



The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.

Do Safety Inspections Improve Safety? Evidence from the Roadside Inspection Program for Commercial Vehicles

Yuanning Liang*

January 10, 2023

Abstract

Regulatory efforts to reduce public hazards could be undermined when agents respond strategically to enforcement. This paper quantifies changes in road safety from a nationwide truck inspection regulation using a comprehensive data set on inspections and accidents. I show that there is a sharp *increase* in the truck's accident rate immediately following an inspection, which aggregates to 1,803 additional accidents per year. These accidents are caused by reduced caution in driving from recently inspected drivers, such as speeding and less maintenance, as regulators rarely conduct repeated inspections. Policy comparisons show that less predictable inspection schedules and practices could reduce accidents.

Keywords: Enforcement and compliance; regulatory design; road safety; trucking industry

*Liang: School of Economics, Peking University, Beijing, China (email: ynliang@pku.edu.cn). I am very grateful to the suggestions and guidance by my committee members, Shanjun Li, Catherine Kling, Ivan Rudik, Crocker Liu, and Andrew Waxman. I also thank Michael Anderson, Lucas Davis, Eyal Frank, Todd Gerarden, Catherine Hausman, Guojun He, Jonathan Hughes, Akshaya Jha, Matthew Kahn, Ryan Kellogg, Christopher Knittel, Ashley Langer, Arik Levinson, Joao Macieira, Matthew Neidell, Paulina Oliva, Ariel Ortiz-Bobea, Nicholas Sanders, Jay Shimshack, Brad Shrago, Katherine Sims, Eric Zou, and seminar and conference participants at Cornell University, London School of Economics, Stanford SIEPR, Hong Kong University of Science and Technology, Jinan University, Peking University, NBER Summer Institute Environmental Energy Economics Workshop, AERE Conference, AAEA Meeting, OSWEET workshop, SEA Meeting, UC Boulder Environmental & Resource Economics Workshop, NTA Conference, Northeast Workshop on Energy Policy and Environmental Economics for helpful comments.

1 Introduction

The effectiveness of regulation depends on economic agents' responses to the enforcement design. The expected utility model by [Becker \(1968\)](#) predicts that violations will fall when the probability of detection rises. Although enforcement mechanisms often aim to maintain a high probability of detection in general, in practice this probability varies over time for individuals. One overlooked dimension that affects the efficacy of enforcement is how people respond strategically to temporary reductions in the risk of being caught: for instance, under limited resource to carry out regulations, regulators rarely conduct repeated inspections on recently checked agents, resulting in a sharp drop in the probability of detection following an inspection. Such practices could undermine the effectiveness of the regulation if people react by reducing their compliance. Hence, understanding responses to changes in the probability of detection is crucial for informing policy design.

This paper provides new insights into the study of regulatory enforcement and compliance by evaluating a key, long-standing national road safety regulation: the North American Standard Inspection Program for commercial motor vehicles (CMVs).¹ I quantify the effect of the regulation on all trucks inspected over two decades. I analyze the causes of the change in accidents, highlighting driver's compensating responses to problems in the inspection design, and suggest ways to improve the regulation.

The US trucking industry represents a multi-billion-dollar sector that transports the nation's freights and provides essential services to passengers across North America. The value of shipments transported by trucks represents 65% of all shipped goods from all modes of transportation. However, the sector is a major source of safety externality ([Muehlenbachs et al. \(2021\)](#)): Every year, more than 130,000 people are involved in large truck accidents; about 4,000 people die from those accidents.² To ensure that both vehicles and their drivers are safe for travel, the Department of Transportation (DOT) implemented a national safety inspection program for all CMVs. The inspections are performed at weigh stations along US highways by law enforcement personnel. When the weigh stations are open, all trucks must enter, then some trucks will be selected for inspections. One special aspect of the inspection design is that the inspectors will focus less on vehicles that recently pass an inspection, so that inspection efforts are not duplicated for trucks already inspected. Such practice could result in a problem in the regulatory design that causes a significant drop in inspection probability following an inspection. Though the program has been in place for decades, empirical studies have not evaluated the program's impacts on road safety.

This paper evaluates the effectiveness of the roadside safety inspection program on trucks in-

¹Section 204 of the Motor Carrier Safety Act of 1984 (MCSA) (Pub. L. 98-554, Title II, 98 Stat. 2832, at 2833) define a "commercial motor vehicle" as one that has a gross vehicle weight rating (GVWR) of 10,001 pounds or more; is designed to transport more than 15 passengers, including the driver; or transports hazardous materials in quantities requiring the vehicle to be placarded.

²From National Highway Traffic Safety Administration (NHTSA)'s calculation.

spected, using the most comprehensive data available on trucks, roadside safety inspections, and accidents assembled. The set of inspection files was obtained from the US Department of Transportation's national network of highway weigh stations and patrol inspectors. It contains a complete inspection record of more than 23 million trucks inspected in all 50 states from 1996 to 2018. Crash files are collected from state police crash reports over the same period, covering all accidents involving CMVs. These files have an advantage over other accident analysis records (e.g., those from the Fatality Analysis Reporting System and the National Highway Traffic Safety Administration (NHTSA) State Data System) since they record the full vehicle identification number (VIN) for trucks involved. Linking inspection and accident history using each truck's unique VIN, the data allow me to match the trucks involved in accidents with their inspection records. Thus, I can implement an event study research design that tracks a given truck's accident rate shortly before and after its inspection.

I provide several empirical tests to support the identification assumption, rule out important confounding factors, and discuss potential selection issues. To identify the impact of inspection on accidents, the event study framework uses time-series variation in the accident rate of a truck around its inspection. Throughout the analysis, a valid identification relies on the assumption that the trend in the accident rate remains flat in the absence of an inspection. First, I show that there is no pre-trend in accidents in the 14 days preceding the inspection, which supports the identification assumption. Second, I perform a placebo test, which randomly assigns inspections to trucks in the sample. The test shows that the placebo inspections do not lead to a change in accident rates. Third, I address issues related to sample selection that arises from the endogeneity in inspections to trucks that are on the road only during inspections. Actually truck drivers are on the road most of the time: A typical trucker drives 500 miles a day for 300 days in a year. To further alleviate the concern, I test the robustness of the results by dropping trucks that had a serious accident and need repairs, and those that were pulled off the road due to violations found in inspections. In addition, I argue that even if inspections are not random, it does not affect the identification because of the truck-level fixed effects that control for differences in truck characteristics due to inspection selections. Fourth, I rule out omitted variables, including changes in weather patterns or traffic conditions, as potential causes of the changes in accidents following inspections. Lastly, to eliminate concerns regarding reverse causality from the post-accident inspections (those caused by accidents), I drop trucks that crash immediately before inspections. These identification tests and sub-sample results underscore that any changes in accidents are caused by the inspections due to problems in enforcement design.

The main finding of this paper is a sharp, 44.6% *increase* in accident rates (2.84/6.37 accidents per 100,000 trucks) immediately following an inspection - an effect that lingers for at least 14 days. In the longer term, the increase in accident rates can be detected up to 12 months after inspections; moreover, there is no compensating reduction in accidents after 12 months. The current inspection

design leads to an estimated 1,803 additional accidents per year, costing roughly \$1.6 billion.

Next, I investigate why accident rates increase following inspections. First, I find that a truck is much less likely to be reinspected for at least a quarter following an initial inspection. This is caused by the regulatory design that focuses more on trucks which have not been inspected in a while, trucks recently inspected face a much lower inspection probability. Correspondingly, in the wake of an inspection, there is a corresponding increase in the involvement of recently checked trucks in single-vehicle accidents, which serve as a good proxy for driver behaviors. While multi-vehicle accidents can have a number of causes, single-vehicle accidents, such as non-collision accidents or collisions with unmovable objects, stem almost exclusively from the behavior of an individual trucker. Further, my findings show that immediately following inspections, accidents increase for both reckless driving and lack of vehicle maintenance, such as driving under the influence, and loose cargo or trailer. The findings suggest an explanation consistent with the [Becker \(1968\)](#) model: knowing that the trucks are unlikely to be reinspected in the near term, hence the implicit cost of violations decreases, drivers conduct fewer pre-trip checks of their trucks, reduce caution in driving, and drive longer hours, offsetting the potential benefits and intention of the safety program.

This paper also seeks to answer a few policy relevant questions with the data: Consider an interstate truck driver who travels from coast to coast, how should the states in her route design the inspection regime? How can the states increase deterrence considering drivers' behavioral responses to changes in the risk of detection? Would it be better to inspect trucks randomly so that the probability of detection does not depend on past inspection timing? To answer these questions, I first check that states with high probability of reinspection indeed have lower accident rates following inspections, which is consistent with the main mechanism found in the paper. I then look for different inspection designs across states that drive the difference in probability of reinspection in two dimensions. I first explore the heterogeneity in inspection schedules across states. I show that there is a smaller increase or even a reduction in post-inspection accident rates in states that use more unpredictable inspection schedules. I then explore the difference in inspection selection practices across states. I look at the extent to which the inspectors rely on time passed since the last inspection to determine which trucks are chosen for inspection today. I find that states that rely on randomness to a greater degree achieve better results than states that are less likely to reinspect trucks that have undergone a recent inspection. These findings help shed light on problems with the design of the national safety inspection network.

This paper contributes to the literature in the following four ways: It contributes to a broad literature on regulatory enforcement through periodic, random, or targeted inspections. Regulatory efficacy hinges on increasing the probability of detection through designing the most suitable form of inspections. On the one hand, the literature has found that violations decrease when the probability of detection increases through targeted inspections under various contexts ([Muehlenbachs et](#)

al. (2016), Duflo et al. (2018), Johnson et al. (2020), Blundell et al. (2020),). On the other hand, violations rise when the probability of detection decreases due to high predictability in enforcement schedules (Zou (2018)) or corruption in inspections (Oliva (2015)) . This paper adds to the literature by illustrating that the probability of detection decreases when regulators focus less on recently inspected individuals. This loophole in inspection design leads to violations and hinders the regulatory effect as people exhibit compensating behaviors afterwards. Thus, this paper finds that randomized inspections could prevent such strategic responses, similar to Okat (2016), Banerjee et al. (2019), and Gonzalez-Lira and Mobarak (2021), albeit in a different context.

