W.E. UPJOHN INSTITUTE
FOR EMPLOYMENT RESEARCH

2017

Three Essays on Nonwage Compensation

Matthew S. Johnson
Boston University

Three Essays on Nonwage Compensation

Matthew S. Johnson

An employment relationship consists of many dimensions other than monetary compensation. Textbook economic theory implies that employers and employees will agree on an efficient level of such nonwage compensation based on an employee's preferences and the employer's cost. At the same time, most types of nonwage compensation are set in a context of substantial regulation, legal restrictions, and other interventions. This dissertation investigates how the institutional environment—including regulation, media coverage of corporate actions, and the strength of the labor market—affects firms' decisions regarding two important types of nonwage compensation: workplace safety and health, and employment mobility.

Chapter 1 investigates how media coverage of employers caught violating workplace safety and health regulations affects future compliance and injuries. Using quasi-random variation in media coverage induced by a policy change at the Occupational Safety and Health Administration (OSHA), I find that media coverage about one employer leads to significantly higher compliance, and fewer serious injuries among other employers likely exposed to the coverage. The results are most consistent with employers acting defensively to avoid costly responses from workers, and they suggest there is information asymmetry between employers and workers regarding employers' safety and health performance.

Chapter 2 examines how workplaces respond to health and safety regulatory enforcement inspections. Using establishment-level data on injuries, illnesses, and other business outcomes, we find that randomly assigned OSHA inspections led to significantly fewer injuries and had no detectable effect on business outcomes at inspected workplaces. We then use new machine learning methods to estimate heterogeneous treatment effects of inspections, and we use the results to simulate how many additional injuries OSHA could avert if it targeted its limited inspection resources to the workplaces where they are most beneficial.

Chapter 3 investigates why employers have employees sign noncompete agreements (NCAs), which contractually limit where the employee can work in the event of a job separation. NCAs may solve hold-up problems that limit incentives to invest in transferrable assets (e.g., general human capital), but they impose costs on employees who sign them. Recent evidence that NCAs are frequently used in many traditionally lower-paying occupations has raised questions about the rationale for, and effects of, NCA use in this setting. We develop a model of how labor market conditions and wage frictions (e.g., the minimum wage) can jointly determine the decision to include an NCA in a hiring

contract, even if NCAs reduce the firm's total surplus. We find strong support for the model's predictions using a survey we conducted among employers in the high-end hair salon industry. Furthermore, we generate a test for identifying when NCAs do not maximize a firm's surplus, and we identify a subset of firms in our sample, characterized by limited access to credit, for which this is the case.

Note: Since this dissertation was accepted in August 2016, I have updated the three chapters to varying degrees (e.g., by adding new analysis) that strengthen the papers but do not change the motivations, implications, or main results. The chapter summaries below correspond to the current versions. After each summary, I briefly describe how the current versions differ from my dissertation.

Chapter 1

Regulation by Shaming: Deterrence Effects of Publicizing Violations of Workplace Safety and Health Laws

Ratings, scores, disclosure, and other means of informing a firm's stakeholders about an aspect of its quality or performance have proliferated in recent years (Dranove and Jin 2010). Such policies are guided by the basic economic insight that, when quality is imperfectly observed, providing information mitigates a moral hazard problem that distorts firms' incentives to invest in quality. Indeed, a growing empirical literature has found that providing information about quality to the public leads rated, scored, or otherwise disclosed firms to improve the quality of the attributes under scrutiny.[1]

Many sources seek to disclose information only about firms whose quality or performance is particularly low: that is, "shaming." For example, nongovernmental organizations and media outlets compile lists of firms that fail in some dimension according to objective data sources, such as "Least Green Companies in America" (*Newsweek* 2010), or the campaigns against companies that used sweatshop labor in the 1990s (Harrison and Scorse 2010). Increasingly, technology and social media have enabled customers, former workers, and other stakeholders to expose companies' actions ranging from tax avoidance (Barford and Holt 2013), high medical drug prices,[2] and sexual harassment of employees.[3] While one intent of such tactics is to pressure the entity being targeted to improve its behavior ("specific deterrence"), a broader and perhaps more important intent is to encourage improvements in quality at other entities that wish to avoid being the target of their own future negative publicity ("general deterrence"). Despite the growing prevalence of these policies, little is known about how firms respond to such information disclosure targeted only at the worst performers. Estimating their effects poses substan-

tial empirical challenges due to the purposely nonrandom selection of entities that are publicized, to the difficulty in knowing which other entities are the most likely respond to general deterrence, and to a dearth in data on outcomes typically under scrutiny.

