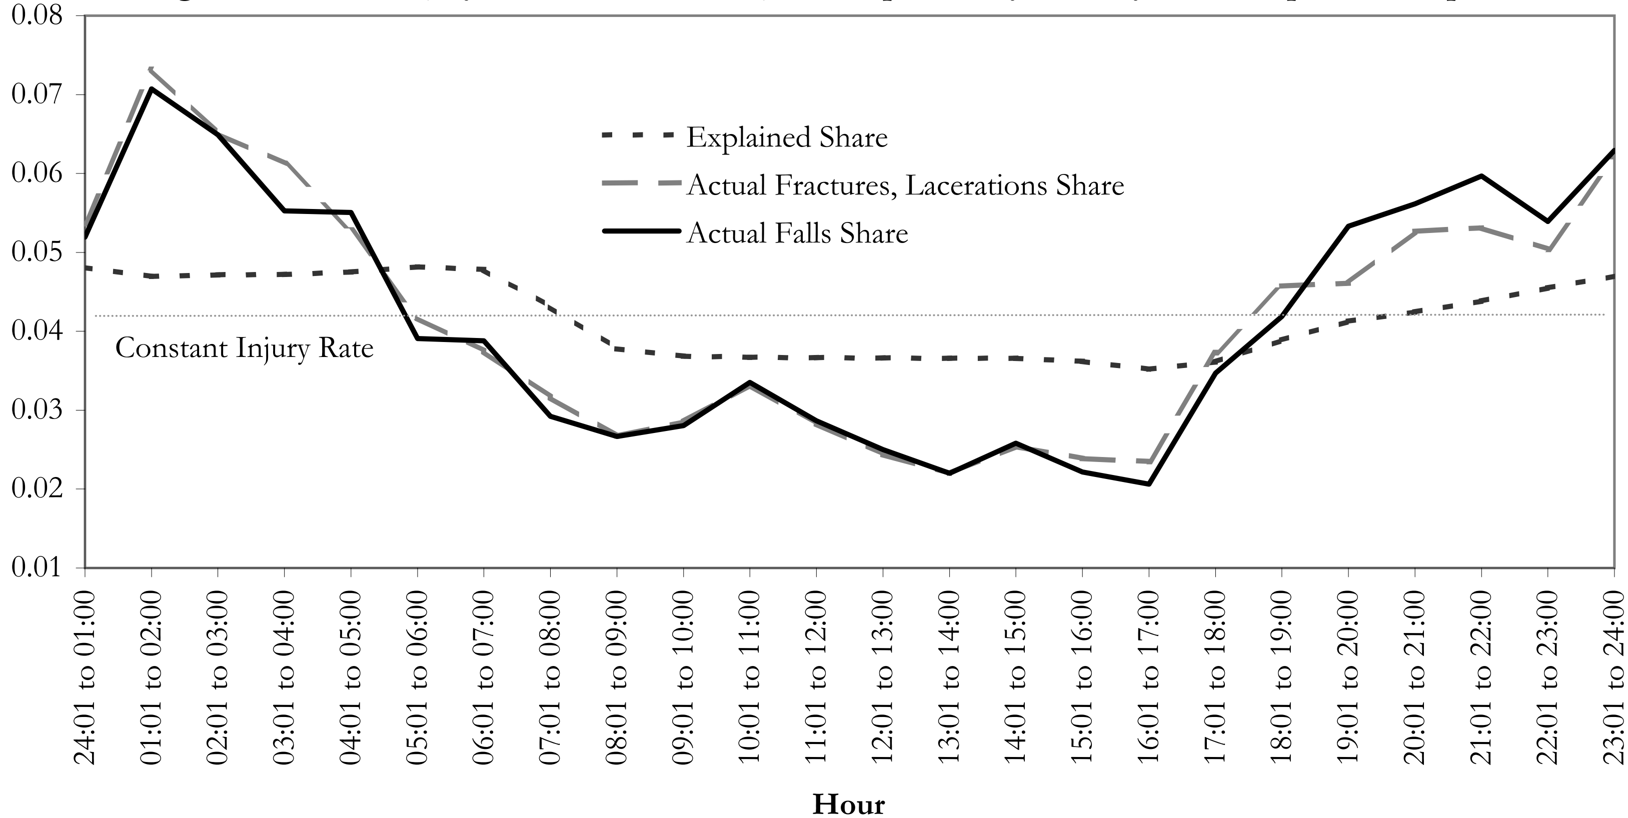

Source: Author's calculations from Current Population Survey May 2001 Workers' Schedules Supplement, Texas Workers' Compensation Commission unpublished data, and Current Population Survey March 2001 Annual Demographic Survey.