Second, this paper builds on the literature that studies offsetting behavioral responses to safety regulations. Peltzman (1975) first shows that drivers respond to seat belt laws by driving more recklessly since they feel safer, a phenomenon known as the Peltzman effect. Controversies remain about whether and to what extent drivers' responses offset effect of regulations (Cohen and Einav (2003), Lv et al. (2015)). More broadly, similar debates exist in the context of product safety regulation (Viscusi (1984)), occupational health and safety regulations (Viscusi (1986)), and cigarette taxes (Adda and Cornaglia (2006)). This paper provides additional evidence supporting the Peltzman effect; that is, in this case, when the risk of reinspection falls following an inspection, drivers exhibit offsetting behaviors that undermine the intended policy goals.

Third, this paper emphasizes the importance of studying the safety impacts of heavy vehicles, which could potentially lead to severe accidents due to their sizes and weights, but generally attract less discussion in the literature (e.g. Li (2012), Anderson and Auffhammer (2014), Graham et al. (2015), Muehlenbachs et al. (2021)).

On the policy front, this is the first paper to examine the impact of the commercial motor vehicle safety inspection program on accidents at the national scale over a long and up-to-date time frame (1996-2018). Previous studies on this program are limited by data availability or empirical methodology (e.g. GAO (2005), Loeb and Gilad (1984)). Similar to this paper, a few studies have also found that inspections are not effective. Keeler (1994) finds that the program was successful in 1970 but not in 1980 based on county-level data from those two years. Kwigizile et al. (2016)) find that current inspections are not economically beneficial using data from Michigan. This paper provides evidence that the effect from inspections is attenuated by the strategic responses of drivers using most granular data at the individual truck level.

The paper proceeds as follows. Section 2 reviews the institution details of the trucking industry and the inspection program. Section 3 describes the data used and samples constructed. Section 4 presents the empirical framework, main findings, and identification challenges. Section 5 discusses the mechanisms behind the increase in accidents following inspections, heterogeneity of the effects, and the longer-run impacts. Section 7 discusses alternative regulatory designs to improve safety. Section 8 concludes.

2 Regulatory framework

2.1 The trucking industry

The trucking industry is an important sector in the United States. In 2016, commercial motor vehicles (CMV) (or, for simplicity, trucks and buses) transported goods worth \$700 billion, which, as measured by value, represented 65% of all shipped goods from all modes of transportation. Among the 269 million total registered vehicles at that time, 11 million (4% of the total) were trucks,³ and 1 million (0.4% of the total) were buses. In 2016, all vehicles traveled 3,174 billion miles, with large truck journeys accounting for 9.1% (287.9 billion miles) and buses for 0.5% (16.3 billion miles) of the total. The trucking industry employs millions of workers. There were 6 million truck and bus drivers, and 543,061 active motor carrier companies operating in the US in 2017.⁴

2.2 Commercial motor vehicle safety

The trucking sector has a large cost on human lives. Every year in the US, more than 130,000 people are involved in CMV-related accidents, more than 60,000 people were injured, and more than 4,000 people died from these accidents.⁵ Large trucks and buses together accounted for 7% of the total number of vehicle fatalities in the US in 2016. Over the past 10 years, the number of fatalities involving large trucks increased 17%, while the number of vehicle miles traveled increased by only 10%.⁶

Because truck safety has significant implications for both truck drivers and other motorists sharing the road, it is a major consideration for the trucking industry. The Federal Motor Carrier Safety Administration (FMCSA) is the regulatory agency to improve safety and prevent CMV-related accidents. Approximately 7,000 certified roadside inspectors throughout the US assist the agency by conducting approximately 3.5 million roadside inspections per year. The inspections are known as the North American Standard Inspection Program.

2.3 The safety inspection program

Commercial vehicle roadside safety inspections are checks performed at fixed weigh stations or mobile locations along the major highways to ensure both vehicles and drivers are safe for travel. Despite its potential to improve road safety, the CMV safety inspection program generates controversy over whether its costs exceeded its benefits. First, the benefit of the program is unclear due to

³Trucks include single-unit trucks (8.7 million) and combination trucks (2.8 million).

⁴From Federal Motor Carrier Safety Administration (FMCSA) 2018 Pocket Guide to Large Truck and Bus Statistics.

⁵From National Highway Traffic Safety Administration (NHTSA) calculation.

⁶Statistics sources are from NHTSA and Bureau of Transportation Statistics (BTS). The percentage increase of crash is calculated by author.

a lack of empirical evidence in how many deaths it has prevented. Second, the costs of enforcement by the regulators and compliance from the private trucking companies are enormous. Enforcement agencies themselves question the program's efficacy, noting that inspections are "no longer leading to annual increases in the industry-wide level of compliance with safety regulations" (GAO (2005)). In the private sector, anecdotal evidence shows that for each in-and-out at the inspection stations, trucks lose \$2.50 per minute they spend at the stations.⁷ Although large efforts have been made in conducting the inspections, the current inspection design cannot be justified unless it can demonstrate that it significantly enhances safety.

2.3.1 Weigh station inspections

All trucks passing by weigh stations must enter them when they are open.⁸ I recorded a video showing that it is indeed the case that all trucks in sight complied. Figure A.5 in the appendix shows the procedures of the inspections: the first picture (Figure A.5a) shows a line of trucks waiting at the ramp to enter the weigh station. After entering, all trucks drive onto a scale to be weighed (Figure A.5b). While they are weighed, the inspectors determine whether the truck will receive a closer inspection based on the truck onsite as well as the truck's safety record (Figure A.5c). Therefore, all trucks must enter the weigh station, but only a selected number of trucks will go through full inspections; others are allowed to proceed immediately. The inspections in the data files used in this study include only those full inspections.

What is it that determines whether a truck will be inspected? First, trucks undergo checks if the inspector suspects any problem with the truck or the driver based on her observations on site. Second, trucks may also be selected for inspections if they have bad safety histories or have not been inspected in a long time. Carriers will be targeted for inspections if they had bad safety records including information collected from previous inspections, crash records, company reviews and violations from investigations. The algorithm of the safety records system will also flag carriers that have not been inspected in the past 12 months. In addition, inspectors will focus less on vehicles that recently pass an inspection so that inspection efforts are not duplicated for trucks already inspected. The selection of truck inspection based on recent inspection histories is a critical detail in the main mechanism of the findings in this paper.

⁷The estimates are reported from Doug Johnson, director of marketing for Drivewyze, which offers a pre-clearance weigh station bypass service. The company analyzed around 13 million individual "site visits" across the US between September and October of 2015.

⁸Only a small number of carriers with good safety scores are allowed to participate in Weigh Station By-Pass system. The bypass system on board will show whether the truck need to enter into the weigh station for inspection or not.

What kind of things are checked during inspections? Depending on the level of scrutiny determined on site, inspections (Figure A.5d) can include six types of Behavior Analysis and Safety Improvement Category (BASIC) examinations. The inspections cover both the driver and the vehicle. Driver inspections check daily log books that record time and distance traveled, hours-of-service compliance, controlled substances/alcohol consumption, and driver fitness. Vehicle inspections cover all aspects of vehicle maintenance, including the brake, tires, and cargo securement, etc. For hazardous material carriers, inspections also cover the hazmat compliance.

Different violations are given for problems found during the inspections. Appendix A describes the various inspection outcomes in details. If no critical violations are assigned to the vehicle during a thorough vehicle inspection, then the vehicle could be issued a CVSA decal, and those trucks will not be reinspected during the three-month time frame in which the decal is valid. Trucks with any violations or defects must be corrected within 15 days of receiving the violation. Another 20% of inspections result in out-of-service violations, indicating that the vehicle or driver presents an imminent hazards to the public. The effect is immediate; the vehicle may not be driven until all necessary repairs are made, and all the violations are corrected.

3 Data description and raw data patterns

In this research, I compile a comprehensive database on trucks, safety inspections, and accidents in every state across the US from 1989 to 2018. The set of files are from the Federal Motor Carrier Safety Administration (FMCSA). FMCSA maintains a complete record of inspections and accidents for commercial motor carriers (trucks and buses) and hazardous material shippers. The records are electronically transmitted from the states to the FMCSA using a reporting system (SAFETYNET). Therefore, a truck's safety records can be checked at different states in the central system. I also observe the firm characteristics of the currently active carriers in the company census data described in appendix B. In addition to the data files obtained from FMCSA, I use data on fatal accidents from the Fatality Analysis Reporting System (FARS), and accident data directly from the Texas Department of Transportation (DOT). I also use highway traffic volume and daily weather records to construct the external condition covariates (details about the data in appendix B).

Inspection files. The inspection records are collected by state and federal inspectors. The inspection files contain information on the time, location, inspection facility, and outcome of inspections for more than 69 million inspections from 1989 to 2018. The records show that 23 million trucks received inspections. Figure A.6 and A.7 in the appendix shows that, over the past decade, as the total number of commercial vehicles has increased, the total number of inspections conducted has been relatively constant. So, the number of inspections that any given truck could receive has decreased. For weigh station inspections, Figure A.8 and A.9 in the appendix show that the inspec-

tion program covers the whole US, but that the intensity of inspections varies considerably among the states.

The inspection data files also contain information on each truck's unique VIN, license plate number, and carrier company. This information allows me to link each truck's inspection record with other records involving the same truck or the same company.

Crash files. The main crash files I use from FMCSA are collected from state police crash reports. The files contain information on the time, location, number of injuries and fatalities, accident event type, VIN, and carrier company for all trucks involved in accidents from 1989 to 2018. Figure A.4 in the appendix shows that the annual total number of truck-related accidents increases over time, and that the trend is procyclical to the general economy. The vehicle miles traveled by CMVs⁹ exhibits a much mild increase over the period.

The files also record the number of vehicles involved as well as the circumstances of the accidents, allowing me to infer the factors that contributed. This is the key information that allows me to tease out the mechanism of the findings in the paper.

In addition to the FMCSA files, I explore two other sources of crash records. One is FARS, which include all fatal crash records maintained by National Highway Traffic Safety Administration. The FARS data used in the paper cover from 2000 to 2017. There were 472 fatal crashes identified within 14 days before or after an inspection for the trucks in the sample. The other crash record is directly obtained from Texas Department of Transportation from 2010 to 2017.¹⁰ There are 756 crashes that involved inspected trucks within the event window.

Both FARS and Texas crash data contain detailed information on the factors contributing to each accident, for example, the driver-related factors, such as speeding, changing lanes recklessly, or driving while intoxicated; or vehicle-related factors, such as not having brakes or functioning lights; or information indicating that truck drivers are not at fault. Therefore, I use the two crash data to examine the exact causes of accidents, which provides important support to the underpinning mechanisms I find.