This paper overcomes these challenges. Specifically, I investigate a policy dubbed "regulation by shaming" implemented by OSHA (Michaels 2010). In 2009 OSHA, the regulatory agency charged with setting and enforcing workplace safety and health standards in the United States, began issuing press releases about facilities found to be violating safety and health standards in a recent inspection. The policy was intended to expose egregious violators to public scrutiny and to publicize OSHA's enforcement actions. These press releases described the violations found in a recent inspection of a facility and the financial penalties levied, and they implied that the employer was exposing its workers to substantial safety and health hazards.

The initiation of OSHA's press release policy provides an ideal setting to understand the scale, scope, and persistence with which publicizing poor performance affects firms' behavior. First, OSHA used a cutoff rule whereby it issued a press release about a facility's violations if the financial penalties it levied at a recent inspection were above a certain threshold. This rule provides quasi-random variation in publicity among otherwise similar facilities that lends itself to a Regression Discontinuity design. Second, OSHA distributed these press releases to local newspapers and industry trade publications, meaning that other facilities in close geographic proximity and in the same industry were most likely to be exposed to publicity of a press release. The policy was only announced internally within OSHA and not made known to the general public. As a result, it led to a sharp and unexpected increase in media coverage of OSHA violations, and meant that a well-defined set of facilities were made aware of this new threat of media coverage. Third, OSHA routinely inspects a broad set of workplaces to detect health and safety violations and collects the results in an internal database, providing a timely and systematic data source to measure facility improvements in response to press releases.

Understanding the extent to which publicity like this affects workplace safety and health is not only useful to understand how firms respond to targeted information disclosure, but it is also an important question for public policy; although U.S. workplace injury rates have declined in recent decades, they continue to have substantial welfare costs, with one recent study estimating that they cost the United States $250 billion per year (Leigh 2011).

I find that press releases revealing OSHA noncompliance lead to substantial improvements in workplace safety and health. A press release about one facility leads to 1.7 fewer violations at other facilities in the same sector within a 5 kilometer radius ("peer facilities"), a decrease of 73 percent. To put the magnitude of this deterrence effect in perspective,

an OSHA inspection has been estimated to lead to between 28 and 48 percent fewer violations at later inspections of the same facility (Ko, Mendeloff, and Gray 2012). Thus, this paper's estimates imply that publicizing violations committed by one facility leads other peer facilities to improve compliance by two to three times as much as if OSHA inspected each of these facilities instead. Given that inspections are relatively costly and that OSHA's budget constraints—like those in many other regulatory agencies—dictate that it can inspect only a small subset of regulated workplaces, this publicity appears to be a highly effective policy tool to improve workplace safety.

Furthermore, using the occurrence of OSHA inspections triggered by a fatal, or otherwise very serious, workplace injury, I find that press releases lead not only to improved compliance with OSHA regulations, but also to fewer injuries. An inspection with penalties just above the press release cutoff leads to significantly fewer inspections triggered by a serious accident among other peer facilities. The magnitude of the effect, as with compliance, is substantial.

The paper then tests for mechanisms through which OSHA's press releases lead facilities to improve their safety and health performance. One theory is that facilities improve compliance to avoid costly responses from stakeholders, especially workers. Workers who have more bargaining power may have more scope to leverage a press release to demand safer working conditions, or a larger compensating differential for job hazards, from an employer. Drawing from literature on how the presence of labor unions affects workers' bargaining power (both at unionized and nonunionized workplaces), I measure workers' bargaining power using two proxies for the strength of labor unions: whether a facility is in a right-to-work state, and the baseline unionization rate of a facility's county. Using either measure, facilities in areas where unions are strong improve compliance by a substantial amount following a press release about a peer (regardless of their own unionization); those in areas where unions are relatively weak display no improvement. In other words, press releases lead to improvements in safety and health conditions only when workers are most likely to be able to use information about an employer's safety record to respond in a costly way.

