

Modern US Workers' Compensation and Work-Related Injury: No Evidence of Moral Hazard*

Emiliano Huet-Vaughn
Pomona College

Youssef Benzarti
UCSB and NBER

November 2020

Abstract

We exploit recent decades of US state-level reforms to the generosity of workers' compensation programs to estimate the associated moral hazard, utilizing an event-study design and analyzing 9 separate reform categories. The reforms vary - some affecting benefit size, some the probability of receiving benefits; some paid in cash, some in-kind; some constituting increases, some decreases, in generosity. Across the board, we find no evidence of resulting changes in workplace injuries, and, generally, can rule out even moderate moral hazard responses for severe and less severe injuries, suggesting a key moral hazard cost of workers' compensation benefits is negligible.

*Thank you to ProPublica for their data. Anna Newman provided excellent research assistance. Funding for this research was provided by the Pomona College. *Email:* Benzarti: benzarti@ucsb.edu. Huet-Vaughn: emiliano.huet-vaughn@pomona.edu.

1 Introduction

Workers' compensation programs are ubiquitous around the world and virtually all US states mandate some form of insurance against workplace injury. This type of insurance is thought to be critical because workplace injuries can prevent workers from being able to work, thus, exposing them to a loss of income at a moment when health expenses are likely to be large.

While such a benefit is clear, one common argument for making workers' compensation less generous is that it is rife with moral hazard, and, in fact, in recent years many US states have rolled back workers' compensation generosity, making it harder to file claims or sometimes cutting benefits.¹ Moral hazard in the workers' compensation case can be due to workers or to firms.² On the workers' side, the possibility of receiving compensation in the event of an injury could lead workers to be less careful, consciously or not, leading to more accidents. Workers may also be more likely to claim or continue benefits conditional on an accident when the benefits are generous. Some even argue that workers could injure themselves intentionally at work in order to qualify for these benefits, or, falsely claim worker injuries to get them. Firms also play a role in moral hazard. By investing in safety, they can reduce the extent of injuries and thus the total benefits disbursed, and, they have more incentive to do so (as the party eventually paying for experience rated insurance) when their payouts go up with a more generous workers' compensation system.

In this paper, we estimate the combined effect of firms' and workers' contribution to moral hazard on the real injury margin of behavior. To do so, we use the last two decades of state-level reforms to workers' compensation program generosity. We implement a standard event-study design to estimate the effect of these staggered, state-specific reforms on both nonfatal and fatal work-related injuries (separately).³ The main identification assumption of this design is

¹This follows a long period of increases in workers' compensation generosity in the middle and late twentieth century (Allen (2015)). For details on the recent reforms see Yuan and Fishback (2020) and reporting from NPR and ProPublica (both of which we discuss further in our data section) found here: <https://www.propublica.org/article/the-demolition-of-workers-compensation>

²Some may find it odd to apply the terminology of moral hazard to the firm since, unlike the worker, the firm is not the beneficiary of the workers' compensation insurance payout as in the textbook moral hazard consideration; however, this is done in the prior workers' compensation literature (e.g. Butler and Worrall (1991); Guo and Burton Jr (2010)), and, we follow precedence by using this terminology broadly here to refer to worker and firm level behavioral responses to workers' compensation insurance.

³By using governmental administrative record of these injuries we can isolate real injury responses from false reports and claims-reporting behavior. The fatal injury outcome, in particular, is immune to false reports (for this outcome the moral hazard consideration is not intentional deadly injury, but, rather, generally relaxed caution and safety on the job in the presence of more generous injury compensation, leading to inadvertent death.)

the same as that for difference-in-differences regressions: absent the change in generosity to a state’s workers’ compensation program, there would have been no change in the number of fatal and nonfatal injuries relative to states where no change in workers’ compensation occurred. A common test of this assumption is to ensure that the outcomes of interest follow parallel trends prior to the reforms taking place. We systematically test this assumption using a variety of control groups and find that our outcomes indeed follow parallel trends, giving confidence in the research design. We also implement specifications that address recent concerns with staggered event-study designs when there is heterogeneity in the treatment effect across treated cohorts (see a formalization of the issue in [Abraham and Sun \(2020\)](#)) and find robust results regardless of the methodological approach.

Most of the reforms taking place in our sample period resulted in less generous workers’ compensation though some did the opposite. We separately estimate the increases and decreases in generosity, motivated by recent work suggesting that asymmetries along this margin exist in a variety of contexts (see [Benzarti et al. \(2020\)](#) and [Gentry and Viscusi \(2019\)](#)). Because the reforms were varied, affecting different characteristics of the insurance programs in different places and times, we can also separately identify the effect of different categories of change to workers’ compensation. We investigate 9 groupings of categories of reform. Of special interest, motivated by the salience literature (see [Chetty et al. \(2009\)](#) and [Huet-Vaughn \(2019\)](#)), we compare reforms that directly affect benefit size to those that instead affect the probability of receiving benefits (through the tightening or relaxing of legal standards, burdens of proof, or exclusion criteria). Additionally, our data also allows us to compare the moral hazard effects of changes to cash benefits versus changes to in-kind benefits (used for medical care), exploring the importance of differences in the form of the benefit structure ([Thurow \(1974\)](#); [Nichols and Zeckhauser \(1982\)](#)).

Our main finding is that across all of these dimensions of difference, the last twenty years of US workers’ compensation reform have had no detectible effect on workplace injuries, suggesting that this type of moral hazard is not a major factor in aggregate in the modern workers’ compensation system. The null effects we find are robust and precisely estimated, generally ruling out much more than single digit percent changes in injury rates resulting from the reforms. While we do not estimate whether workers’ welfare was affected by added restrictions on the programs and therefore cannot speak to the optimal level of workers’ compensation, we can say that since our

estimated moral hazard cost is small, it is likely that the recent trend of restrictions were not warranted from a welfare maximizing perspective if real injury margins are the main perceived moral hazard cost.

Our paper contributes to the literature on moral hazard in workers' compensation insurance programs by being the first to use reforms that occurred in this century and throughout the entire United State to estimate their effect on workplace injuries. The findings differ from much of the prior work using older US reforms (without event-study research designs) which document evidence of significant moral hazard in the workers' compensation program along a variety of margins of response (Moore and Viscusi (1989); Krueger and Burton Jr (1989); Krueger (1990); Butler and Worrall (1991); Butler (1994); Ohsfeldt and Morrissey (1997))⁴ Other work using recent years of data and focusing on the state of Texas (Cabral and Dillender (2020)) finds evidence of moral hazard along different margins of behavioral response than the real injury margin we consider here (specifically, the intensive margins of the length of the claim spell and the amount of medical spending while claiming). While the above works are the most related to our own, taken together, the scope of reforms we consider (US-wide), the recency of such reforms (the 21st century), the event-study identification strategy pursued, and, the robust and precise null effect of our findings make the contribution of this paper significant not only in the workers' compensation literature but also more broadly in the public finance literature studying moral hazard in social insurance programs (Chetty and Finkelstein (2013); Schmieder and Von Wachter (2016)).

2 Data

In this paper we analyze 9 separate categories of workers' compensation reform, each taking place across multiple states at different times during the last two decades. These reforms are varied, and within the 9 categories are some reforms that are mirror images of each other in that we categorize increases and decreases separately to allow us to determine if there are asymmetric effects of changing workers' compensation generosity.

The 9 reforms can be grouped into 3 types. The first type is a direct change to monthly

⁴Guo and Burton Jr (2010) find more modest moral hazard effects in this period than the other works cited. Gentry and Viscusi (2019), who we discuss more within, use more recent years in their analysis but still finds evidence of moral hazard in some case cases.

benefit amount or maximum compensation amount. In this group we have 3 separate categories of reform: those reforms that led to increased compensation for some or all workers covered by the workers' compensation program (CI); those reforms that instead led to lower compensation (CL); and, those reforms that increased the compensation a workers' family receives upon an on-the-job death (typically in the form of funeral payments or payouts to a surviving family member), effectively increasing death benefits (DBAI).⁵

The second type of reform consists of indirect changes to workers' compensation generosity made not by changing the size of the benefit amount itself but rather by changing the probability of receiving the workers' compensation benefit or the rules determining who is eligible. Our first category of this type (LS) brings together three distinct changes identified by [Yuan and Fishback \(2020\)](#). One of these changes constitutes the rejection of the use of the liberal construction rule, a rule which, as the default, interprets workers' compensation-related statutes in a way that is favorable to the workers when there are ambiguities about the merits of a workers' compensation claim. Over our time period several states revoked such liberal construction rules, lowering the odds of workers' winning a workers' compensation judgment. The second legal change in the LS category involves the change of law to presume worker intoxication caused a worker's injury (thus invalidating the worker's claim) and requiring the worker to prove otherwise (rather than the burden of proof being on the employer to prove intoxication). This, of course, also lowers the odds that a worker wins a workers' compensation judgment. Finally, the third legal change from [Yuan and Fishback \(2020\)](#) included in our LS category involves the requirement to apportion certain workers' compensation benefits in accordance with the degree to which a pre-existing condition was not also a factor in the injury (requiring a worker to meet a higher burden of proof than otherwise in order to receive full benefits, since they must demonstrate no pre-existing condition).