Figure 3.4. Average Cumulative Hours by Hour Worked
All Workers Under 70

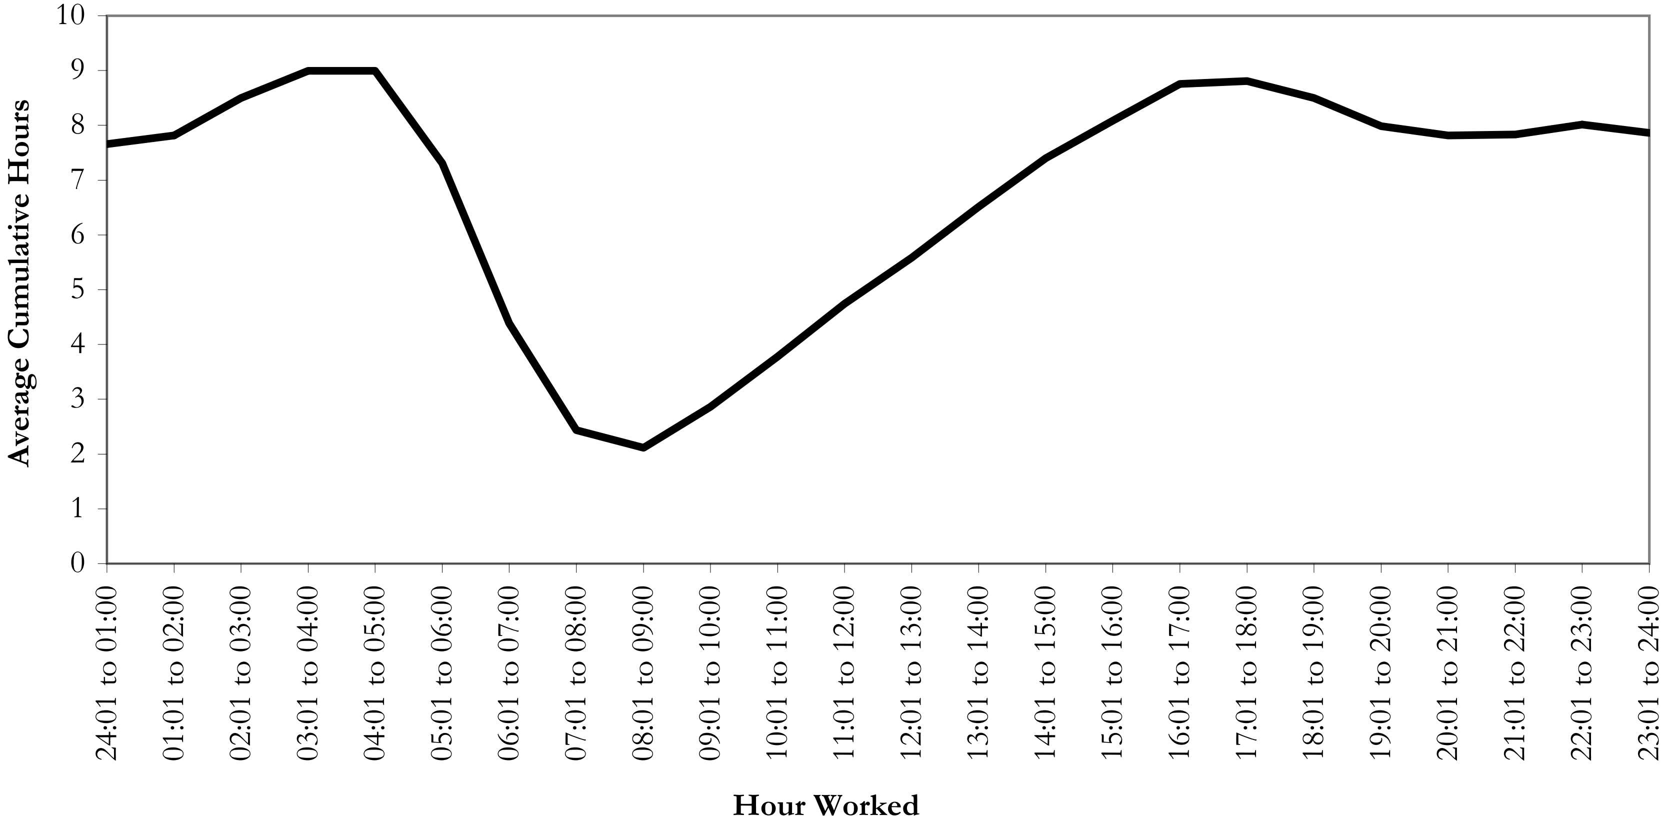

Source: Author's calculations from Current Population Survey May 2001 Workers' Schedules Supplement.