This paper makes multiple contributions. First, it provides a novel contribution to a literature on the disciplinary effects of information provision. While a growing body of work (such as those papers cited in note 1) has investigated the extent to, and conditions under, which information disclosure leads firms to improve their performance or quality, this paper is one of the first to identify how providing information about some targeted firms can have broader effects on the behavior of other firms. A separate literature has explored the effect of "shaming" in other domains, such as public release of criminal records (Lee 2013) and tax delinquency (Perez-Truglia and Troaino 2016). In the realm of politics, media coverage has been shown to affect politicians' incen-

tives to engage in malfeasant behavior (Snyder and Strömberg 2010). This paper builds on these literatures by exploring how shaming—and targeted information disclosure in general—affects firm behavior in a regulatory environment.

Second, this paper contributes to the literature on the determinants of regulatory compliance in firms. Many prior studies have investigated the specific deterrence effects of OSHA inspections on future compliance of inspected facilities (Gray and Jones 1991; Ko, Mendeloff, and Gray 2010; Weil 1996), as well as effects in other regulatory domains such as by the Environmental Protection Agency (see Alm and Shimshack [2014] for an overview). At least in the environmental domain, the consensus in this literature seems to be that "rigorous monitoring and enforcement remains the number one motivator for many facilities' environmental compliance decisions" (Gray and Shimshack 2011, p. 1). This paper's findings suggest the media and "shaming" have been overlooked as powerful forces governing firms' compliance decisions, at least for safety and health.

Note: The version of this chapter submitted for my dissertation used shared zip code, rather than geographic distance, as a measure of geographic proximity when estimating general deterrence effects of press releases (the substance of the results is unchanged). That version also did not analyze the effects of press releases on occupational injuries, and it also did not analyze how labor union strength moderates the effect of press releases on compliance.

Chapter 2

Improving Regulatory Effectiveness through Better Targeting: Evidence from OSHA

(with David I. Levine and Michael W. Toffel)

Government agencies spend billions of dollars to send inspectors to assess compliance with regulations governing, for example, worker safety, environmental protection, and product safety.[4] Budget constraints require almost all regulators to monitor a small subset of regulated units. For example, in 2016 OSHA and its state counterparts inspected less than 1 percent of the 8 million workplaces they regulated. (U.S. Department of Labor 2017b). Other agencies face similar discrepancies between the scale of their resources and scope of their jurisdiction.[5]

Such constraints require regulators to make difficult choices about how to prioritize their inspections. Agencies prioritize inspections based on various statutory requirements, heuristics, and algorithms. For example, OSHA allocates most of its inspections to facilities that recently experienced serious accidents, had employee complaints, or dangerous workplaces in high-hazard industries.[6]

Assessing agencies' effectiveness is a challenging task. First, it is difficult to know if agencies' inspections are

furthering their goals at all; agencies choose which establishments to inspect, which makes it hard to find a credible comparison group of uninspected establishments for evaluation purposes. Second, it is even more difficult to know if agencies are maximizing their effectiveness. Even if one could credibly evaluate the average effects of inspections as agencies currently target them, conventional econometric methods do not tell us if current targeting is optimal. The lack of robust evidence of inspection effectiveness has left government agencies susceptible to criticism that their efforts waste taxpayer dollars, that they target establishments to promote politicians' agendas (Weisman and Wald 2013), or that they serve the interests of those they regulate (Stigler 1971).

In this paper, we combine a large randomized experiment with machine learning methods, which yields an approach to assess the extent to which regulatory agencies are maximizing the benefits of their limited inspection resources. We do so by first identifying establishments where inspections do the most to accomplish the agency's objectives, and we then simulate outcomes if the agency adopted various policies that targeted those establishments.

We conduct our analysis in the context of a major inspection program of OSHA, which has been highly controversial since its creation in 1970. While its supporters argue the agency saves lives at little to no cost to employers, critics charge that its regulations add costs but "don't add value to safety in the workforce" (Heitkamp 2016), or that its penalties for noncompliance are too low for OSHA to have any effect at all (Bartel and Thomas 1985).