In addition to these three specific reforms, we also define two catchall categories identifying other sorts of reforms that similarly tighten or loosen, respectively, the eligibility rules and affect the probability of a worker receiving workers' compensation payments.⁶ The first of these,

⁵In our data period there are no recorded decreases to death benefits.

⁶We keep the LS category separate from the LSH category introduced here because the [Yuan and Fishback \(2020\)](#) data underlying the LS category allows us to use a slightly longer data window than does the ProPublica data used to form the LSH category, as discussed further below (the two categories also, as noted in the text, refer to different kinds of tightening of the eligibility rules). A similar consideration underlies the use of distinct in-kind categories of MC and MCH discussed below.

representing a tightening of legal standards (LSH) that effectively lower the chance of receiving benefits, includes the following as common examples grouped together in this category of reform: changes to worker compensation laws that increase the causation standard necessary to prove a claim and raise the burden of proof on the worker to show harm; the capping of attorney fees; the revocation of program eligibility for certain categories of worker; the limiting of the timespan for which a worker can claim benefits, among other sundry reforms. Reforms in the opposite direction, which ease legal standards and make it easier to qualify for workers' compensation, are grouped together in a separate category of reform (LSE).

The third and final type of reform to the workers' compensation program that we consider involves changes to the medical benefits paid on behalf of injured workers. The reforms of this type involve effective changes to in-kind benefits rather than to cash transfers. There are three categories of reform of this type that we analyze. The first category (MED) again groups together two distinct reforms identified by [Yuan and Fishback \(2020\)](#). These are changes to the laws governing medical fee schedules paid to doctors for workers' compensation claims and changes to the physician network that a beneficiary can access. In the first case, the reforms specifically limit the amount that can be paid for medical care by the workers' compensation programs on behalf of an injured worker (reducing effective in-kind benefits). In the second case, reforms specifically limit the physicians that a workers' compensation beneficiary can visit (also effectively reducing the value of the workers' compensation). Other sundry reforms that make preferred medical care harder to get are grouped into a second catchall category (MCH), and include a variety of other types of reforms (e.g. limitations on the ability to change a doctor; the setting of windows of time after which medical care terminates; deference to company medical experts, among other changes). Reforms that make medical easier to get are grouped into a third category (MCE), and have the opposite effect as MCH reforms.

The LS and MED reform data comes from [Yuan and Fishback \(2020\)](#) and cover years 1997 to 2016. The authors in part draw on a ProPublica database containing a more comprehensive set of all types of state-level reforms to the workers' compensation program.⁷ We use the additional information in the ProPublica database to define the event times for the remaining non-Yuan and Fishback reforms (i.e. the remaining 7 categories listed above, CI, CL, DBAI, LSH, LSE,

⁷ProPublica's workers' compensation reform database can be found here <https://projects.propublica.org/graphics/workers-comp-reform-by-state>

MCH, MCE). The ProPublica database covers the years 2002 to 2014 for all 50 states.

For work-related injury data we take data from the Bureau of Labor Statistics (BLS). Hours-based nonfatal injury rate data for the private sector (measured per 100 workers) was obtained from the BLS’ state occupational injury files at the state-year level.⁸ While BLS does not provide a comparable fatal injury rate for our time span, they do provide the raw counts of fatal workplace injuries at the state-year level. We use these to calculate an employment-based fatality rate, dividing this raw count by yearly employment in the state (which was computed by averaging monthly employment numbers for a state, also provided by BLS, to obtain average yearly employment, and then multiplying that by 100,000 to get the fatality rate per 100,000 workers).

3 Results

To estimate the effect of the generosity of workers’ compensation programs on workplace fatal and nonfatal injuries, we use an event-study design (Evans et al. (1999), Bailey and Goodman-Bacon (2015), Fuest et al. (2018)) that leverages changes in workers’ compensation programs over time and across states. Specifically, we apply the following econometric model to study one by one the causal effect of each of the 9 workers’ compensation reform categories:

$$\log(y_{st}) = \sum_{k=-4, k \neq 0}^{k=6} \beta_k \mathbb{1}_{kst} + \lambda_t + \psi_s + X_{st} + \epsilon_{st}, \quad (1)$$

where $\log(y_{st})$ represents the log of the fatal or nonfatal (depending on the specification) workplace injury measure in state s at time t (see Section 2 for details on the two measures), λ_t represents time fixed effects, ψ_s represents state fixed effects, and ϵ_{st} is the error term, clustered at the state level initially. The treatment dummy $\mathbb{1}_{kst}$ equals 1 if in state s there is a reform for the category of reform under study k years from date t . This definition implies that $k = -1$ represents the year just prior to the workers’ compensation reform, and, $k = 1$ is the year in which the reform happened. As is common in such specifications, we omit the $k = -1$ period, making it the reference category. Note that we use a smaller event window when estimating the effect of changes identified from reforms in the ProPublica data set because the panel dataset

⁸This rate is equal to $\frac{N}{EH} 200,000$ where N = the number of workplace injuries and illnesses recorded in the state-year, EH = total hours worked by all employees during the calendar year, and, 200,000 represents the base for 100 full-time equivalent workers (working 40 hours per week, 50 weeks per year). For data see <https://tinyurl.com/yymebv6j>

covers fewer years than is the case with the data derived from [Yuan and Fishback \(2020\)](#). In the former case, the event window is instead 4 years prior and 3 years post-reform (rather than 4 and 6, respectively). When $k < -4$ we additionally include a $\beta_{<-4}$ event dummy for years outside the event window that are well before the reform, and, similarly, we include $\beta_{>6}$ (or, $\beta_{>3}$, depending on which data source is used) for years outside the event window that are well past the reform. Finally, indicator X_{st} controls for post-event periods for other categories of reform in state s and period t .

Thus, in equation (1), β_1 identifies the average percent change in workplace fatal or nonfatal injury (depending on the specification) between the year of reform in the workers’ compensation program and the year just prior to treatment among those states treated by the reform category relative to the same change over time among those states that do not experience the category of reform. β_2 identifies the same for the year after the reform and so on for the other event dummies. To interpret the lagging, post-event β_k terms as measures of the causal effect of the workers’ compensation reform, the main identification assumption needed in this generalized event-study design is the same as that for difference-in-differences regressions: absent the workers’ compensation program change in a state, there would have been no change in the relative trajectories of workplace injuries in treated and untreated states (i.e. the workplace accident series in treated and untreated states would have moved in parallel). We assess the credibility of this assumption below when discussing our graphical evidence, using the leading terms in (1) to measure evidence of any pre-existing trends, as is standard.

We estimate three different versions of equation (1) to ensure that our results are not driven by our choice of specification and to address some concerns described below. Reassuringly, all three approaches yield very similar results, as we subsequently document.

First, we estimate equation (1) while dropping any state where the policy reform analyzed involves multiple changes over multiple years in our data set (leaving us with the set of “treated” states that go from “no reform” status to “reform” status in a single year during our time period, with no subsequent changes to the policy). Therefore, the control group for *Approach 1* is any remaining states that did not experience that particular category of policy reform in our data window.

Our second approach differs from the first by including states with staggered policy reforms that take place over multiple years, while only attending to the first change in the category of

reform. Specifically, for *Approach 2* we assign treatment status of value 0 before the first reform and of value 1 for the year of the first reform and all subsequent years (regardless of whether additional later reforms in the category take place in the state). This means that for all but the estimate of the $t + 1$ post-treatment event-time dummy for a reform category, the post-treatment dummies may be representing a combined effect of delayed or sustained consequences of the original change *plus* any contemporaneous effect of a subsequent change that may exist (leaving the $t + 1$ event-time dummy as the only actual causal estimate of a single reform).