However, there are also disadvantages for using FARS or Texas crash data over the crash data maintained by FMCSA. FARS data set only contains the first 12-digit of VIN for privacy purposes, but the FMCSA has access to the full (17-digit) VIN, which is critical to link to the truck inspection data files.¹¹ So, I can directly compare the crash rate for the same truck before and after it receives an inspection, which is the key to identification. Second, the FMCSA crash files contain all truck-related accidents including non-injury or non-fatal ones. Since I am interested in road safety in general, crashes that result in property damage and/or life loss are both relevant. As for Texas DOT crash data, although it has full records on all types of crashes, it only tracks events in its own state.

⁹The vehicle miles traveled by CMVs is obtained from the Bureau of Transportation Statistics.

¹⁰The accident data maintained by Texas DOT are only available from 2010 to 2017 to the public.

¹¹The last six digits of VIN are the serial number of a vehicle. So, they are critical for identifying a vehicle.

Therefore, I use the crash data from FMCSA to estimate the effect of inspection on accidents, but use FARS and Texas crash data to provide supporting evidence for the mechanisms of the effect; and the two also add robustness to my findings with multiple sources of data.

3.1 Summary statistics

I summarize the key variables used in the paper in Table 1. Panel A of Table 1 shows the summary statistics of inspection files. There are 69,549,512 inspections recorded for 23 million trucks from 1996 to 2018 across the US.¹² Nearly half of the inspections are conducted at fixed weigh stations, and I focus on these in this paper. On average, a truck is inspected 3 times in its lifetime (10 years) at weigh stations. Regulators reinspect 32% of trucks. On average, 791 trucks pass through a county from both directions during inspection hours.

The inspection and crash files are the two main data sets combined to create an truck-inspection-crash daily panel that traces all accidents that happened to the same truck inspected within the time frame of interest. Appendix B illustrates the sample construction. As shown in panel B of Table 1, there are 670,402,180 observations in the panel constructed. There are on average 6.37 accidents per 100,000 trucks inspected in a day. Among these accidents, 58.5% result in injuries, and 3.8% result in fatalities. For all trucks in the sample, a truck on average experiences 0.005 accidents in its lifetime.

To analyze the long-term impact of inspections, I construct a monthly event panel that is similar to the daily panel, but the unit of observation is a month. The time frame of interest here is 12 months before to 24 months after inspections. Table 1 panel C shows the summary statistics of the monthly event panel. There are on average 160 crashes per 100,000 trucks inspected in a month.

Raw data patterns: Recently inspected trucks had more crashes Before presenting the empirical framework in the next section, here I plot the relationship between number of crashes and crash-inspection time interval using the raw data to reveal some simple patterns. For all accidents in the crash data file, I look for the time interval between the crash and the latest inspection of the same truck occurred before the crash. In Figure A.10, I simply count the number of crashes by the time interval in 1 to 12 months. The left figure adds up all accidents, while the right figure separately plots accidents by state. Both figures are raw data patterns without any control variables. Both figures reveal the same pattern that more accidents occurred for recently inspected trucks compared to trucks that have not been inspected in a while: more than 190,000 crashes happened in the first

¹²The data requested from DOT are from 1989 to 2018. However, the data quality before 1996 is significantly poorer. For example, all records of VIN are missing for truck inspected before 1996; the records of inspection facility are largely missing before 1994; and all inspection records in the year 1995 were not available. As a result, this paper uses data starting from 1996 for consistency of data availability and quality.

Table 1: Summary Statistics

	(1) Observations	(2) Mean	(3) SD	(4) Min.	(5) Med.	(6) Max.
<i>Panel A: Inspection file, 1996 to 2018</i>						
Total inspections	69,549,512	-	-	-	-	-
Weigh stations inspections	30,101,086	-	-	-	-	-
Trucks ever inspected	23,078,901	-	-	-	-	-
OOS violations	6,158,151	1.45	1.02	1	1	56
Driver violations (not OOS)	4,953,818	1.43	0.94	1	1	80
Vehicle violations (not OOS)	9,322,306	2.34	2.05	1	2	85
Truck volume	3,354,901	790.98	1957.23	1	410	1,280,036
<i>Panel B: Daily event study panel: (-14,13) days of inspection</i>						
Crashes	670,402,180	6.37E-05	0.008	0	0	2
Injuries	670,402,180	3.64E-05	0.010	0	0	48
Fatalities	670,402,180	2.43E-06	0.002	0	0	10
<i>Panel C: Monthly event study panel: (-12,24) months of inspection</i>						
Crashes	191,668,177	0.0016	0.058	0	0	32
Injuries	191,668,177	0.0009	0.0508	0	0	49
Fatalities	191,668,177	0.00005	0.008	0	0	11
<i>Panel D: Carrier companies</i>						
No. of trucks	1,669,661	20.91	3420.75	0	1	2,699,990
No. of drivers	1,669,661	5.15	225.72	0	1	110,690

Note: In Panel A, *Total inspections* is the total number of inspections conducted in the whole sample from 1996 to 2018; *Weigh station inspections* are those conducted at the weigh stations, as opposed to at roadside. *Trucks ever inspected* is the number of trucks that ever receive an inspection. *OOS violations* is the number of out-of-service (OOS) violations found in inspections. *Driver violations (not OOS)* is the number of driver violations (but not OOS violations) found in inspections, similarly with *Vehicle violations (not OOS)*. *Truck volume* is the total number of trucks that pass through the inspection county at the inspection hour from both directions.

In Panel B, *crashes* is the number of accidents resulting in at least some property damages for a given truck in a day within the event window, excluding trucks that receive out-of-service violations in inspections, similarly with *injuries* and *fatalities*.

In Panel C, *crashes* is the number of crashes for a given truck in a month within the event window, similarly with *injuries* and *fatalities*.

In Panel D, *No. of trucks* is the number of trucks owned by each carrier company recorded in the company census file, which is a snapshot of all active carrier (and shipper) companies as of October 2018, similarly with *No. of drivers*. It shows that there are 1,669,661 active companies registered as of October 2018.

quarter following inspections, and only 70,000 crashes happened in the fourth quarter. Indeed, one cannot draw conclusive evidence from the raw plots due to inspection selection that trucks repeatedly receiving inspections (thus having shorter time intervals) could be more dangerous. However, this negative relationship between number of crashes and the crash-inspection time lapse does suggest that, under the current regulatory design which mostly focusing on inspecting trucks that have not been inspected in a while, it leaves out a large fraction of trucks which caused more accidents. The misalignment of regulatory efforts and targeted truck groups motivates the research question in the paper. In the next section, I empirically identify the causal effect of inspection on accidents, and discuss potential concerns in identification.

4 The impact of inspection on truck accidents

4.1 The econometric framework

This section describes two econometric specifications and the identification assumption required to consistently estimate the effect of an inspection on truck accidents. The first specification is an event study, which allows for the effect of an inspection to vary over the 14 days afterwards, and tests the zero pre-trend assumption. In the second specification, I estimate the effect using a post-inspection indicator for time after the inspection to test for differences in accident rate 14 days before and after the inspection, which efficiently estimates the average effect size of one inspection on accidents up to 14 days later.

In the event study framework, I regress the number of crashes for the same truck that receives the inspection on a set of inspection indicators from 14 days before to 14 days after inspection, controlling for individual truck fixed effects, year, month, and day-of-week fixed effects.

Specifically, the regression equation I use is the following:

$$Crash_{it} = \sum_{\tau=-14, \tau \neq -1}^{13} \beta_{\tau} Insp_{it}^{\tau} + \beta_{-15} Insp_{it}^{-15} + \beta_{14} Insp_{it}^{14} + u_i + \eta_t + \varepsilon_{it}, \quad (1)$$

where $Crash_{it}$ is the number of crashes made by truck i on day t .¹³ $Insp_{it}^{\tau}$ is an inspection indicator, $Insp_{it}^{\tau} = 1$ if as of time t , truck i experiences an inspection τ days ago. Since inspection happens at event day 0 ($\tau = 0$), the 28-day event window is between event day $\tau = [-14, 13]$. $Insp_{it}^{-15}$ and $Insp_{it}^{14}$ are indicators equal to 1 if there is any inspection for truck i that happens before or after the 28-day event window, respectively, and 0 otherwise. u_i is the individual truck fixed effect. Each truck is identified by using its VIN primarily, or, if VIN is missing, by using the license

¹³Most of the values of $Crash_{it}$ are 0 or 1 at the daily level because a truck normally is involved in an accident once in any given day in the estimation sample.

plate number. The truck level fixed effects ensure the identifying variation comes from day-to-day changes in accidents from the same truck.¹⁴ η_t includes year, month and day-of-week fixed effects. The day-of-week fixed effects are included because the inspection schedules largely follow a day-of-week pattern that more inspections are conducted during weekdays, which is consistent with truck traffic being heavier during weekdays and lighter on weekends. Location fixed effects, such as county fixed effects, are not included in the baseline estimation since they barely change the results here. I combine the daily inspection indicators into two-day bins to increase the power of estimation, such as (-14,-13), (-12,-11), ..., (-2,-1), (0,1), (2,3), ..., (12,13) relative to the day of inspection at day 0.¹⁵ The effect of an inspection on accidents happening in the two-day bin (-2,-1) is normalized to 0. The standard errors are clustered at the truck level.¹⁶

To estimate the average effect of an inspection on accidents throughout the 14 days after the inspection, I use the following econometric framework, which regresses the number of crashes on a post-inspection indicator, also controlling for inspections outside the event window, individual truck fixed effects, year, month, and day-of-week fixed effects. The estimation equation is the following:

$$Crash_{it} = \gamma post-insp_{it} + \beta_{-15} Insp_{it}^{-15} + \beta_{14} Insp_{it}^{14} + u_i + \eta_t + \varepsilon_{it}, \quad (2)$$

where $post-insp_{it}$ is a post-inspection indicator; it equals zero before an inspection happens on truck i , and it equals one on and after the inspection. The other variables used in this framework are the same as in equation 1.

Throughout the analysis, both econometric frameworks identify the impact of inspections on accidents using time-series variation in the accident rate before and after a given truck undergoes an inspection. The validity of such identification relies on the assumption that the trend in the accident rate remains flat in the absence of an inspection. In Section 4.3, I discuss why this assumption holds by addressing all potential confounders and selection issues. I also discuss potential treatment effect heterogeneity across different trucks inspected at various timing (De Chaisemartin and d’Haultfoeuille (2020), Sun and Abraham (2020)) in appendix C.

4.2 Main Results


The event study framework described above allows me to flexibly estimate the effect of an inspection on accidents beginning with the 14th day prior to the inspection and ending with the 14th day after the inspection.