We first evaluate the extent to which a subset of inspections that OSHA randomly assigned under its Site-Specific Targeting (SST) program affected the injury rates of inspected workplaces. SST, one of OSHA's largest inspection programs from 2001 to 2010, prioritized for inspection establishments that had experienced high injury rates two years earlier, a rule intended to allow "the most effective use of OSHA's limited resources" (U.S. Department of Labor 2009). By focusing on random assignment, our estimates are free of selection bias that plague most comparisons of inspected and uninspected establishments.[7] The roughly 16,000 establishments at risk of randomized SST inspection over this period employed nearly 2.5 million workers, making this the largest group-level randomized control trial of which we are aware.

On average, we find that randomly assigned SST inspections reduced serious injuries (those leading to days away from work [DAFW]) by 9 percent over the five years following the inspection. This equates to 2.4 fewer DAFW-injuries over the five years following inspection, amounting to a social benefit of $105,000 per inspection, based on a prior estimate that DAFW-injuries typically cost $44,000 (Waehrer et al. 2007). We find no evidence that these benefits of inspections came with a cost to business outcomes such as establishment survival, employment, sales, or credit rating.

We then investigate the potential benefits of alternative targeting policies that OSHA could adopt. Conducting this analysis is challenged by the fact that a particular establishment's treatment effect of an inspection on its number of injuries—that is, the difference in the number of injuries it would experience if had been inspected versus had it not been inspected—is unobservable. When there are many factors that could moderate an inspection's effect, conventional methods to estimate such heterogeneous treatment effects (such as adding interaction terms) risk estimating spurious interactions that have poor out-of-sample predictive power.

To overcome this challenge, we employ a machine-learning approach that yields data-driven estimates of each establishment's treatment effect based on its baseline characteristics called Targeted Maximum Likelihood Estimation (TMLE) (van der Laan and Rose 2011). We first obtain an initial estimate of each establishment's outcome if inspected and outcome if not inspected, using a procedure called "super learner" (van der Laan, Polley, and Hubbard 2007). Super learner is an ensemble machine-learning prediction technique that uses cross-validation to create an optimal linear combination of machine-learning prediction algorithms. TMLE then fluctuates this initial estimate to minimize bias of the estimated average treatment effect. As a result, TMLE provides estimates to overcome the fundamental problem of causal inference (Holland 1986), estimating the counterfactual number of injuries if inspected among those that were not inspected, and number of injuries if not inspected among those that were inspected.

We estimate substantial heterogeneity in the degree to which SST inspection affects serious injuries. Armed with these estimates, we examine how different targeting policies could improve OSHA's effectiveness. For example, if OSHA reallocated all its SST inspections to target those establishments with the largest predicted treatment effects each year, which we consider our "benchmark" policy, it could have averted 58 percent more injuries at inspected establishments than the historical rule.

While this benchmark policy could avert many more injuries at inspected establishments, elements of it may be undesirable. For example, shifting the allocation of inspections may increase OSHA's cost of conducting inspections, requiring extra resources OSHA may not have. Furthermore, OSHA may want to preserve a degree of randomization to maintain effects that the threat of inspection has on uninspected establishments (Cohen 2000; Shimshack and Ward 2005). When we examine a policy that maintains OSHA's historical overall cost of conducting inspections, and which preserves a degree of randomization, we find it still increases the injuries averted among inspected establishments by 42 percent relative to the historical policy.

Another potential concern with our benchmark policy is that a government agency might be leery of targeting inspections based on the output of a "black box" machine-learning algorithm like TMLE. We thus consider policies in which OSHA bases its targeting policy on two transparent measures of hazardousness: predicted number of DAFW injuries or predicted noncompliance with OSHA regulations. We estimate these predicted values using a Lasso regression, a standard (and simpler) machine learning tool for prediction. OSHA would still avert 70 percent as many injuries if targeting based on the predicted number of DAFW injuries as it would by targeting based on the largest treatment effects. Targeting based on predicted noncompliance, in contrast, only increased averted injuries slightly compared to the historical rule.

In sum, considering a policy that accounts for each of these potential concerns of our benchmark policy, our estimates imply that targeting OSHA's SST inspections differently could have created social value of roughly $220 million over a decade.