Finally, our third approach is designed with the goal of addressing some common issues with dynamic event-study designs like ours. As described in [Abraham and Sun \(2020\)](#), these issues arise when using a staggered event-study design when there is heterogeneity in the treatment effect across treated cohorts.⁹ In such a case, the estimated event-time dummies can be contaminated by effects from other periods and may produce causally uninterpretable results for the event-time dummy estimates (as the decomposition of the β_k estimates can be shown to potentially assign negative weights to cohort-specific treatment effects). We follow the approach of [Cengiz et al. \(2019\)](#) to address this issue. Intuitively, under *Approach 3*, we proceed as follows: we identify for a given category of reform each instantiation of the reform, then create event-time windows around each one of these events, with a window-size that corresponds to the one we use to estimate equation (1). Next, for each of these individual instances of reform we identify all control states that do not experience this category of reform in the years within the events' event-time window, and, we only keep this treated state and its associated control states and for the years within the event window, thus, creating a unique event-specific dataset for each reform event. Finally, we stack each of the event-specific datasets we create for each instantiation of a given category of reform, and estimate equation (1) on this stacked dataset.¹⁰ We repeat this process for each of the categories of reform we consider, generating many event-specific data sets and a total of 9 resulting stacked data sets, one for each category of reform.

Reassuringly, all three approaches yield very similar results. Figures 1 and 2 show the results of estimating equation (1) using *Approach 1* for each category of reform and nonfatal and fatal

⁹A series of related recent papers make similar points, including [Borusyak and Jaravel \(2017\)](#), [Callaway and Sant'Anna \(2019\)](#), [De Chaisemartin and d'Haultfoeuille \(2020\)](#), [Goodman-Bacon \(2018\)](#).

¹⁰For this approach, standard errors are now clustered at the state by event-specific-data set level, the dummies for years above and below the event window are necessarily excluded given the contraction of the sample to include only the event window years, and, an additional control is added for the post-event period of any additional same-category reform whose 6-year post-treatment period takes place within the event-specific data set, accounting for the staggered nature of some reforms in a state.

workplace injury outcomes, respectively. Figures 6 through 9 in the Appendix show the corresponding results using *Approach 2* and *3*. Importantly, trends appear reasonably parallel across specifications, and in most cases, we cannot reject that trends are parallel. Specifically, out of the 48 separate specifications run, in 39 cases there is not a single pre-event period dummy that is significantly different from zero at the 5% level of significance, providing compelling support for the parallel trends assumption and the lack of pre-trend differences in treatment and control states across these 39 specifications.¹¹ In the remaining 9 cases, the identification assumption may not hold, making inference regarding causal effects in such cases questionable, and, we, thus, do not discuss these cases further (though we report their results along with the others for completeness and note that they do not include significant post-period dummies at all).¹²

Turning from design considerations to the estimates of our main moral hazard results, we systematically find no evidence of an effect of workers' compensation program generosity on either fatal or nonfatal injuries. This is true across the many categories of workers' compensation reform considered and across specifications. In fact, with only 4 exceptions out of the 48 specifications, the various event-study figures do not show a clear discontinuity at the time of the reform, with no significant break from the pre-trend at the 5% level of significance.¹³ The null results are characterized by insignificant point estimates on the $t + 1$ event dummy, for instance, that are often essentially zero and typically no greater than 0.05 in magnitude. Figure 3 summarizes these $t + 1$ term point estimates across the various specifications run for all 9 categories of workers' compensation reform with each outcome and using *Approach 1*, *2*, and *3*. 95% confidence intervals are also included in the figure. While the confidence intervals around these estimates vary depending on the reform and specification, they remain relatively small in most cases, with upper bounds usually around 0.10 in absolute value, representing no more than about a 10%

¹¹There are only 48 and not 54 specifications because the specifications for the LS and MCE reforms with *Approach 1* and *Approach 2* are the same given the absence of staggered reforms among the states implementing either category of reform in our time period (something that is therefore true for specifications with either the fatal or nonfatal outcome), and, because *Approach 3* for the MCE reform category (again, for each outcome) results in a single state receiving treatment status, because of the requirement to have the full event window in the sample period, resulting in no estimate with this method due to collinearity.

¹²The 9 instances include CI fatal with *Approach 1*, MED fatal with *Approach 1/2*, MCE nonfatal with *Approach 1* (which for this reform category is the same specification as *Approach 2* as already noted), DBAI fatal with *Approach 1/3*, LS nonfatal with *Approach 3*, MCH fatal and MCH nonfatal with *Approach 3*.

¹³The exceptions are those figures derived from the MED nonfatal specification via *Approach 1/2* (Figure 1 Panel (b) and Figure 6 Panel (b)) and the LSE nonfatal and DBAI nonfatal specifications via *Approach 3* (Figure 8 Panels (c) and (h)). Only the MED nonfatal and DBAI nonfatal cases have additional lagging terms that approach significance and suggest a sustained effect of the reform. Even in these cases, the the additional significant terms are on the margin of significance, with very small coefficient point estimates in all cases.

change in the relevant form of workplace injury due to policy changes, allowing us to rule out even moderate changes in most cases.

While some of the changes to workers' compensation programs cannot be easily translated into dollars amounts (such as stricter application requirements, stricter burdens of proof etc.), we can quantify the changes to compensation and death benefits. On average, and across all reforms in all the states we consider, compensation increases by \$555 a month, while the cap on compensation increased by \$62,500. Similarly, death benefits increase by \$3,754, while the cap on death benefits increases by \$44,500. Thus, the null results exist for non-trivial changes in benefits.

Further, the variety of reform categories we study allow us to make comparisons across different types of reform. We consider three comparisons of theoretical interest.

First, we assess whether there are asymmetric effects of increases versus decreases to workers' compensation generosity, motivated by prior work indicating such asymmetry may exist. [Gentry and Viscusi \(2019\)](#) find that increases in workers' compensation benefits reduce fatality rates (presumably because firms increase safety-improving investments to avoid higher overall payments to workers' compensation funds) while decreases have no effect (presumably because these investments are not variable inputs and there is no reason to reduce the stock of safety-improving changes so that accidents don't change). However, this prior analysis, while using a similar state-year panel does not present any evidence that the parallel trends identifying assumption is valid in their sample. In our results, on the other hand, across the numerous specifications where this identifying assumption holds, we see no such asymmetry and in fact find null results for both increases or decreases in generosity. This is clearly seen in the comparison of the CI and CL reforms; the effects of increases versus decreases in monthly benefit amount are presented in [Figure 4](#), for ease of comparison, across *Approach 1-3* (and for both outcomes). But, the pattern can also be seen in other reform category comparisons that separate increases and decreases of the same type.

Second, we compare whether there are different effects from changes to generosity that come from direct changes to the amount of the monthly benefit payment rather than from indirect changes that affect the probability of getting the workers' compensation benefits by tightening or relaxing legal standards, burdens of proof, or exclusion criteria. The former can be found in our CI and CL categories for increases and decreases, respectively, while the later reforms are found

in our LSE and LS/LSH categories for more and less stringent standards, respectively. While a direct change in benefit amount might *a priori* be thought to be more salient and thus lead to greater responsiveness, we find no systematic difference across these sets of reforms, with null effects being the norm for both reforms that directly affect the payment amount (conditional on being determined to be qualified) and those that indirectly affect payment by making it harder to win workers' compensation claims, thus affecting expected payment by changing the odds of qualification status. This can be seen in the comparison of columns 1 and 2 of Figure 4, presenting the aforementioned estimates of direct increases in compensation, to columns 1 and 2 of Figure 5, presenting the equivalent estimates of the indirect relaxation of legal standards that make qualifying for benefits easier.

Finally, we are also interested in knowing whether benefits received in cash affect moral hazard differently than benefits received in kind, motivated by long-standing interest in the distinction between the form of transfer in public economics (see, for example, [Thurow \(1974\)](#) and [Nichols and Zeckhauser \(1982\)](#)). In the workers' compensation case, the in-kind benefits are medical care (represented by our MED/MCH and MCE categories for effective decreases and increases, respectively, of the in-kind transfer). Comparing in-kind to cash transfers within the reforms that decrease generosity, using MCH versus CL, for instance, we yet again, see no difference in the basic result: there is no evidence of an effect of changes to workplace compensation generosity on workplace accidents, as can be seen in the last two columns of Figure 4 and Figure 5 which collect the CL and MCH estimates, respectively, across all specifications and outcomes.