¹⁴Gray and Shimshack (2011) discuss the endogeneity of selected inspections. Here I resolve the issue by including individual truck fixed effects. I discuss sample selection issues in detail in Section 4.3.2.

¹⁵Single day inspection indicators generate similar estimation results.

¹⁶I perform a placebo test in Section 4.3.4. The t-statistics (Figure A.13) from the placebos follows the t-distribution, which suggests that the clustering at truck level is appropriate.

Figure 1: Event study: the impact of an inspection on truck accidents



Note: The coefficients plotted in this figure are estimated using equation 1 where I regress the number of accidents for the same truck that receives the inspection on a set of inspection indicators from 14 days before an inspection to 14 days after inspection, controlling for individual truck fixed effects, year, month, and day-of-week fixed effects. I illustrate the level shift in accidents before and after the inspection by fitting two horizontal lines using the average of percentage changes respectively. The shaded area is the 95% confidence interval for the estimates. The standard errors are clustered at the truck level. I combine the daily inspection indicators into 2-day bins, such as (-14,-13), (-12,-11), ..., (-2,-1), (0,1), (2,3), ..., (12,13) relative to the day of inspection at day 0. The horizontal axis is labeled by the first day of the 2-day bins. The effect of an inspection on accidents happening in the two-day bin (-2,-1) is normalized to 0.

The result is summarized in Figure 1. I find that accidents involving trucks inspected *increase* by 44.6% immediately following an inspection, and that the increase lasts for at least two weeks after the inspection. The percentage change in accidents is relative to a mean of 6.37 accidents per 100,000 trucks inspected per day. I illustrate the level shift in the accident rate in Figure 1 by fitting two horizontal lines using the average of percentage changes before and after the inspection, respectively. In addition, the figure shows that there is almost no pre-trend during the 14 days before the inspection; this suggests that the identification assumption is valid.

To gauge the effect of an inspection on accidents more concisely, I estimate equation 2 using one post-inspection indicator instead of 14 leads and lags indicators. Table 2 column 1 shows the regression result, which confirms the finding from the event study. It shows that the post-inspection increase represents 2.84 more accidents per 100,000 trucks inspected in a day comparing to the pre-inspection level. The increase in accident rate is 44.6% relative to the average accident rate.

4.3 Tests for identification

In this section, I discuss how I rule out all potential confounders to the identification, as well as concerns regarding sample selections. As mentioned previously, both econometric frameworks rely on the identification assumption that the trend in the accident rate remains flat in the absence of an inspection. There are three concerns regarding the identification strategy: first, **omitted variables** that are correlated with both inspections and accidents such as, weather conditions and traffic conditions; second, **sample selection issues**, including trucks being idle before or after inspection, trucks switch drivers after inspections, and inspection selection on truck observables; and third, the possibility of **reverse causality** that inspections are triggered by accidents. Lastly, I implement a **placebo test** to show that there would not be an increase in accidents without the inspections.

4.3.1 Potential confounders: weather and traffic

Weather conditions, or traffic conditions, are potential confounders because adverse weather conditions and high traffic volumes contribute to more accidents. At the same time, if inspections are also scheduled at times when the probability of accidents in the area is high, then the observed increase of accidents following inspections might be caused by inspections that precede periods of heavy rain or heavy traffic.

I rule out that weather conditions are confounders by showing that the inspections are not chosen at times of worse weather conditions (such as rain or snow). On the contrary, I find that the inspections take place on days with better weather conditions. Therefore, even if adverse weather conditions do lead to higher probabilities of accidents, the observed *increase* in accidents after an inspection is not caused by adverse weather conditions. Details are included in appendix C.

I can also rule out higher traffic volume as a potential cause of the increase in accidents after an inspection. Table B.1 in the appendix shows that the effect of inspections on accidents stays the same after controlling for traffic conditions in the inspection county at the inspection hour in different ways. Details are also shown in appendix C.

4.3.2 Sample selection issues

In this section, I address several concerns (in bold) related to sample selections. The selection problem arises because of the *endogeneity of inspections to a truck being on the road*. I break down the discussion of potential issues caused by selections in four ways, and discuss the solutions to them or the implications regarding the interpretation of the estimated effects.

Are trucks on the road only during and after the day of inspection? First, a typical truck driver drives 500 miles a day for 300 days a year. Even if they are not on the road every single day, most drivers won't be idle for too many days consecutively. And the idle periods are not likely to happen all before the inspections. In addition, I drop trucks that receive out-of-service violations (20% of all inspections) from the estimation because those trucks are pulled off the road after inspection until all repairs/corrections are made.¹⁷

What if a truck crashes in days before inspections and needs to be repaired, so it is less likely to be on the road? To address this concern, I perform a robustness test that only looks at no-injury or nonfatal accidents, so trucks who have a serious accident then stop operating for a while are dropped. Table B.8 in the appendix shows that, for this subsample of trucks, there is a 45.6% increase in accident rates after inspection. Therefore, the truck selection issue described does not bias the main result.

Because trucks are not randomly selected for inspections, would it affect the identification and/or the interpretation of the estimated effect size? There is inspection selection on trucks when all trucks enter the weigh station. But under the research design of this paper, sample selection here only affects the interpretation of the estimation results but not the identification of the econometric framework. One aspect in the design of the regulation indicates that trucks appear to be problematic or have worse histories (inspection, crash, company review records) will receive more inspections. As a result, trucks appear more often in my sample are likely to be those having potentially higher probability of crash.

It is not a threat to identification since all estimations in this paper include truck level fixed ef-

¹⁷I test that if I do not drop out-of-service trucks in the estimation, the effect size is still similar (43.5%). This sample selection practice aims to increase the chance that the trucks are always on the road, which alleviates endogeneity concerns. Regulation requires that trucks receive out-of-service violations must immediately be taken off the road. Nonetheless, in my data, I still observe accidents following inspections that involve these same trucks. This suggests that some drivers either disregard the citation, which is a serious crime, or that enforcement of such regulations is not strict.

fect using each truck’s unique VIN, which accounts for differences in trucks with different baseline accident rates or other unobservable characteristics. While it is not possible to find a strict comparison group, which consists a sample of trucks not inspected, due to the data collection process that only records the inspected trucks, I generated a “comparison” group by performing the placebo test shown in the next section. The test randomly assigns placebo inspections to trucks, and estimates the change in accidents following the placebo inspections. The result shows that the post-inspection accident profile of the uninspected trucks who were given the placebo inspections is different from the inspected trucks.

The non-random truck selection process do affect the interpretation of the estimated effect size. The correct way of interpreting the result is how inspections affect accidents for those trucks inspected, not for any given truck on the highway. But all trucks will receive inspections at some point of their lifetime trips, discussing the regulatory impacts on trucks’ inspected is still relevant for all trucks on the road. In addition, this study finds that accident rates increase for trucks inspected through evaluating the inspection regime. The finding implies that the regulatory effect of inspections is attenuated by the strategic responses of drivers to the current design, therefore, the current regulatory design needs to be improved.

What if trucks switch drivers after inspections so that truck level fixed effects cannot control for driver behaviors? The fourth concern stems from the discussion above: whether truck VIN fixed effects is sufficient in controlling for difference characteristics across trucks (or truck drivers). I cannot identify the driver-vehicle pairs because I lack information on drivers’ identities in the data; thus, if drivers in these trucks do change following an inspection (sometimes long-haul trucks have two drivers in the vehicle), the different driver before and after inspections would create problem for identification. I address this possibility by looking at a subsample of firms with only one truck and one driver (around 50% of firms), because these firms cannot switch drivers after inspections. The estimated effect of inspection on accidents is 44.53% as shown in Column 3 of Panel A in Table B.5. This potential concern is addressed by the confirmation that a similar effect emerges among both types of firms – those with only one truck and one driver, and those with multiple trucks and drivers. Other related concerns can also be addressed in this comparison, including the possibility that some firms decided to use more of the recently inspected trucks relative to other trucks in the firm.

4.3.3 Eliminate reverse causality

Inspections induced by accidents should be excluded in the estimation to eliminate reverse causality. The issue warrants attention in this particular case because it is not uncommon for a police officer to call a nearby inspector to check a truck that has been involved in an accident. In this case, there is reverse causality because the accident itself invites the inspection. To eliminate that

concern, I drop all trucks that are inspected within 18 hours after accidents. Figure A.2 illustrates how the time interval is chosen, and shows that the coefficients on the other days in the event study are not affected when varying the length of the time interval - except for the coefficients on days (-2, -1) and (0, 1). This also suggests that dropping those trucks does not change the increase in accidents that we observe in the days following the inspections. Details are shown in appendix C.

4.3.4 Placebo test

To provide additional evidence supporting the identification assumption, I perform a placebo test that reshuffles the inspections for a randomly selected group representing 1% of all trucks in my sample.¹⁸ The placebo test illustrates that no other factors could generate such an increase of accidents after inspections except for the observed inspections themselves, which is exactly the identification assumption of the event study.

I randomly select 205,423 trucks from all trucks inspected from 1996 to 2018. For each truck selected, I reshuffle the inspections it receives during its operating years for 500 times.¹⁹ Then I estimate 500 event studies using the same framework in equation 1 and equation 2. The null hypothesis for the placebo event studies is that there is no change in accident rates after the inspections comparing to before. I compare the estimated effect sizes in the 500 placebo tests with the observed effect size estimated using the real sample. The results are shown in Figure 2 and 3. Figure 2 shows that, first, the placebo effect sizes center around zero; second, the observed post-inspection coefficient lies outside of the 95% confidence interval of the distribution of the coefficients from the 500 placebo tests.²⁰ Figure 3 is an event study figure that compares each lead and lag coefficient in the 28-day event window from the observed real sample to those from the 500 placebo tests. It shows that, after inspections, almost all coefficients of the observed sample lie outside of the 95% confidence interval of the placebos. The mean of the placebos is almost a flat line, which implies that the inspections in the placebo tests would not cause an increase in accidents after inspections. Both results suggest that there will not be changes in accident rate to the trucks without the observed inspections, which supports the identification assumption.

¹⁸There are 23 million trucks in the full sample, so I randomly select a smaller but sufficient sample to save computation time.

¹⁹For each truck selected, I calculate the total number of inspections it receives during its lifetime. I then randomly redistribute the same number of inspections for the given truck in between its first and last inspection observed in the real sample. This ensures that the same truck gets the same number of inspections in the placebo test so that truck operating patterns remain the same.