This paper makes multiple contributions. We add to a large literature examining the effects of OSHA inspections on injuries, which to date has yielded a wide range of estimates. While many early studies find that OSHA inspections have little or no correlation with subsequent workplace injury rates (Ruser and Smith 1991; Smith 1979b; Viscusi 1986), others find that OSHA inspections are associated with a decline in injury rates (Gray and Scholz 1993; Haviland et al. 2012; Levine, Toffel, and Johnson 2012). We depart from this literature in two important ways. First, by evaluating the average effects of a subset of OSHA inspections that were randomly assigned, our approach yields credible causal estimates for the subset of inspections we study. Second, we go beyond evaluating historical effects of inspections to examine the implications of alternative inspection policies, allowing us to determine the extent to which OSHA could reallocate inspections to avert more injuries.

Note: The version of this chapter submitted for my dissertation did not include the results that use TMLE to estimate heterogeneous treatment effects of inspections. Rather, that version used a different machine-learning procedure (Lasso) to estimate establishments' predicted injuries and predicted noncompliance, and analyzed the extent to which these measures accounted for heterogeneity in the effectiveness of inspections.

Chapter 3

Cutting Out the Competition: Labor Market Conditions and the Use of Noncompete Agreements

(with Michael Lipsitz)

When a new worker receives his or her employment contract, it may include a noncompete agreement (NCA), which

contractually limits the worker's ability to enter a professional position in competition with his or her employer in the event of a job separation. Economic theory of the hold-up problem (Grossman and Hart 1986) suggests that NCAs can potentially enhance efficiency by aligning incentives to invest in various assets, such as general human capital training, trade secrets, or client lists. At the same time, NCAs may also impose significant costs on workers by limiting their ability to pursue outside employment opportunities.

Recent evidence suggests that our understanding of the reasons behind—and implications of—NCA use remains incomplete. For one, while NCAs are most prevalent in higher-skill, knowledge-intensive industries and occupations, they are also frequently used in many traditionally lower-paying occupations (Starr, Bishara, and Prescott 2015), even among fast food workers (Irwin 2014), leading some to question what benefit NCAs could be bringing to these employment relationships. Furthermore, the use of NCAs has been growing in recent years (Greenhouse 2014), which, absent corresponding changes in the importance of training, trade secrets, client lists, or other facets of production technology, is difficult to rationalize with the theory of the hold-up problem alone.

These developments have captured policymakers' attention: in Congress, the Mobility and Opportunity for Vulnerable Employees (MOVE) Act,[8] introduced on June 4, 2015, would prohibit NCAs for workers earning less than $15 per hour, and bills with similar intents have been introduced by several state legislatures.[9] Despite this policy interest, little is known about the efficiency of NCAs in this context, let alone the rationale for their use in the first place.

In this paper, we show that NCAs arise when employers and employees are limited in their ability to transfer utility via the wage. When the market-clearing wage is constrained, NCAs may be used as a tool to transfer additional surplus to the employer, even if NCAs do not maximize an employer and employee's joint surplus. In fact, such constraints on wages will only affect NCA use if NCAs are not surplus-maximizing. Thus, we provide a simple method that generates a sufficient condition to determine when NCAs do not maximize surplus in firms: if a change in the bindingness of a wage constraint affects NCA use, NCAs cause a joint surplus loss (relative to a contract without an NCA) for at least a subset of firms. We implement this test using data from a survey we conducted of employers in the hair salon industry.

We start with a simple, perfectly competitive model of the labor market in which NCAs provide a benefit to the employer and impose a cost on the employee. If utility is fully transferable between the employer and employee via the wage, NCAs will be used only when the firm's net benefit of NCA use is positive: when NCAs maximize joint surplus. However, when utility transferability via the wage is limited, the terms of trade in the labor market may dictate that NCAs

are used as a tool to transfer surplus from the employee to the employer, even if NCAs do not maximize firms' surplus. NCA use will therefore increase when the terms of trade become more favorable to the employer or when transferability of utility decreases.