4 Conclusion

Using several workers' compensation changes over time and across US states, we find that these have null effects on the number of fatal and nonfatal injuries experienced by workers on the job, which we interpret as evidence that moral hazard effects for workers' compensation programs are limited along the real injury margin. Importantly, this null result holds across all types of changes to the workers' compensation programs, including increases and decreases in compensation amounts, adjustment of the restrictiveness of application requirements, changes to in-kind and cash benefits, and, changes in death benefits.

Are our null results sufficiently precise to rule out large moral hazard responses? While there is limited prior work to reference regarding what constitutes a large magnitude response,

our estimates' 95% confidence interval lower and upper bounds generally allow us to rule out effect sizes of approximately a 10% change in injury rate or greater. In the reform instances where the benefit increases are easily calculable, our estimates allow us to exclude such modest effects sizes even when the change in economic benefits brought on by the workers' compensation reform is sizable. We conclude that there is very little evidence that more generous workers compensation programs will result in significant real injury moral hazard, and, that optimal workers' compensation policy should be designed without serious concern for this "cost" of more generous workers' compensation laws.

Finally, while we study one leading kind of moral hazard, the kind that is associated with more generous workers' compensation insurance causing an increased rate of entry into the adverse state of workplace injury, we note that this is not the only margin of possible moral hazard in the workers' compensation program. For example, [Cabral and Dillender \(2020\)](#) study the effect of changes in the workers' compensation system that cause a change in duration and program costs while in the adverse state. Obviously, these additional margins of possible response, outside the purview of the current work, are also important for optimal program design.

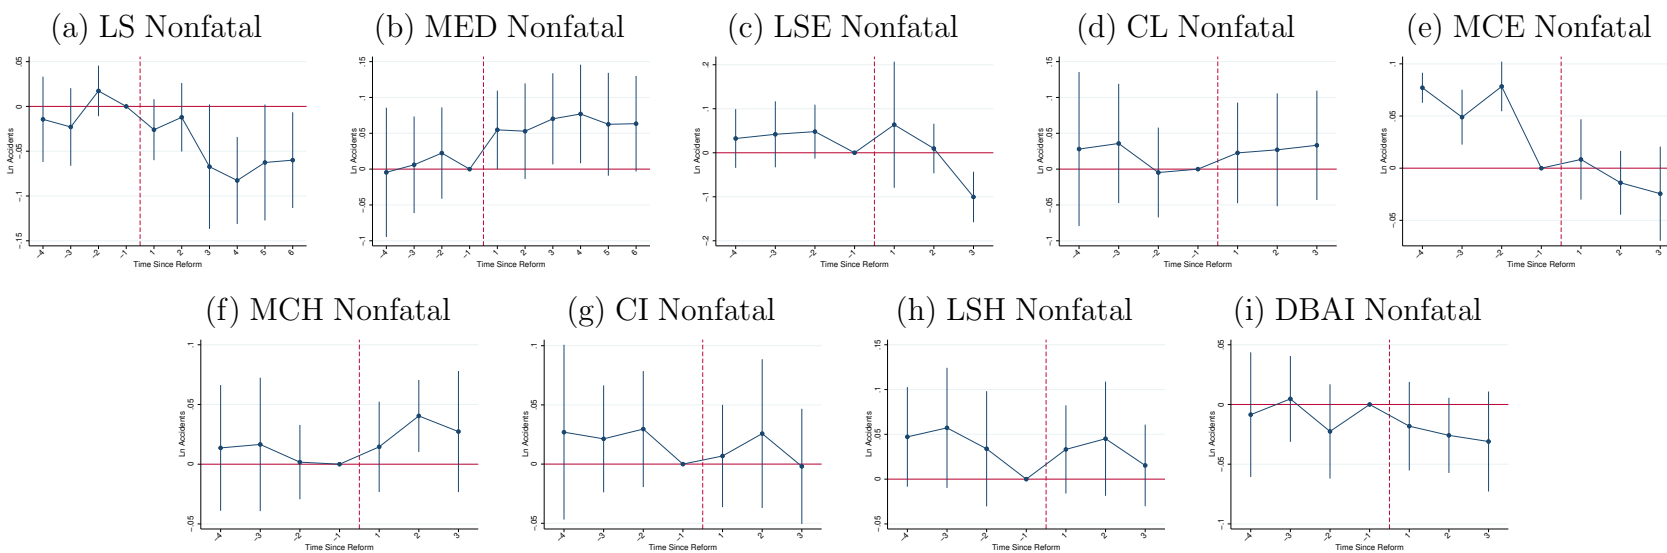
References

- Abraham, Sarah and Liyang Sun**, “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” Technical Report, ArXiv Archive 2020.
- Allen, Samuel K**, “Struggle for Regulatory Power between States and the US Federal Government: The Case of Workers? Compensation Insurance 1930-2000,” *Journal of Economics and Political Economy*, 2015, 2 (3), 351–373.
- Bailey, Martha J and Andrew Goodman-Bacon**, “The War on Poverty’s experiment in public medicine: Community health centers and the mortality of older Americans,” *American Economic Review*, 2015, 105 (3), 1067–1104.
- Benzarti, Youssef, Dorian Carloni, Jarkko Harju, and Tuomas Kosonen**, “What goes up may not come down: asymmetric incidence of value-added taxes,” *Journal of Political Economy*, 2020, 128 (12), 000–000.
- Borusyak, Kirill and Xavier Jaravel**, “Revisiting event study designs,” *Available at SSRN 2826228*, 2017.
- Butler, Richard J**, “Economic determinants of workers’ compensation trends,” *Journal of Risk and Insurance*, 1994, pp. 383–401.
- **and John D Worrall**, “Claims reporting and risk bearing moral hazard in workers’ compensation,” *Journal of Risk and Insurance*, 1991, pp. 191–204.
- Cabral, Marika and Marcus Dillender**, “The Impact of Benefit Generosity on Workers? Compensation Claims: Evidence and Implications,” Technical Report, National Bureau of Economic Research 2020.
- Callaway, Brantly and Pedro HC Sant’Anna**, “Difference-in-differences with multiple time periods,” *Available at SSRN 3148250*, 2019.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer**, “The effect of minimum wages on low-wage jobs,” *The Quarterly Journal of Economics*, 2019, 134 (3), 1405–1454.

- Chaisemartin, Clement De and Xavier d’Haultfoeuille**, “Two-way fixed effects estimators with heterogeneous treatment effects,” *American Economic Review*, 2020, *110* (9), 2964–96.
- Chetty, Raj, Adam Looney, and Kory Kroft**, “Salience and taxation: Theory and evidence,” *American economic review*, 2009, *99* (4), 1145–77.
- **and Amy Finkelstein**, “Social insurance: Connecting theory to data,” in “handbook of public economics,” Vol. 5, Elsevier, 2013, pp. 111–193.
- Evans, William N, Jeanne S Ringel, and Diana Stech**, “Tobacco Taxes and Public Policy to Discourage Smoking,” *Tax policy and the Economy*, 1999, *13*, 1–55.
- Fuest, Clemens, Andreas Peichl, and Sebastian Siegloch**, “Do higher corporate taxes reduce wages? Micro evidence from Germany,” *American Economic Review*, 2018, *108* (2), 393–418.
- Gentry, Elissa Philip and W Kip Viscusi**, “Asymmetric effects on fatality rates of changes in workers’ compensation laws,” *American Law and Economics Review*, 2019, *21* (2), 307–345.
- Goodman-Bacon, Andrew**, “Difference-in-differences with variation in treatment timing,” Technical Report, National Bureau of Economic Research 2018.
- Guo, Xuguang and John F Burton Jr**, “Workers’ compensation: Recent developments in moral hazard and benefit payments,” *ILR Review*, 2010, *63* (2), 340–355.
- Huet-Vaughn, Emiliano**, “Stimulating the vote: Arra road spending and vote share,” *American Economic Journal: Economic Policy*, 2019, *11* (1), 292–316.
- Krueger, Alan B**, “Incentive effects of workers’ compensation insurance,” *Journal of Public Economics*, 1990, *41* (1), 73–99.
- **and John F Burton Jr**, “The employers’ cost of Workers’ Compensation Insurance: magnitudes, determinants, and public policy,” Technical Report, National Bureau of Economic Research 1989.
- Moore, Michael J and W Kip Viscusi**, “Promoting safety through workers’ compensation: The efficacy and net wage costs of injury insurance,” *The RAND Journal of Economics*, 1989, pp. 499–515.