²⁰The t-statistics of the post-inspection coefficients from the 500 placebo test follows the t-distribution, as shown in Figure A.13 in the appendix, which suggests that the clustered standard errors are specified correctly.

Figure 2: Placebo test: compare effect sizes





Figure 3: Placebo test: compare 14 days before and after an inspection



Note: Figure 2 shows that the observed post-inspection coefficient (the red solid line) lies outside of the 95% confidence interval of the distribution of the coefficients from the 500 placebo tests (the histogram).

In Figure 3, the red solid line indicates the coefficient estimates from the observed sample, while the black dashed line is the mean of coefficients from the 500 placebos. The shaded area around each line is the 95% confidence interval of the coefficients.

5 Mechanism

After eliminating all potential confounders, this section discusses the main mechanism behind the increase in accidents following inspections. The mechanism is consistent with the [Becker \(1968\)](#) model: because inspectors rarely conduct repeat inspections of recently checked trucks, the drivers of these trucks drive more recklessly and undertake fewer safety checks after inspections, leading to more accidents. This section first describes two pieces of evidence that point to the mechanism, then provides direct evidence using two supplementary data. Last, I looks at the longer-term impact of inspection and estimate the total economic cost of the accidents.


5.1 Drivers' strategic responses to a low probability of reinspection

First, I find that a truck is much less likely to be reinspected following a recent inspection. [Figure 4](#) shows that for a given truck, the probability of reinspection is only 1.3% within the first week and 8.9% within the first quarter (13-weeks) following an inspection. The probability of reinspection gradually increases to 19.9% after one year (52 weeks). [Appendix D](#) shows the detailed calculation of the probability of reinspection. Because a truck on average receives one to two inspections per year, the average time since last inspection for most trucks will be around six months or even longer. We can then infer from [Figure 4](#) that, before an inspection, a given truck faces more than 15% chance of an inspection, while right after the inspection event, a truck faces only 1.3% chance of an inspection. Hence, there is a significant drop in the probability of inspection following a recent inspection. As mentioned in [Section 2.3.1](#), inspections cover both the driver and the vehicle, then if truck drivers learn that they are much less likely to experience reinspection once they pass an inspection, the implicit cost of violations found during inspections as well as other traffic violations drops.

Reasons for the low probability of reinspection following a recent inspection are as follows. Enforcement through inspections is costly for both the regulatory agency and the trucking industry. To save regulatory efforts and to potentially inspect more trucks, the inspectors rarely conduct repeat inspections on recently inspected trucks. This rationale is embedded in the truck inspection selection process. As mentioned in [Section 2](#), when a truck enters a weigh station for inspection, inspectors will review its inspection and accident history to look for trucks that have a bad history or have not been inspected for a long time. Those trucks are more likely to be selected for inspection. In addition, a vehicle displaying a valid CVSA decal²¹ will not be reinspected during the three-month time frame in which the decal is valid. These two inspection selection algorithms result in a low probability of reinspection for trucks with a recent history of inspection and relatively good records.

²¹If no critical violations are assigned to the vehicle during inspection, the vehicle could be issued a CVSA decal.

Figure 4: Mechanism: probability of reinspection



Note: The probability of reinspection plotted in this figure represents the probability of getting another inspection within N weeks after the most recent inspection. Details on the construction of the probability is in Appendix D.


Table 2: The impact of an inspection on truck crashes

	(1) Baseline (all crashes)	(2) Single-vehicle crashes	(3) Multi-vehicle crashes
post_insp	2.84*** (0.06)	1.64*** (0.06)	1.17*** (0.05)
Average crash rate	6.37	2.19	4.09
Effect size	44.58%	74.89%	28.60%
Observations	670,402,180	670,402,180	670,402,180

Note: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All three columns use the same estimation framework following equation 2, which looks at the impact of an inspection on crashes by comparing 14 days before to 14 days after inspection. *post_insp* is a post-inspection indicator, it equals to 0 before an inspection happen on truck *i*, and it equals to 1 on and after the inspection. All regressions include dummies for inspection happened outside of the time window of interest, as well as individual truck, year, month, and day-of-week fixed effects. Standard errors are clustered at the truck level. Column 1 looks at the impact of an inspection on all kinds of crashes. Column 2 looks at single-vehicle crashes only. Column 3 looks at multi-vehicle crashes only. The average crash rates are calculated using the number of any kinds of crashes, or only single-, or multi-vehicle crashes per 100,000 trucks inspected.

Correspondingly, I find a larger increase in the involvement of these recently inspected trucks in single-vehicle accidents following an inspection compared to multi-vehicle accidents. This difference suggests that truck drivers are responsible for the increase in accidents. Single-vehicle accidents (35% of all accidents) include all non-collision accidents, collisions involving parked motor vehicles, fixed objects, and all other accidents involving only the truck itself. Meanwhile, multi-vehicle accidents are collisions with other motor vehicles in transport (64% of all accidents). As shown in Figure 5, by estimating the same framework using equation 1 with different accident categories, single-vehicle accidents increased by 74.9% following an inspection, and multi-vehicle accidents increased by 28.6%. Both effects occur immediately on the day of inspection and last for at least 2 weeks after the inspection. There is no pre-trend in both cases. Table 2 columns 2 and 3 show the regression results by estimating equation 2 for single- and multi-vehicle accidents. The baseline probability of a single-vehicle accident is smaller than that of a multi-vehicle accident. Figure A.11 in the appendix breaks down the single-vehicle accident categories into non-collision ran-off road, non-collision overturn (rollover), non-collision cargo loss or shift, non-collision equipment failure (brake failure, blown tires, etc.), and collision involving fixed objects. The figure shows that, for most single-vehicle accident categories, percentage increases in accidents are all larger than for multi-vehicle accidents.

Figure 5: Mechanism: comparing effect sizes in single- vs multi-vehicle accidents



Note: This figure compares the two event studies estimated using the same framework in equation 1 with different accident categories. I plot the percentage change in each type of accident using the change in number of accidents with respect to the average number of corresponding accidents. Both event studies look at 14 days before to 14 days after an inspection, controlling for individual truck fixed effects, year, month, and day-of-week fixed effects. The standard errors are clustered at the truck level. The shaded areas are the 95% confidence intervals respectively.

I compare the difference in factors contributing to single-vehicle accidents versus multi-vehicle accidents. The factors contributing to a single-vehicle accident can only be related to the truck drivers' driving behaviors, cargo securement, or truck equipment failures, which are also due to low maintenance efforts by drivers. These factors could be affected by the inspection probability. Factors contributing to a multi-vehicle accident, on the other hand, are either related to the truck or to the other vehicles involved in the accident. Thus, a larger percentage increase in single-vehicle accidents following an inspection indicates that the reason for the increase on average is attributable to the truck driver's behaviors following an inspection. Furthermore, I find that the increase in single-vehicle accidents is even higher when the external conditions are worse, as shown in appendix E. This suggests that the increase is caused by drivers paying less attention to driving conditions.

To summarize, first, I find evidence that a truck is much less likely to be reinspected for at least a quarter following an initial inspection. This is due to the fact that inspectors rarely conduct repeated inspections on recently inspected trucks. Correspondingly, I find a larger increase in single-vehicle accidents, which serves as a good proxy for driver behaviors, as compared to multi-vehicle accidents. Among single-vehicle accidents, I find an even larger increase in crashes under adverse external conditions, which highlights the changes in truck drivers' behaviors. All evidence indicates that, knowing that the truck will not be reinspected in the near term, the driver might conduct fewer pre-trip checks and drive more recklessly on the road for longer hours, which in turn causes more accidents and offsets the potential benefits of the safety program.

I discuss alternative mechanisms, including the possibility of drivers feeling safer after passing inspections that causes them to reduce caution, and the incentives for speeding to make up for time lost due to inspections, in Appendix F.

5.2 How drivers' behaviors contribute to accidents

Using two supplementary crash data sets, I find direct evidence that drivers' compensating behaviors contribute to the increase in accidents following inspections. Section 5.2.1 brings in FARS (Fatality Analysis Reporting System) data maintained by NHTSA, and Section 5.2.2 brings in crash data maintained directly by the Texas Department of Transportation (TxDOT). Both data sets have the advantage of identifying the exact factors contributing to crashes, which allows me to tease out what causes the rise in crashes following inspections. In addition, I find very comparable results to the main findings using these two separately maintained data sources. They provide more robustness to my main findings.

5.2.1 Evidence from FARS

The Fatality Analysis Reporting System (FARS) data contains information on the exact factors contributing to fatal accidents, as well as the demographics of the drivers involved in the accidents, which are not included in the FMCSA crash files. Therefore, using FARS data from 2000 to 2017, I analyze the causes of fatal accidents involving those trucks inspected.

First, Figure 6 shows that there is a large increase in the number of fatal accidents involving trucks following inspections.²² There are 15 fatal accidents among all trucks in each two-day bin on average before inspections. The number increases to 52 following inspections.

Second, from analyzing the factors contributing to the fatal accidents, I find that there are more accidents caused by driver- or vehicle-related factors following inspections. The examples of fac-

²²Here I present the results by plotting the raw total number of fatal accidents instead of performing regressions, because accidents resulting in fatalities are rare events. The trucks ever inspected are involved in a total of 472 fatal accidents in the 28 days before and after their inspections.