While firms plausibly have many possible nonpecuniary instruments for surplus transfer at their disposal, there are reasons to believe NCAs would be a prominent one. First, prior research has identified clear ways NCAs benefit employers. For example, because NCAs lead to longer worker tenure (Starr, Bishara, and Prescott 2015), they can reduce employee replacement costs, which can be substantial (Dube, Freeman, and Reich 2010). Additionally, in industries for which production depends on transferable assets such as client lists and general human capital, ensuring retention of these assets is extremely valuable. Second, relative to other nonpecuniary job attributes such as provision of health care, it is relatively easy for employers to switch in and out of using NCAs; NCA use requires straightforward changes to employment contracts, whereas adjustments to health insurance benefits requires much more coordination, time, and resources.

To test the empirical predictions of our model, we surveyed owners of independent hair salons in April 2015 via the Professional Beauty Association, a trade association for the industry. The benefits of NCAs are clear in this setting due to the importance of client attraction and retention in production, and the prevalence of on-the-job training. At the same time, due to state-level occupational licensing laws that make mobility costly, the costs of NCAs to workers are also potentially high. We find that NCAs are widely used: 30 percent of our sample had their most recently hired stylist sign an NCA, and 39 percent have had at least one stylist sign an NCA in the past.

Taking the model to the data, we find strong empirical support that limitations on transferability of utility via the wage affect NCA use. First, we test the prediction that NCA use is higher when the terms of trade in the labor market are more favorable for the employer. We find that outward shifts in labor supply (proxied by the number of applicants an owner received for her most recent vacancy), and increases to the local unemployment rate—both of which will be associated with a lower market-clearing wage—are related to higher NCA use. We estimate that one additional applicant for a vacancy leads to a 4 percent increase in the probability that the hired worker signed an NCA. We also find that salons in counties that experienced higher increases in the unemployment rate between 2006 and 2012 (roughly the period spanning the Great Recession) were more likely to have their most recently hired worker sign an NCA.

Second, we find that increases in the minimum wage, which limit transferability of utility, also have a strong effect on NCA use. Owners in states with a higher minimum wage for tipped employees are more likely to use NCAs. Because

cross-sectional variation in the minimum wage might be driven by other, unobservable differences across states, we separately estimate the effect of the minimum wage on NCA use for salons that hire workers as employees versus those that hire as independent contractors. The latter group is not covered by the Fair Labor Standards Act and thus acts as a "placebo group" for the minimum wage. The effect only holds for the employment-based salons in our sample and is small and statistically insignificant for contractor-based salons. Among employment-based salons, a $1.00 increase in the minimum wage is associated with an 8 percentage point increase in the probability that an owner has used a NCA in an employment contract.

Combined with the implications of our model, these results imply that NCAs do not maximize surplus for at least some firms in our sample. However, NCAs may still be surplus-maximizing contracts for a subset of firms. For example, if the benefits of NCAs are heterogeneous across employers, NCAs may maximize surplus for those firms with the highest benefit.

To investigate the extent of variation in the benefit of NCAs in our sample, we first corroborate existing evidence that one benefit of NCAs is to enhance incentives for employers to invest in production assets, and we then utilize a measure of employers' ability to invest in production assets originating in the corporate finance literature: access to a line of credit with a bank (Sufi 2009). We find evidence consistent with NCAs being surplus-maximizing for employers with high capacity for investment, but not for those with low capacity. Employers with high capacity use NCAs at a high rate, regardless of whether the market-clearing wage is likely constrained. On the other hand, employers with low capacity are highly unlikely to use NCAs in an unconstrained environment (proxied by a low minimum wage, low level of labor supply, or low local unemployment rate), but this likelihood increases as the wage becomes constrained.

Overall, these results highlight a potential explanation for the growth of NCAs among lower-wage occupations and industries in recent years. Between 2007 and 2009, the federal minimum wage rose from $5.15 per hour to $7.25 per hour, and several states have increased their minimum wage in more recent years. Furthermore, in the wake of the Great Recession, there is a consensus that the labor market has deteriorated dramatically, especially for low-wage workers. Our results imply that employers leveraged this weak labor market to use NCAs as a tool to extract additional surplus from workers.