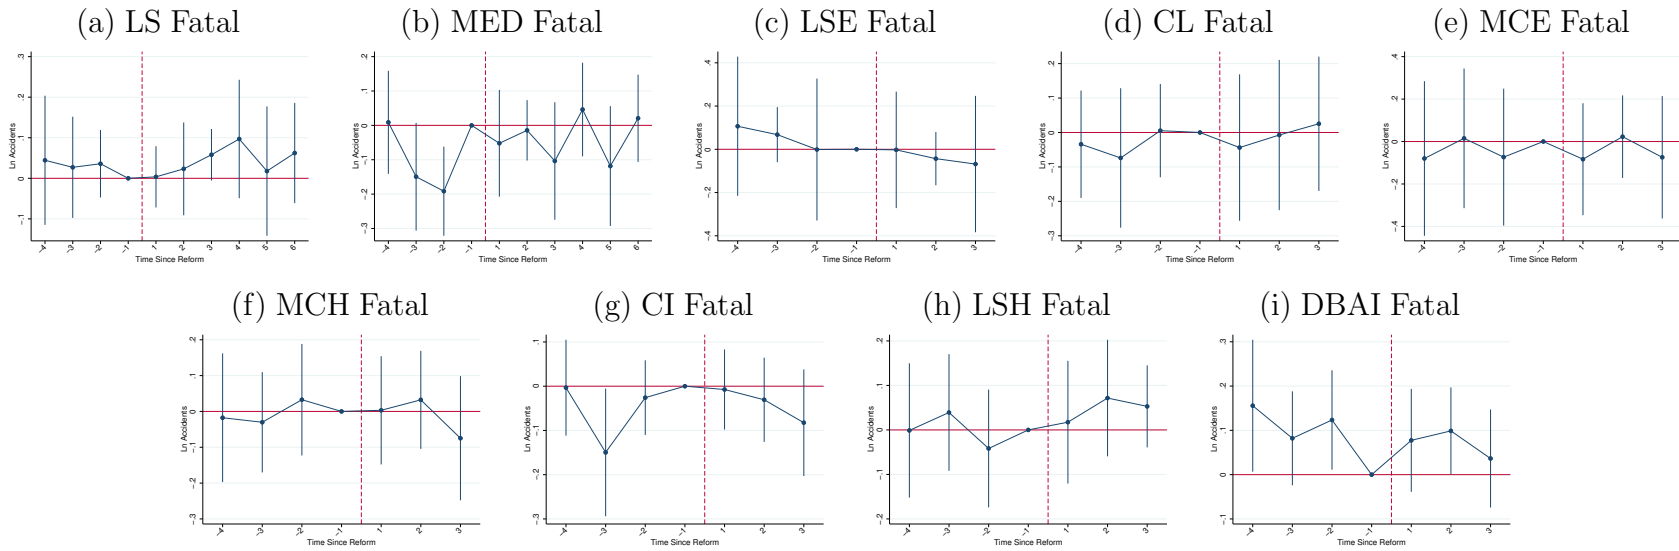
- Nichols, Albert L and Richard J Zeckhauser**, “Targeting transfers through restrictions on recipients,” *The American Economic Review*, 1982, *72* (2), 372–377.
- Ohsfeldt, Robert L and Michael A Morrissey**, “Beer taxes, workers’ compensation, and industrial injury,” *Review of Economics and Statistics*, 1997, *79* (1), 155–160.
- Propublica**, “Workers’ Compensation Reforms by State Dataset,” (accessed September, 2020). 2002-2014.
- Schmieder, Johannes F and Till Von Wachter**, “The effects of unemployment insurance benefits: New evidence and interpretation,” *Annual Review of Economics*, 2016, *8*, 547–581.
- Thurow, Lester C**, “Cash versus in-kind transfers,” *The American Economic Review*, 1974, *64* (2), 190–195.
- Yuan, Andy and Price V Fishback**, “Rising Burdens of Proofs and The Grand Bargain of Workers? Compensation Laws,” Technical Report, National Bureau of Economic Research 2020.

Figure 1: Moral Hazard and Workers' Compensation Reforms: Approach 1 Nonfatal Injuries




Notes: The figure reports the treatment estimates and 95% confidence intervals from separate estimation of Equation (1) under *Approach 1* for all nine categories of workers' compensation reform (with nonfatal injuries as the outcome, in logs). Details of the category definitions and outcome measure are in Section 2. The red dashed vertical line represents the onset of the workers' compensation reform. *Approach 1* drops any state from the sample where the relevant policy reform analyzed involves multiple changes over multiple years in our data set. This means the remaining sample analyzed in these panels includes only "treated" states that go from "no reform" status to "reform" status in a single year during the time period, with no subsequent changes to the policy (with the control group for *Approach 1* being any remaining states that did not experience that particular category of policy reform in our data window).

Figure 2: Moral Hazard and Workers' Compensation Reforms: Approach 1 Fatal Injuries



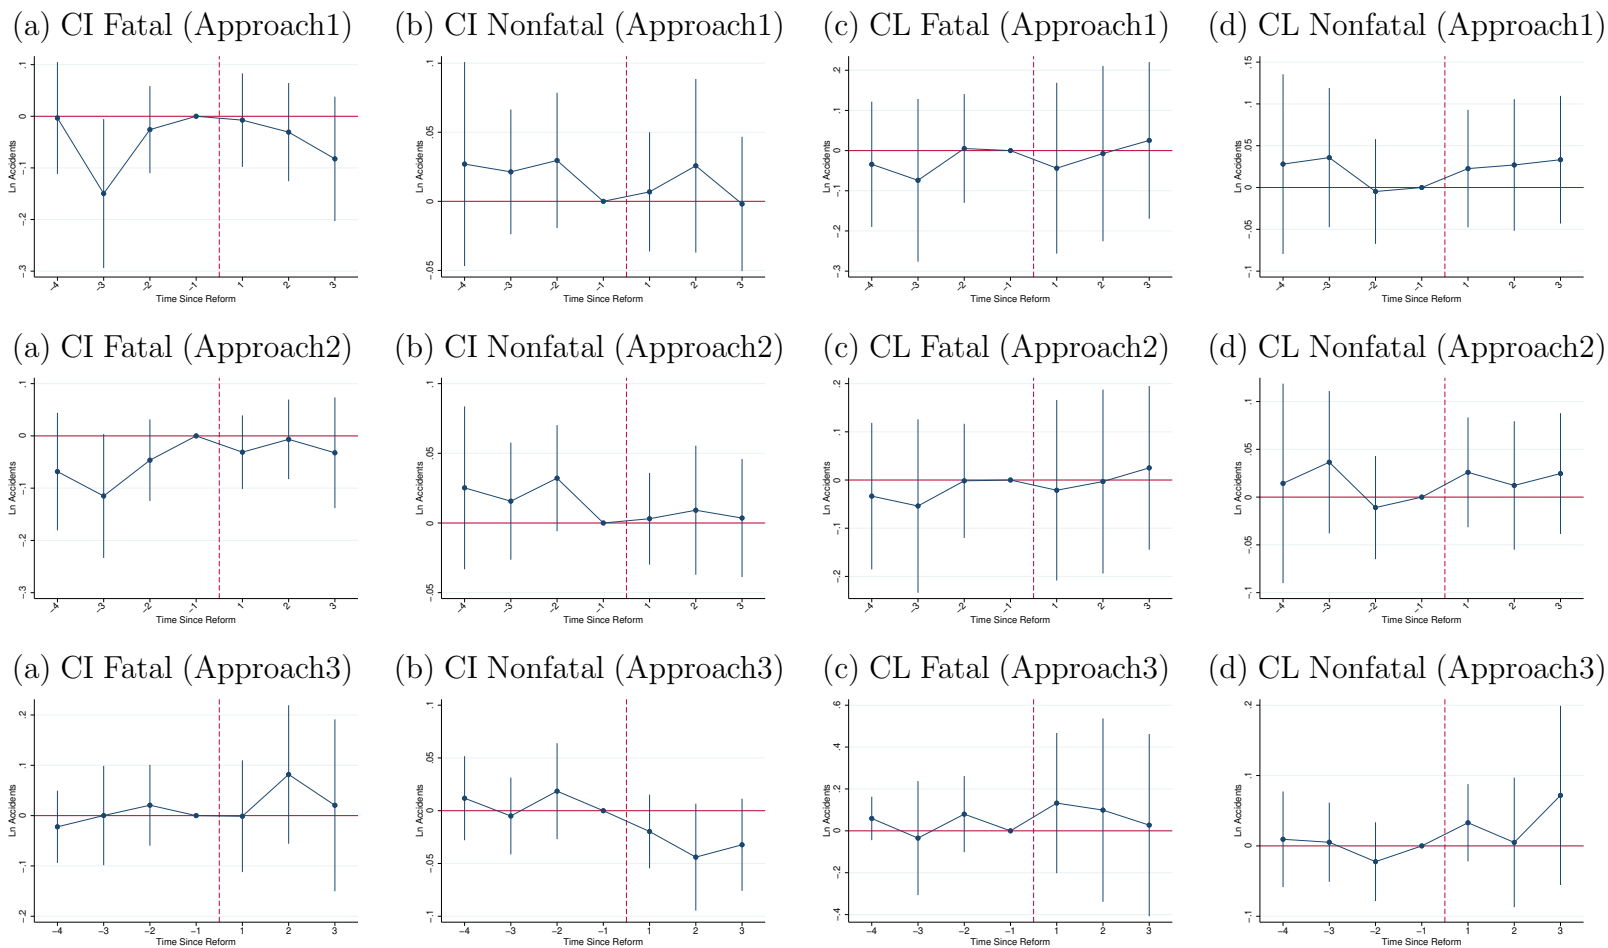
Notes: The figure reports the treatment estimates and 95% confidence intervals from separate estimation of Equation (1) under *Approach 1* for all nine categories of workers' compensation reform (with fatal injuries as the outcome, in logs). Details of the category definitions and outcome measure are in Section 2. The red dashed vertical line represents the onset of the workers' compensation reform. *Approach 1* drops any state from the sample where the relevant policy reform analyzed involves multiple changes over multiple years in our data set. This means the remaining sample analyzed in these panels includes only "treated" states that go from "no reform" status to "reform" status in a single year during the time period, with no subsequent changes to the policy (with the control group for *Approach 1* being any remaining states that did not experience that particular category of policy reform in our data window).