Figure 6: Additional evidence: increase in **fatal** accidents


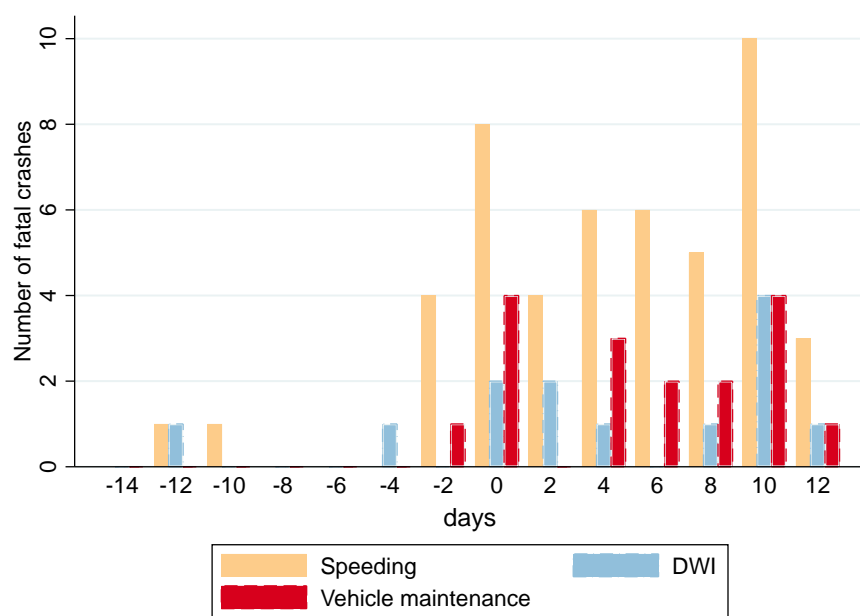


Figure 7: Additional evidence: increase in **fatal** accidents due to **truck** violations



Note: In Figure 6, I plot the raw number of fatal crashes involving trucks around inspection days (event time = 0) in the FARS sample. Figure 7 looks at the total number of fatal accidents caused by truck violations before and after inspections. There are 15 fatal accidents among all trucks in each two-day bin on average before inspections. The number increases to 52 following inspections. The increases in accidents are due to drivers' reckless driving behaviors, including speeding and driving while intoxicated, and due to a lack of vehicle maintenance, such as loose cargo or trailer, which could be caused by reduced pre-trip checks.

tors contributing to an accident include driver-related factors, such as speeding, changing lanes recklessly, or driving while intoxicated; vehicle-related factors, such as loose cargo, not having brakes or functioning lights; or factors where truck drivers are not at fault and the other vehicle involved in the accident is. Figure 7 shows that there are increases in accidents due to drivers' reckless driving behaviors, including speeding and driving while intoxicated. A possible explanation is that the implicit cost of a traffic violation is reduced because drivers know that if they are stopped by police for speeding or drinking, they would be less likely to have an impromptu roadside inspection in addition to the speeding ticket if they were recently inspected at a weight station. In addition, the figure shows that there is also increase in accidents due to a lack of vehicle maintenance, such as loose cargo or trailer, which could be caused by reduced duration and frequency of pre-trip checks at the vehicles. As now drivers face lower probability of reinspection, hence they are less likely to be in violation of vehicle safety regulations.

5.2.2 Evidence from Texas DOT crash data

In this section, I provide additional supporting evidence using data containing all types of accidents obtained from the Texas DOT. The Texas DOT maintains a set of crash files under the Crash Records Information System (CRIS) collected from the Texas Peace Officer's crash reports from 2010 to present. The benefits of using this set of files for supplementary analysis are three-fold: first, the CRIS includes crashes that result in property damage only, injury, and fatalities; second, there are detailed descriptions of the factors contributing to each accident; third, the crash records contain the full VIN.²³

Figure A.12 in the appendix shows that there is a 39.7% increase in truck crashes following inspections, which is comparable to the effect size estimated using data from FMCSA on the whole US.²⁴ The effect size is slightly attenuated because I can only observe crashes in Texas involving trucks inspected in Texas; therefore, crashes that happened outside of Texas for trucks inspected in Texas are not included in the analysis.

Figure 8 shows that there is an increase in crashes due to driver-related factors, including all moving and parked violations assigned to the truck driver. Among those crashes, trucks involving speeding violations exhibit a significant increase in the number of crashes. Figure 9 serves as a falsification test that shows that there is no increase in crashes that are *not* due to truck-related factors.²⁵

²³To find the federal recordable crash accidents as those from the FMCSA files, I use the VIN and date of crash to find crash records in the CRIS that matches up with those in FMCSA. A "federal recordable" crash has occurred when at least one person dies, at least one person experiences bodily injury that requires immediate medical treatment away from the scene of the crash, or a vehicle is towed away according to the FMCSA file documentation.

²⁴I use the exact same framework in equation 1 for all inspections and crashes happening in Texas. There are 1,204,497 inspections used in this analysis.

²⁵The test for differences in accident rates before and after is 1.63e-06, with p-value = 0.063.

Figure 8: Additional evidence: increase in **driver-related** accidents

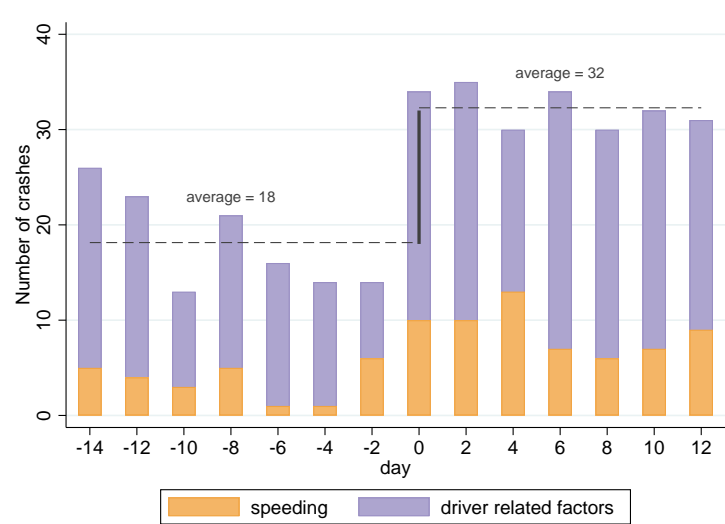



Figure 9: Additional evidence: no effect on other crashes (not driver related)



Note: Figure 8 shows that, before inspections, there are on average 21 crashes due to driver-related factors for each 2-day bin, but the number increases to 37 following inspections. Driver-related factors include all moving and parked violations assigned to truck drivers. Among those crashes, crashes that have a speeding violation exhibit a significant increase. Figure 9 shows that there is no increase in the number of crashes which are *not* due to driver behavior related factors. Test for differences in average crash rates before and after inspection: $1.63e-06$ ($P = 0.063$).

In summary, the FARS data and the Texas DOT data provide additional evidence of drivers' offsetting behaviors to the safety regulation. First, I find that there is an increase in the number of crashes following inspections using both separately maintained data sets, which supports the main finding of this paper. As FARS data analyze fatal crashes, the finding also makes the current study even more imperative. Second, from analyzing the factors contributing to crashes, I find that, after inspections, drivers indeed drive more recklessly (speeding, DWI) and perform fewer vehicle maintenance checks. So the mechanism presented in the previous section is validated with the facts presented in this section.

5.3 Longer-term impact on accidents

In this section, I explore whether the impact of inspection is economically significant by quantifying the cost of accidents that result from behavioral responses to the regulation. In order to estimate the total increase in accidents following an inspection, I look at the overall effect of inspection on accidents in the longer term, or 24 months after an inspection. I find that the increase in accidents is largest in the first two months after inspection; it gradually declines by month 12 to the pre-inspection accident rate and remains at that level afterwards. The total cost of accidents attributable to drivers' strategic responses is \$1.6 billion per year.


I implement a monthly event study that enables me to estimate the impact of an inspection on truck accidents in the longer term. I use a similar framework as the daily event study while changing the unit of time from 1 day to 1 month. The event window is from 12 months before to 24 months after an inspection (or inspections) in month 0 for any given truck. Again, I control for individual truck fixed effects and year and month fixed effects. The estimation equation is as follows:

$$Crash_{it} = \sum_{\tau=-12, \tau \neq -1}^{24} \beta_{\tau} Insp_{it}^{\tau} + u_i + \eta_t + \varepsilon_{it}, \quad (3)$$

In this event study, for each truck in the sample, I use only the inspections that occurred 12 months after its first inspection and 24 months before the last inspection. So trucks entering or exiting the market will not affect the estimation result.

The result of the monthly event study is shown in the top panel of Figure 10. There is no significant pre-trend in the 12 months before inspections. There is a significant increase in accidents during the first two months after an inspection/inspections in event month 0. After month 2, accidents gradually decline. In the longer term, the increase in accidents can be detected up to the 12th month after the inspection. The accident rate goes back to the pre-inspection level after 12 months and remains at that level until the 24th month. There is no compensating reduction in accidents in the 24 months after inspections. In the lower panel of Figure 10, I plot the probability of reinspec-

Figure 10: In longer term: the impact of an inspection on truck accidents



Note: In the upper panel (left y-axis), this plot shows the estimation result of the monthly event study that looks at the impact of inspections on truck accidents in the 12 months before to 24 months after an inspection (or inspections) in month 0 for any given truck. I control for individual truck, year, and month fixed effects. Standard errors are clustered at the truck level. In the lower panel (right y-axis), I plot the probability of reinspection for any given truck from month 1 to 24 after the inspection.

tion for any given truck from month 1 to 24 after the inspection. The probability of reinspection is only 4.0% in the first month after inspection, and it grows over time. The correspondence between the top and lower panels of Figure 10 shows that as the probability of reinspection increases over time, the driver’s compensating behaviors — and so the probability of a crash - decreases.

The total effect of one inspection on accidents is calculated by adding up the first 12 coefficients after inspections estimated from equation 3. As the crash rate goes back to the pre-inspection level after the 12th month, variation in accidents after month 12 is mostly noise that should not be counted. The total effect is 0.00103 additional accidents per inspection. The number of additional accidents is compared to the counterfactual scenario of having no change in accidents after inspections so that the accident rate remains at the same level before inspections. It is not compared to the counterfactual scenario of no inspection regulation. There are on average 1.75 million inspections conducted at weigh stations every year. Therefore, a back-of-the-envelope calculation of the total number of crashes caused by the behavioral responses to the inspection program is roughly 1803 (=0.00103*1.75 million). According to the National Safety Council’s cost calculation of motor vehicle injuries in 2018, the weighted average cost of a motor vehicle accident is around \$0.9 million. Hence, the total loss due to drivers’ compensating responses to the current inspection regulatory design is roughly \$1.6 billion per year.

6 Heterogeneity in inspection effects

6.1 Effects by inspection outcomes

To see how the effects vary with heterogeneous inspection outcomes, I compare the magnitudes of the increase in accidents following inspections with different types of violations found. In Table B.4, columns 1 and 2 represent two subsamples: those that do not find any violation (38%), and those that result in violations but not the most serious, out-of-service violations (41%).²⁶ These two groups exhibit similar increases in accidents after inspections, and they are comparable to the average effect size estimated using the full sample. In columns 3 and 4, I find that the post-inspection increase in accident rate among trucks that receive driver violations is larger than trucks that receive vehicle violations.