Thus, our results yield nuanced implications for policy. On the one hand, our survey results suggest that NCAs may arise as a tool for employers to extract surplus from workers in weak labor markets, reducing joint surplus and leaving workers worse off. At the same time, even within a narrowly defined industry, we find that NCAs do not maximize surplus for some firms but do for others. This finding

stresses the need for future research to further investigate the benefits NCAs provide to firms, which can aid policymakers by pinpointing where NCAs are most likely to be surplus-diminishing.

This paper contributes to a growing literature examining the rationale for NCAs and the effects of their use. Using variation in the enforceability of NCAs across states, an increase in NCA enforceability has been found to increase firm-sponsored training (Starr 2017) and decrease employee mobility (Marx, Strumsky, and Fleming 2009). We add to the literature by empirically demonstrating how forces external to the firm influence the decision to use NCAs in the first place, and by providing a method to identify the presence of NCAs that do not maximize a firm's joint surplus. We also conduct the first survey on NCA use with employer information, allowing us to explore determinants and effects of NCA use not available through worker surveys or variation in enforceability.

Note: The version of this chapter submitted for my dissertation did not include the analysis of how changes to the local unemployment rate affect NCA use.

Notes

1. Some examples are restaurant hygiene report cards (Jin and Leslie 2003), disclosure of drinking-water quality (Bennear and Olmstead 2008), and environmental ratings (Chatterji and Toffel 2010). See Dranove and Jin (2010) for an overview of the literature.
2. "Social Media Shaming: Can Outrage Be Effective?" Knowledge@Wharton, November 20, 2015. http://knowledge.wharton.upenn.edu/article/social-media-shaming-can-outrage-be-effective/ (accessed October 25, 2017).
3. "Uber Orders Investigation into Sexual Harassment Claims," February 20, 2017. NPR, *All Things Considered*. http://www.npr.org/2017/02/20/516292319/uber-orders-investigation-into-sexual-harassment-claims (accessed October 25, 2017).
4. Three examples are illustrative. OSHA and its state counterparts spent more than $300 million on enforcement in 2016 (U.S. Department of Labor 2017a). Second, the U.S. Environmental Protection Agency spends more than $600 million per year on enforcement, which does not include inspections performed by state-level environmental agencies, which actually led the efforts to enforce environmental regulation in the United States (Shimshack 2014). Third, the U.S. Food and Drug Administration (FDA) allocated roughly two-thirds of its $1 billion Office of Regulatory Affairs budget to inspections in fiscal year 2015 (U.S. Food and Drug Administration 2016).
5. For example, FDA is charged with ensuring imported products meet FDA standards but in 2010 only physically inspected 2 percent of imported food products (U.S. Department of Health and Human Services 2010).
6. "OSHA Fact Sheet: OSHA Inspections," https://www.osha.gov/OshDoc/data_General_Facts/factsheet-inspections.pdf (accessed February 2017).
7. For example, because many OSHA inspections target establishments with recent accidents or complaints, inspected establishments likely have systematically different character-

istics (both observable and unobservable) than noninspected establishments. Furthermore, establishments experiencing high injury rates in one year (thus triggering an OSHA inspection) may experience fewer injuries the following year simply due to regression to the mean, in which case OSHA inspections would be correlated with lower injury rates without actually causing them.

8. Senate Bill S. 1504. Full text available at https://www.govtrack.us/congress/bills/114/s1504.

9. Some examples include Washington (HB 1926; introduced February 2, 2015), Utah (HB 251; introduced February 1, 2016), and Illinois (Illinois Freedom to Work Act; goes into effect January 1, 2017).

References

Alm, James, and Jay Shimshack. 2014. "Environmental Enforcement and Compliance: Lessons from Pollution, Safety, and Tax Settings." Foundations and Trends in Microeconomics 10(4): 209–274.

Barford, Vanessa, and Gerry Holt. 2013. "Google, Amazon, Starbucks: The Rise of 'Tax Shaming.'" BBC Magazine, May 21. http://www.bbc.com/news/magazine-20560359 (accessed October 25, 2017).

Bartel, Ann P., and Lacey Glenn Thomas. 1985. "Direct and Indirect Effects of Regulation: A New Look at OSHA's Impact." Journal of Law and Economics 28(1): 1–25.