Figure 3: Event-Study $t + 1$ Estimates and 95% CI Across All Reforms and Estimation Approaches



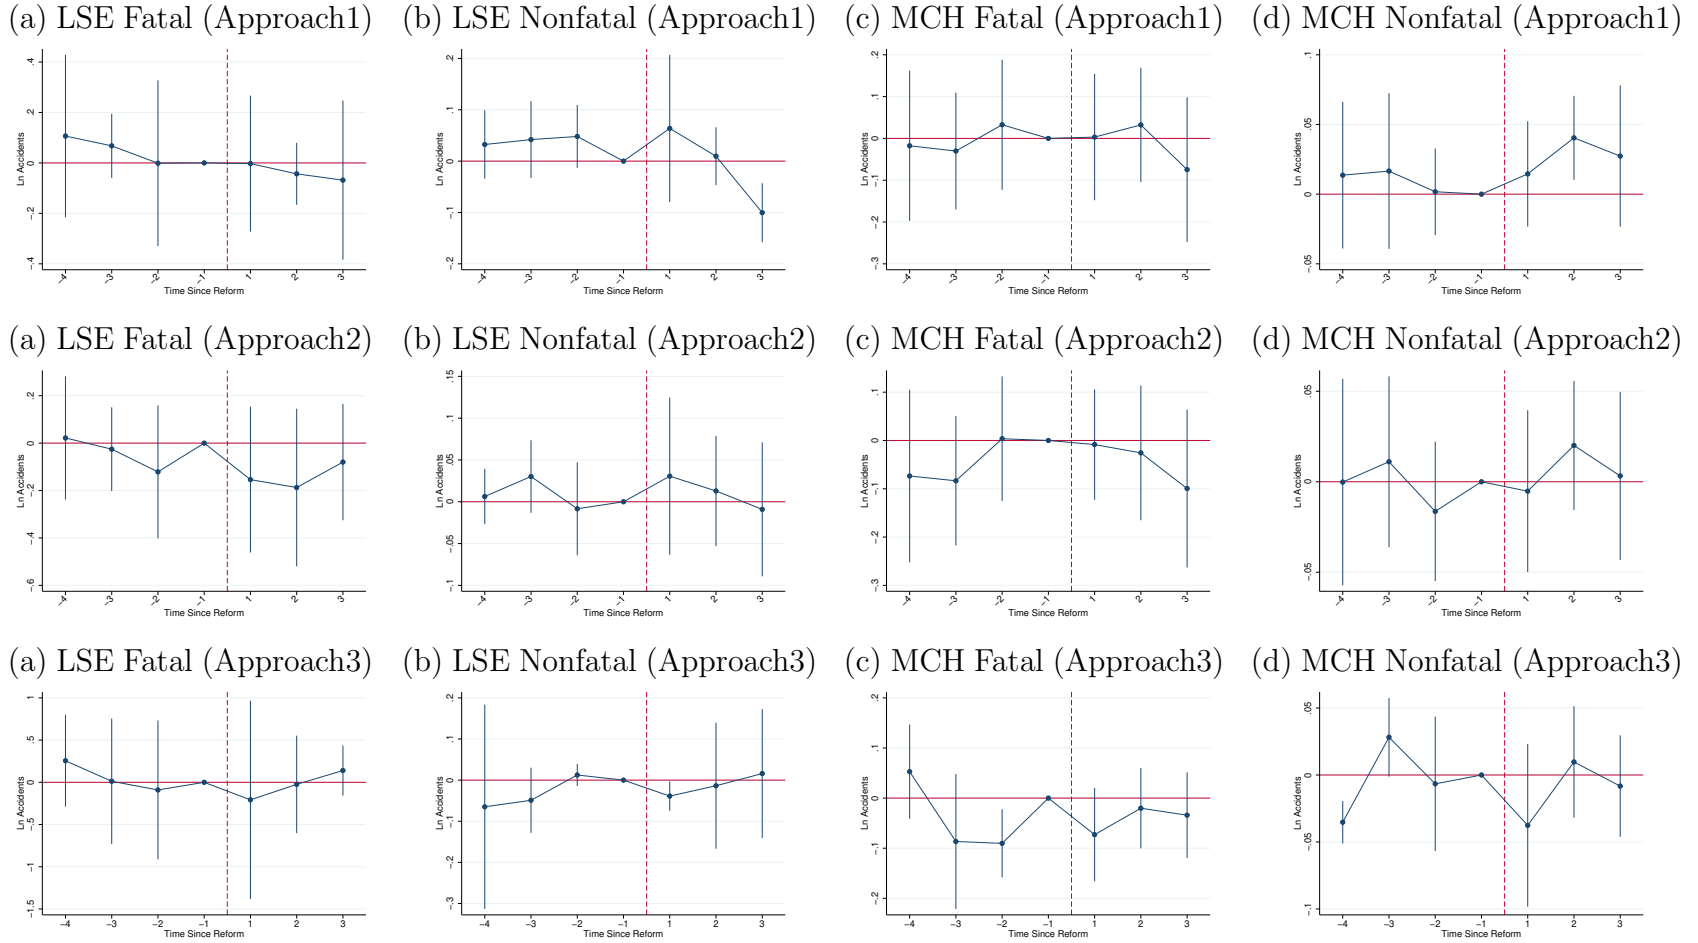
Notes: The figure reports the range of $t + 1$ event-study estimates and associated 95% confidence intervals across all 9 categories of workers' compensation reform (for both nonfatal and fatal outcomes) using Equation (1) with *Approach 1*, *2*, and *3* to estimate the immediate workplace injury moral hazard effects of changes to workers compensation generosity. The full set of event time dummies for each category of reform and each specification can be seen in Figures 1-2 and Figures 6-9 of the Appendix. Because the confidence intervals span such a wide length as to make the scale of the remaining estimates too compressed, the LSE reform analyzed under *Approach 3* and the fatal outcome is not included in Figure 3. The $t + 1$ point estimate for this specification is -0.207 and the 95% confidence interval is (-1.381, 0.966).