The results in Table B.4 indicate that following inspections, truck drivers reduce caution in driving, regardless of whether they receive – or do not receive – a violation. Inspections fail to deter future reckless driving behaviors because inspections lead to a lower reinspection probability for trucks already inspected, which contributes to more crashes. One possible explanation for the weak

²⁶I do not include trucks that receive the out-of-service violations (20%) in my estimation, because these trucks must either be repaired at scene, or must be parked immediately; as a result, they cannot experience any accidents after inspections in principle.

deterrent effect is because minor violations do not increase the probability of future inspection. According to the Safety Measurement System (SMS) score algorithm, carriers that receive violations will see an increased likelihood of inspections. However, note that the intervention threshold on SMS score is around 70-80%, which means that only the worst 20-30% trucks in rank will be targeted for more inspections. Because the out-of-service violation rate is 20%, those targeted carriers are likely to be those that receive these violations during inspections. Most trucks receive less severe violations might not be targeted in inspections. So, the drivers are not deterred by the violations.

6.2 Effects by firm characteristics

To examine how the effect of an inspection varies by firm characteristics, I perform a heterogeneity analysis by looking at the following four characteristics: firm sizes, interstate or intrastate commerce, major vehicle types, and cargo carried. The carrier/shipper census file contains information on the type of cargo transports, number of vehicles owned, and the number of drivers employed for 1.6 million active firms as of October 2018.

I look at whether drivers at large and small firms respond differently after inspections. Firm sizes are very heterogeneous in the trucking industry. Panel A in Table B.5 compares the effect of an inspection on crashes for trucks belong to firms with large or small numbers of power units.²⁷ It shows that the increase in accidents from smaller firms is slightly smaller than large firms. Panel B in Table B.5 shows that defining firm sizes using the number of drivers instead of power units gives a similar result.

The estimated effect sizes suggest that when the reinspection probability is low after inspections, drivers from both large and small firms decrease compliance. However, it does not suggest that large and small firms have the same incentives in doing so. There are as many reasons why large firms exhibit compensating behaviors as small firms. For examples, for 1-truck/1-driver carriers, the driver's earnings are tied with how much work he/she chooses to finish. So these drivers have more incentive to drive longer once their perceived inspection probability is lower. On the other hand, a large carrier, such as FedEx, could potentially have a principal-agent problem (Baker and Hubbard (2004)), which suggests that the drivers would not drive in the best manner to preserve the truck's value. Thus, drivers in large and small firms are living in quite different worlds with different profit structures and liabilities, which in turn affect their risk preferences when driving. In this paper, my ability to tease out the different constraints that firms face is limited by data availability.

²⁷I define a large firm as one having more than the median number of power units (or drivers) among all inspected firms; the small firms are those in the rest of the sample. So, I can split the full sample into two equally sized subsamples. A large firm has more than 48 power units (or more than 47 drivers). This number is larger than the number of power units for a median firm in the carrier census file. This is because large firms with more power units are more likely to receive inspections.

That does not suggest, however, that the dynamics are the same when facing inspections.

In addition to comparing effects across different firms sizes, I also find no remarkable differences surface according to other firm characteristics, including interstate or intrastate commerce, vehicle types, and the type of cargo transported in appendix G.

To summarize, all the findings reported above show that there is very *little* heterogeneity in the size of the effect of inspections on accidents. These results indicate that compensating behaviors exist for all types of drivers after they receive an inspection. Although drivers could potentially face different constraints, incentives and regulations, they all face a small reinspection probability following the current inspection; as a result they are likely to be less careful afterwards, resulting in an increase of accidents.

6.3 Effects by rest area distribution

I find that the presence of rest areas on the highways could reduce the increase in accidents following inspections. Safety rest areas, located at the roadside of highways, feature benefits including parking areas (for commercial trucks), restrooms, etc. Generally speaking, rest areas serve different functions from weigh stations, and the two are not at the same location. I collect location information on all the safety rest areas from Texas Department of Transportation. Figure A.14 shows that there are abundant counties that have weigh stations present but not rest areas.²⁸ There are on average 1.2 (s.d.=1) rest areas per county. Around half of all counties have rest areas.

This exercise uses the distribution of rest areas to see if truck drivers were provided with the opportunity to take a break when tired, or check at their vehicles after inspections, whether the increase in accidents caused by reckless driving and lack of maintenance could be smaller. I first confirm that the Texas sub-sample is representative of the whole country. Column 1 of Table B.9 validates that the Texas sub-sample produces similar baseline effect size (44.0%) as the national sample. Column 2 shows that with each additional rest area in the inspection county, the increase in accidents following inspections drops by 7.9%. Column 3 serves as a robustness test by including county-by-year fixed effects, and shows that the effect is not driven by other unobserved differences of counties with and without rest areas. The presence of rest areas gives truck drivers an option to stop if they feel tired and could no longer stay focused on the traffic. In addition, it gives drivers a second chance to check at their vehicles to make sure, for instance, that the cargo is secured on the truck, and prevent accidents caused by driver carelessness. This exercise suggest that, although drivers might exhibit compensating behaviors after inspections, rest areas serve as a buffer that lowers the number of accidents caused by drivers' reckless behaviors.

²⁸I collect geo-coded locations of rest areas so that I could plot them as points on the map, but the exact locations of weigh station are difficult to pin down. So I highlight the counties where weigh stations are present.

7 Policy recommendations

In this section, I explore two alternative policy options that could reduce the loss resulting from the sub-optimal design of the inspection program. Although the program is conducted nationwide, each state DOT chooses its own way of handling the enforcement strategy. The differences lie in the intensity and the type of inspections, the trucks selected, the inspection schedules, and the location of inspections. Distinct from *little* heterogeneity found across various inspection outcomes or firm characteristics, there is *large* variation in the effect of inspections on accidents across different geographical areas of the U.S. The spatial heterogeneity in effect sizes reveals differences in the regulatory designs across the states.

The mechanism behind the increase in accidents points to a sharp drop in the probability of reinspection following a recent inspection. Therefore, I first explore the difference in such probability across the states. I find that the re-inspection probability is an important driver of the spatial heterogeneity in effect sizes across states.

I calculated the reinspection probability the same way as in Section 5.1. Figure 4 suggests that the probability is around 8.9% at 13 weeks (a quarter) when averaging across all inspections in all the states. Here I separately compute that probability in each state by year. I estimate the impact of inspection on accidents depending on the differences in reinspection probability in a quarter following a recent inspection in the following way.

$$\begin{aligned} Crash_{it} = & \beta post_insp_{it} + \gamma post_insp_{it} \times Pr(reinsp.)_{s,yr} + \theta Pr(reinsp.)_{s,yr} \\ & + \beta_{-15} Insp_{it}^{-15} + \beta_{14} Insp_{it}^{14} + u_i + \eta_t + \varepsilon_{it} \end{aligned} \quad (4)$$

where $Pr(reinsp.)_{s,yr}$ is the average reinspection probability in a quarter for state s in year yr . Equation 4 is a variation of equation 2 by adding the interaction term $post_insp_{it} \times Pr(reinsp.)_{s,yr}$ and the main effect of $Pr(reinsp.)_{s,yr}$. Thus the coefficient of interest that estimates the effect of inspection by different reinspection probability is γ .

The estimation result is in Panel A of Table 3. The result shows that when the reinspection probability increases by 1%, the number of crashes following inspections decrease by 0.89%. The estimate is also robust when including additional controls, such as state fixed effects. The result is consistent with the mechanism found in the previous section that suggests the low reinspection probability leads to more accidents caused by drivers' strategic responses.

The reinspection probability calculated for each state in the exercise above is an outcome from different inspection designs, next I illustrate in detail below how differences in inspection schedules and truck selection criteria across states could influence safety outcomes. I show that if the inspections are designed to be more random, the inspection program could achieve better results.

Table 3: The impact of an inspection on crashes under two alternative regulatory designs

Panel A: Reinspection probability, equation 4	
post_insp	4.15*** (0.16)
Pr(reinsp. in a quarter)	1.28 (0.93)
post_insp \times Pr(reinsp. in a quarter)	-5.67*** (0.65)
Average crash rate	6.36
Panel B: Predictability in schedule, equation 5	
post_insp	2.83*** (0.06)
std. $pred_{s,yr}$	-0.07 (0.07)
post_insp \times std. $pred_{s,yr}$	-0.59*** (0.05)
Average crash rate	6.36
Panel C: Randomness in selection	
post_insp	2.83*** (0.06)
std. $Dtime_{s,yr}$	-0.07** (0.09)
post_insp \times std. $Dtime_{s,yr}$	0.69*** (0.06)
Average crash rate	6.36

Note: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. $post_insp$ is a post-inspection indicator, it equals to 0 before an inspection happen on truck i , and it equals to 1 on and after the inspection. $std_pred_{s,yr}$ is the predictability index of inspection schedule for state s in year yr , standardized to mean 0 and standard deviation 1. $std_Dtime_{s,yr}$ is the average reinspection time interval calculated for state s in year yr , standardized to mean 0 and standard deviation 1. All regressions include dummies for inspection happened outside of the time window of interest, as well as individual truck, year, month, and day-of-week fixed effects. Standard errors are clustered at the truck level.

7.1 Predictability of inspection schedules

In the first alternative inspection policy, I explore how the predictability in inspection schedule affects drivers' behavioral responses. Under highly predictable inspection schedules, drivers could easily forecast if there will be open inspection sites when travelling at a given state on a particular day. On this margin, I find that states with most **un**predictable inspection schedules could achieve crash reduction among trucks inspected after inspections; while states with very predictable schedule have larger increase in accidents after inspection.

More specifically, I use the variation in the day-of-week inspection schedules across states to measure the predictability of inspection schedules. In the illustrative example given in Figure 11, I compare the number of inspections done in each day of the two weeks in State A and B. The average number of inspections per day is the same (mean = 58) for both states. State A has a fixed day-of-week inspection schedule, meaning that State A is conducting 10 inspections every Sunday, 60 every Monday, 80 every Tuesday, etc.; while State B has a random day-of-week inspection schedule so that the number of inspections done on Sunday this week is different from the next Sunday, this Monday is different from the next Monday, etc. Therefore, in State A, it will be very easy for truck drivers who travel the same route frequently to predict whether they will receive an inspection on a given day of the week. But it will be hard for them to predict in State B. As a result, drivers will always be cautious when driving in State B since they are uncertain of the inspection intensity on a particular day, but they could driver more recklessly or spend less time in pre-trip checks in State A if they know that they are very less likely to get an inspection today, or even for the rest of the week.

I develop a predictability index to measure the variation in the day-of-week inspection schedule for each state annually. Appendix H shows how the index is constructed. Figure 12 shows the yearly average of the state \times year predictability index across all states. The index captures the variation in the inspection schedule after accounting for fixed patterns in the schedule at the county-by-day-of-week-by-year level. So, the predictability index is not driven by different inspection frequencies, traffic and road conditions across states.²⁹ The lower the predictability index, the easier the drivers could predict an inspection. In the two-state example in Figure 11, State A has a predictability index of 0, State B has a predictability index of 32.