Bennear, Lori S., and Sheila M. Olmstead. 2008. "The Impacts of the 'Right to Know': Information Disclosure and the Violation of Drinking Water Standards." Journal of Environmental Economics and Management 56(2): 117–130.

Chatterji, Aaron K., and Michael W. Toffel. 2010. "How Firms Respond to Being Rated." Strategic Management Journal 31(9): 917–945.

Dranove, David, and Ginger Zhe Jin. 2010. "Quality Disclosure and Certification: Theory and Practice." Journal of Economic Literature 48(4): 935–963.

Dube, Arindrajit, Eric Freeman, and Michael Reich. 2010. "Employee Replacement Costs." Working Paper No. 201-10. Berkeley: Institute for Research on Labor and Employment.

Gray, Wayne B., and Carol Adaire Jones. 1991. "Longitudinal Patterns of Compliance with Occupational Safety and Health Administration Health and Safety Regulations in the Manufacturing Sector." Journal of Human Resources 26(4): 623–653.

Gray, Wayne B., and John T. Scholz. 1993. "Does Regulatory Enforcement Work? A Panel Analysis of OSHA Enforcement." Law and Society Review 27(1): 177–214.

Greenhouse, Steven. 2014. "Noncompete Clauses Increasingly Pop Up in Array of Jobs." New York Times, June 8. https://www.nytimes.com/2014/06/09/business/noncompete-clauses-increasingly-pop-up-in-array-of-jobs.html (accessed October 26, 2017).

Grossman, Sanford J., and Oliver D. Hart. 1986. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration." Journal of Political Economy 94(4): 691–719.

Harrison, Ann, and Jason Scorse. 2010. "Multinationals and Anti-Sweatshop Activism." American Economic Review 100(1): 247–273.

Haviland, Amelia M., Rachel M. Burns, Wayne B. Gray, Teague Ruder, and John Mendeloff. 2012. "A New Estimate of the Impact of OSHA Inspections on Manufacturing Injury Rates, 1998–2005." American Journal of Industrial Medicine 55(11): 964–975.

Heitkamp, Heidi. 2016. Hearing with Senate Homeland Security and Governmental Affairs Committee, February 11. http://www.safetyandhealthmagazine.com/articles/13662-compliance-assistance-not-fines-should-be-priority-senators-tell-osha (accessed October 31, 2017).

Irwin, Neil. "When the Guy Making Your Sandwich Has a Noncompete Clause." New York Times, October 14. https://www.nytimes.com/2014/10/15/upshot/when-the-guy-making-your-sandwich-has-a-noncompete-clause.html (accessed October 26, 2017).

Jin, Ginger Zhe, and Philip Leslie. 2003. "The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Cards." Quarterly Journal of Economics 118(2): 409–451.

Ko, Kilkon, John Mendeloff, and Wayne B. Gray. 2010. "The Role of Inspection Sequence in Compliance with the U.S. Occupational Safety and Health Administration's (OSHA) Standards: Interpretations and Implications." Regulation & Governance 4(1): 48–70.

Levine, David I., Michael W. Toffel, and Matthew S. Johnson. 2012. "Randomized Government Safety Inspections Reduce Worker Injuries with No Detectable Job Loss." Science 336(6083): 907–911.

Leigh, J. Paul. 2011. "Economic Burden of Occupational Injury and Illness in the United States." Milbank Quarterly 89(4): 728–772.

Marx, Matt, Deborah Anne Strumsky, and Lee Fleming. 2009. "Mobility, Skills, and the Michigan Noncompete Experiment." Management Science 55(6): 875–889.

Michaels, David. 2010. "OSHA at Forty: New Challenges and New Directions." Letter to OSHA staff, July 19.

Newsweek. 2010. "Least Green Companies in America: Photos." October 16. http://www.newsweek.com/least-green-companies-america-photos-68107 (accessed October 25, 2017).

Perez-Truglia, Ricardo, and Ugo Troiano. 2016. "Shaming Tax Delinquents: Theory and Evidence from a Field Experiment in the United States." NBER Working Paper No. 21264. Cambridge, MA: National Bureau of Economic Research.

Ruser, John W., and Robert S. Smith. 1991. "Reestimating Osha's Effects: Have the Data Changed?" Journal of Human Resources 26(2): 212–235.