Figure 4: Direct Compensation Increase versus Decrease



Notes: The figure reports the treatment estimates and 95% confidence intervals from estimation of Equation (1) under *Approach 1* (row 1), *Approach 2* (row 2), and, *Approach 3* (row 3) for the CI (columns 1 and 2) and CL (columns 3 and 4) categories of workers' compensation reform, with fatal (columns 1 and 3) and nonfatal (columns 2 and 4) injuries (in logs) as separate outcomes. The red dashed vertical line represents the onset of the workers' compensation reform. The CI/CL category of reform involves reforms that directly increase/lower the dollar benefits a covered worker receives from the workers' compensation program. Further details of the category definitions and outcome measure are found in Section 2. Details on *Approach 1, 2* and *3* can be found in the notes to Figure 1, 6, and 8, and, in Section 3.

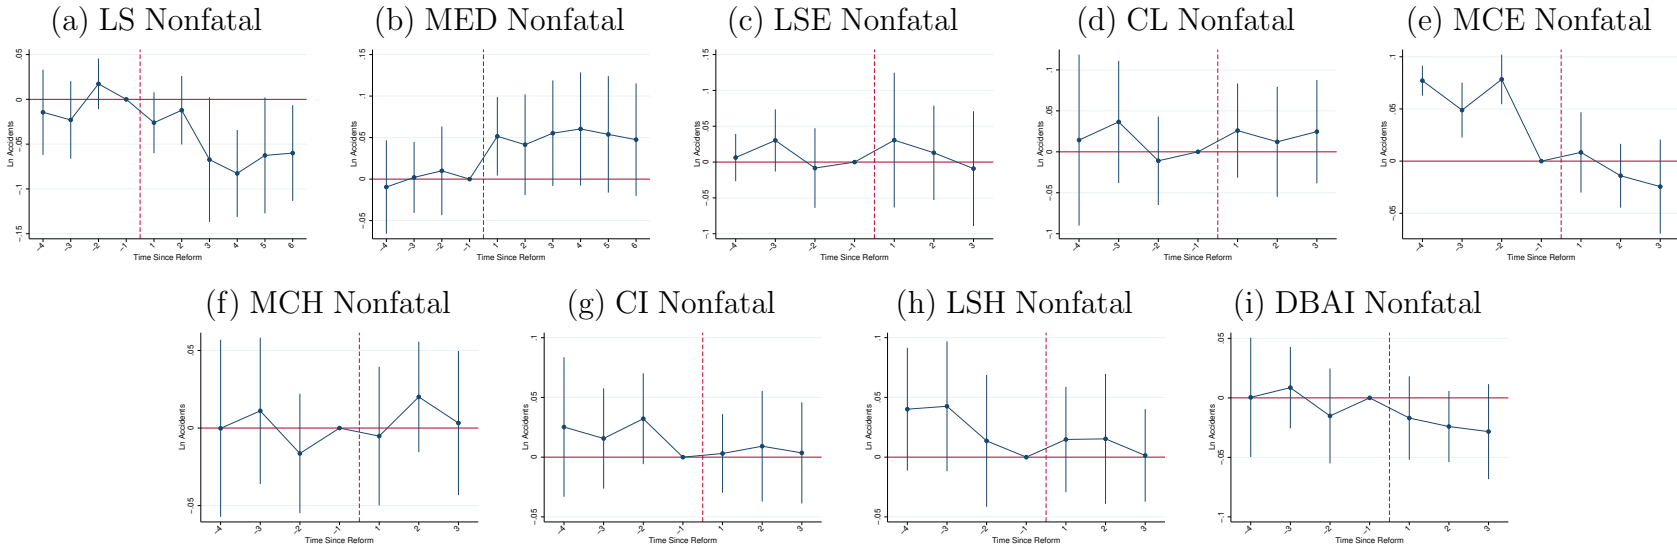
Figure 5: Reforms that (1) Increase Chance of Getting Benefits, or, (2) Reduce In-Kind (Medical) Compensation



Notes: The figure reports the treatment estimates and 95% confidence intervals from estimation of Equation (1) under *Approach 1* (row 1), *Approach 2* (row 2), and *Approach 3* (row 3) for the LSE (columns 1 and 2) and MCH (columns 3 and 4) categories of workers' compensation reform, with fatal (columns 1 and 3) and nonfatal (columns 2 and 4) injuries (in logs) as separate outcomes. The red dashed vertical line represents the onset of the workers' compensation reform. The LSE category of reform includes legal standards reforms that indirectly increase the chance a worker is deemed eligible for workers' compensation in a variety of ways. The MCH category of reform involves various reforms that effectively lower the value of in-kind workers' compensation benefits by reducing quality or access to medical care. Further details of the category definitions and outcome measure are found in Section 2. Details on *Approach 1, 2* and *3* can be found in the notes to Figure 1, 6, and 8, and, in Section 3.

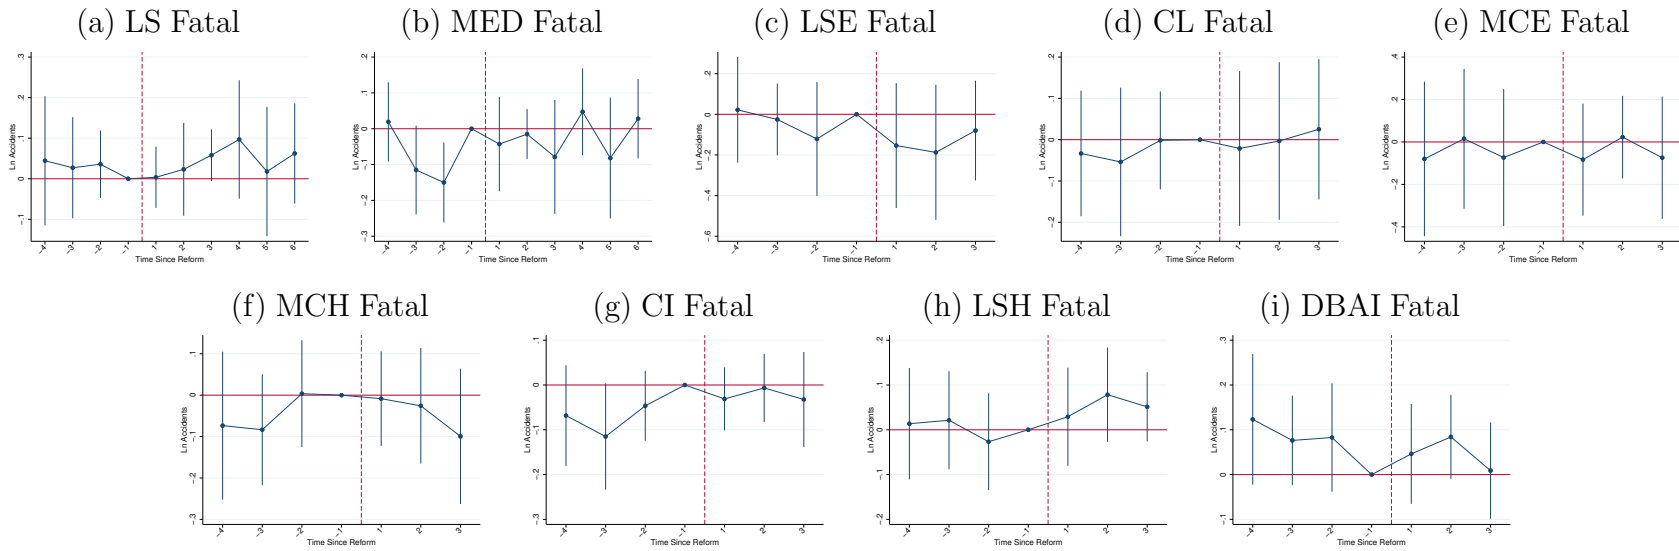
APPENDIX FOR ONLINE PUBLICATION

Figure 6: Moral Hazard and Workers' Compensation Reforms: Approach 2 Nonfatal Injuries



Notes: The figure reports the treatment estimates and 95% confidence intervals from separate estimation of Equation (1) under *Approach 2* for all nine categories of workers' compensation reform (with nonfatal injuries as the outcome, in logs). Details of the category definitions and outcome measure are in Section 2. The red dashed vertical line represents the onset of the workers' compensation reform. Unlike *Approach 1*, *Approach 2* includes all states in the analysis, including those with staggered policy reforms that take place over multiple years, but assigns treatment status of value 0 before the first reform in a category and of value 1 for the year of the first reform and all subsequent years (regardless of whether additional later reforms in the category take place in the state). Note that for the LS and MCE reforms the estimates are identical to those using *Approach 1* reported in Figure 1 as there are no staggered reforms for this category of reform.

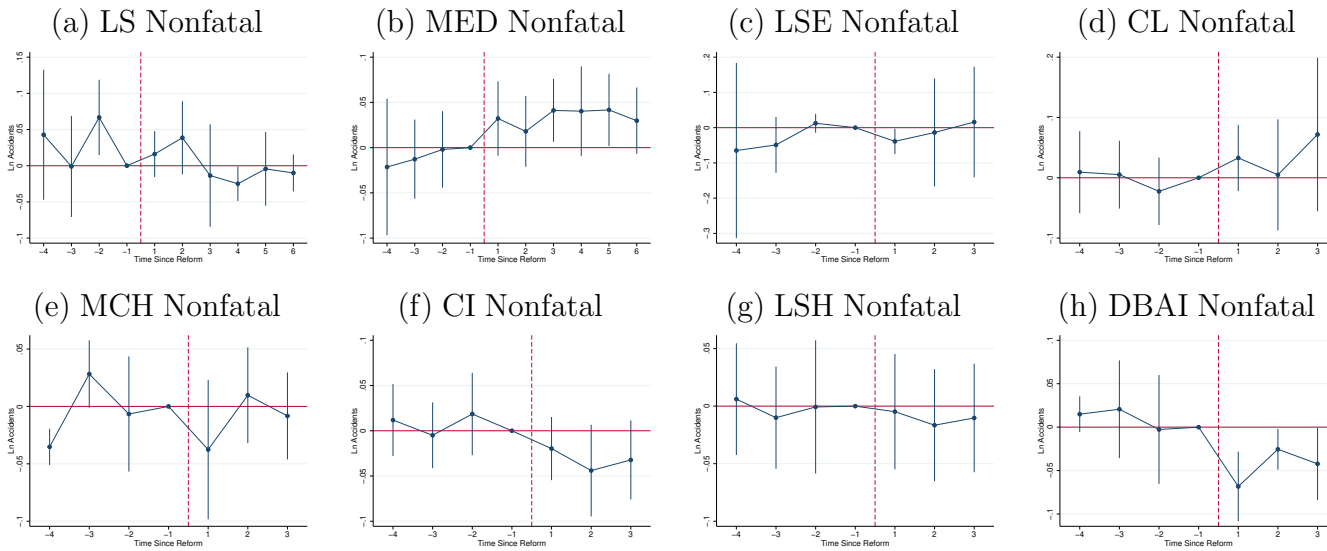
Figure 7: Moral Hazard and Workers' Compensation Reforms: Approach 2 Fatal Injuries



22

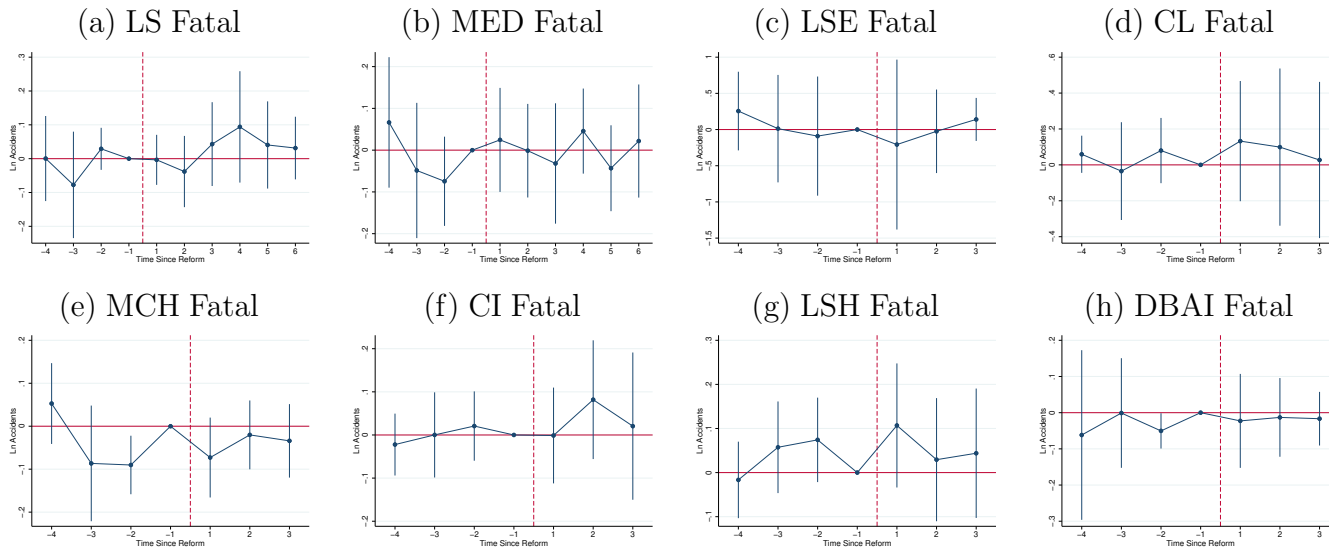
Notes: The figure reports the treatment estimates and 95% confidence intervals from separate estimation of Equation (1) under *Approach 2* for all nine categories of workers' compensation reform (with fatal injuries as the outcome, in logs). Details of the category definitions and outcome measure are in Section 2. Unlike *Approach 1*, *Approach 2* includes all states in the analysis, including those with staggered policy reforms that take place over multiple years, but assigns treatment status of value 0 before the first reform in a category and of value 1 for the year of the first reform and all subsequent years (regardless of whether additional later reforms in the category take place in the state). Note that for the LS and MCE reforms the estimates are identical to those using *Approach 1* reported in Figure 2 as there are no staggered reforms for this category of reform.

Figure 8: Moral Hazard and Workers' Compensation Reforms: Approach 3 Nonfatal Injuries



Notes: The figure reports the treatment estimates and 95% confidence intervals from separate estimation of Equation (1) under *Approach 3* for all nine categories of workers' compensation reform (with nonfatal injuries as the outcome, in logs). Details of the category definitions and outcome measure are in Section 2. The red dashed vertical line represents the onset of the workers' compensation reform. *Approach 3* addresses issues with staggered event-study designs raised in [Abraham and Sun \(2020\)](#) following the approach of [Cengiz et al. \(2019\)](#) using a stacked-by-event approach described further in the text in Section Results. The specifications run on this stacked data set modify Equation (1) by clustering standard errors at the state by event-specific-data set level and adding an additional control for the post-event period of any additional same-category reform whose 6-year post-treatment period takes place within the event-specific data set, accounting for the staggered nature of some reforms in a state (while dummies present in Equation (1) for years above and below the event window are necessarily excluded given the contraction of the sample to include only the event window years). The MCE reform category is absent from the figure because under *Approach 3* a single state receives MCE treatment status, because of the requirement to have the full event window in the sample period, leaving no estimate for this reform under this approach due to collinearity.

Figure 9: Moral Hazard and Workers' Compensation Reforms: Approach 3 Fatal Injuries



Notes: The figure reports the treatment estimates and 95% confidence intervals from separate estimation of Equation (1) under *Approach 3* for all nine categories of workers' compensation reform (with fatal injuries as the outcome, in logs). Details of the category definitions and outcome measure are in Section 2. The red dashed vertical line represents the onset of the workers' compensation reform. *Approach 3* addresses issues with staggered event-study designs raised in [Abraham and Sun \(2020\)](#) following the approach of [Cengiz et al. \(2019\)](#) using a stacked-by-event approach described further in the text in Section Results. The specifications run on this stacked data set modify Equation (1) by clustering standard errors at the state by event-specific-data set level and adding an additional control for the post-event period of any additional same-category reform whose 6-year post-treatment period takes place within the event-specific data set, accounting for the staggered nature of some reforms in a state (while dummies present in Equation (1) for years above and below the event window are necessarily excluded given the contraction of the sample to include only the event window years). The MCE reform category is absent from the figure because under *Approach 3* a single state receives MCE treatment status, because of the requirement to have the full event window in the sample period, leaving no estimate for this reform under this approach due to collinearity.