I estimate the impact of inspection on accidents depending on the differences in inspection schedule predictability in the following way.

$$\begin{aligned} Crash_{it} = & \beta post_insp_{it} + \gamma post_insp_{it} \times std_pred_{s,yr} + \theta std_pred_{s,yr} \\ & + \beta_{-15} Insp_{it}^{-15} + \beta_{14} Insp_{it}^{14} + u_i + \eta_t + \varepsilon_{it} \end{aligned} \quad (5)$$

²⁹For example, in appendix table B.7, I test that the effects from different predictability in schedules on post-inspection accidents are not affected much after controlling for the total number of inspections a state does in a year.

Figure 11: Policy option 1: the variation in the day-of-week inspection schedule





Figure 12: Policy option 1: the predictability index



Note: The set of plots in Figure 11 shows an example of two states that have very different day-of-week inspection schedule in terms of the number of inspections done in each day of the two weeks. In order to measure the different predictability of the day-of-week inspection schedule, as in State A and B, I develop a predictability index *pred_index*. The lower *pred_index*, the easier the drivers could forecast an inspection. Here, *pred_index* = 0 for state A, which indicates that the schedule is highly predictable; *pred_index* = 32 for state B, which indicates that the schedule is not quite as predictable.

Figure 12 shows the yearly average of the state×year predictability index (*pred_index*) across all states. In this figure, lighter colors indicate a lower *pred_index* and a higher predictability in inspection schedule, while darker colors indicate a higher *pred_index* and a lower predictability.

where $std_pred_{s,yr}$ is the standardized predictability index for state s in year yr . $std_pred_{s,yr}$ is mean 0, standard deviation 1. Equation 5 is of the same format as equation 4 by replacing $Pr(reinsp.)$ with $std_pred_{s,yr}$. The estimation result of equation 5 is in Panel B of Table 3.

I then calculate the counterfactual effect sizes from this policy exercise through interpolation using the estimates from Table 3. Under highly predictable inspection schedules, or low predictability indices, accident rates increase by 52% following inspections; under highly unpredictable inspection schedules, or high predictability indices, the effect size of inspection on accidents is -6%.³⁰ One standard deviation increase in the predictability index, or decrease in predictability, *drops* the effect size by 9%. In other words, if states were to adopt highly unpredictable inspection schedules, they could achieve crash reduction after the inspections.

7.2 Randomness in inspection selection

Another policy variation I explore is the extent to which the inspectors rely on past inspection timing to determine which trucks are chosen for inspections. That is, are recently inspected trucks being subject to further inspection to the same degree as those that have not received inspections – as if all were selected randomly. This exercise simulates the situation when regulators only have limited resource to conduct selective inspections, and looks at what would be a better selection mechanism considering drivers' compensating behaviors. I compare states with longer and shorter time intervals from previous inspections, and find that if states choose to inspect trucks at greater randomness, they could reduce the increase in accidents.

I illustrate the policy exercise with an example of two states, A and B, in Figure 13. In this figure, the horizontal axis is the number of days passed since the most recent inspection. The gray bars represent all trucks enter into the weigh stations in the two states today with different past inspection timing. The red lines represent the trucks eventually selected for inspection. Inspectors in State A only choose trucks that have not been inspected for a long time, which resembles the current regulatory design in many states, while inspectors in State B randomly choose trucks across all timing. The two states choose the same total number of trucks for inspection. I then calculate the average reinspection time interval for all trucks selected for inspection in both states. State A has a longer average reinspection interval than State B. Therefore, shorter reinspection intervals imply more randomness in the state's inspection selection mechanisms.

I calculate the average reinspection time interval, $Dtime_{s,yr}$. It is defined as the average number of days that have passed since last inspection (in any state) for trucks inspected in state s in year yr . Figure 14 shows that states' approaches vary widely. California has an average reinspection interval that is around 4 months, but Michigan has an average reinspection interval that is longer

³⁰The increase in accident rates in the two cases are calculated using the lowest and highest 1% in the distribution of $std_pred_{s,yr}$, which is 0.84 and 5.53, respectively.

Figure 13: Policy option 2: differences in inspection selection based on past inspection timing





Figure 14: Policy option 2: time interval between consecutive inspections



Note: In Figure 13, the horizontal axis is the number of days passed since the most recent inspection. The gray bars represent all trucks enter into the weigh stations in the two states today with different past inspection timing. The red (darker) lines represents the trucks eventually selected for inspection.

Figure 14 shows the variation of the average reinspection time interval for trucks inspected in a state. Lighter colors in this figure indicate a shorter reinspection interval.

than 1 year.³¹ Trucks travelling in California today face a higher chance of reinspection – even if they were only recently checked – while those in Michigan do not. As a result, the increase in accidents due to drivers’ offsetting behaviors after inspections is less severe in California than in Michigan.

I estimate the effect of inspection on accidents by different reinspection time intervals using the same framework as in equation 4 and 5, by replacing $\Pr(reinsp.)$ with $std_Dtime_{s,yr}$. $std_Dtime_{s,yr}$ is the reinspection time interval of state s in year yr , standardized to mean 0 and standard deviation 1. The estimated result is in Panel C of Table 3.

I calculate the counterfactual effect sizes with various lengths of reinspection time interval using estimates from Panel C of Table 3. When reinspection time intervals are short (less than 1 month), there is a 25% increase in accident rates post inspections; when reinspection time intervals are long (longer than 1 year), there is a 64% increase in accident rates post inspections. A one standard deviation increase in the time interval (5 months) *raises* the effect size by 9%. All above estimations of effect sizes are based on the current sample. Hence, one could extrapolate from the estimates that greater reductions in accidents could be achieved if states choose to inspect trucks at greater randomness.

8 Conclusion

This paper provides one of the first evaluations of the commercial vehicle roadside safety inspection program on trucks inspected across the US over more than two decades. The primary contribution to the literature of regulatory enforcement design is to show that regulatory efforts to improve public safety could be undermined when people respond strategically to a temporary reduction in the probability of detection. The analysis uses the most comprehensive data available on trucks, inspections, and accidents, linking inspection and accident histories to each truck’s unique vehicle identification number. I find that there is a sharp, 44.6% *increase* in a truck’s accident rate immediately following an inspection, and that the effect lasts for at least 14 days. In the longer term, the increase in accident rates persists for 12 months following inspections. Thus, the loss resulting from the design of the current inspection program is \$1.6 billion per year.

The increase in accident rates is attributable to truck drivers’ strategic responses to a low probability of reinspection after a recent inspection. Assuming that the truck will not be reinspected in the near term, the driver might conduct fewer pre-trip safety checks and drive more recklessly,

³¹Although California has more inspections than Michigan, it does not seem to confound the effect coming from the variations in the length of reinspection intervals. In appendix table B.7, I test that by including the total number of inspections a state does in a year as a control variable in the regression. The result indicates that when all states conduct the same number of inspections, if a state has a shorter reinspection interval, that state will experience more reductions in accidents.

resulting in an increase in accidents caused by speeding, driving while intoxicated, and equipment failures. Such compensating behaviors offset the potential benefits of the safety program.

I explore two alternative policy options that could reduce existing accident rates following inspections. First, states could adopt a more random inspection schedule so that drivers could not anticipate inspection days. Second, states could adopt policies that randomly reinspect trucks regardless of the timing of their most recent inspection.

This paper should not be interpreted as an argument for abandoning the commercial motor vehicle roadside safety inspection program. To the contrary, as implemented, the program removes 20% of inspected commercial motor vehicles with severe violations from the road each year. Without the inspection program, these more dangerous vehicles and drivers would likely remain on the road, leading to more serious safety problems. The findings of the paper instead suggest that the current enforcement mechanism of inspections is sub-optimal, and that changes are needed to improve public safety.

In response to drivers' strategic behaviors, this paper also calls for the integration of more advanced technologies and better monitoring practices from private companies when designing safety regulations. One option is to utilize widely adopted on-board computer systems ([Baker and Hubbard \(2004\)](#)), or the newly implemented electronic logging devices, which can be used as monitoring mechanisms to enforce better driving behaviors both before and after inspections. It would also be interesting to explore the effectiveness of several management practices on truck drivers' performances by conducting a field experiment inspired by [Gosnell et al. \(2020\)](#).

Overall, this paper provides compelling evidence that drivers respond strategically to changes in the probability of detection that are caused by regulation design. Although this study is in the context of transportation safety, it highlights the general importance of understanding compliance behaviors when designing enforcement mechanisms. The lesson we learn from this study also applies to many other regulatory settings where agents could change behaviors between repeating instances of enforcement, such as audit strategies, workplace safety regulations, and criminal deterrence. Well designed and consistently enforced regulatory regimes should aim to increase the detection probability in order to reduce violations.

References

- Adda, Jerome and Francesca Cornaglia**, “Taxes, cigarette consumption, and smoking intensity,” *American Economic Review*, 2006, 96 (4), 1013–1028.
- Anderson, Michael L and Maximilian Auffhammer**, “Pounds that kill: The external costs of vehicle weight,” *Review of Economic Studies*, 2014, 81 (2), 535–571.
- Baker, George P and Thomas N Hubbard**, “Contractibility and asset ownership: On-board computers and governance in US trucking,” *The Quarterly Journal of Economics*, 2004, 119 (4), 1443–1479.
- Banerjee, Abhijit, Esther Duflo, Daniel Keniston, and Nina Singh**, “The efficient deployment of police resources: theory and new evidence from a randomized drunk driving crackdown in India,” *National Bureau of Economic Research*, 2019.
- Becker, Gary S**, “Crime and punishment: An economic approach,” in “The economic dimensions of crime,” Springer, 1968, pp. 13–68.
- Blundell, Wesley, Gautam Gowrisankaran, and Ashley Langer**, “Escalation of scrutiny: The gains from dynamic enforcement of environmental regulations,” *American Economic Review*, 2020, 110 (8), 2558–85.
- Chaisemartin, Clément De and Xavier d’Haultfoeuille**, “Two-way fixed effects estimators with heterogeneous treatment effects,” *American Economic Review*, 2020, 110 (9), 2964–96.
- Cohen, Alma and Liran Einav**, “The effects of mandatory seat belt laws on driving behavior and traffic fatalities,” *Review of Economics and Statistics*, 2003, 85 (4), 828–843.
- Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan**, “The value of regulatory discretion: Estimates from environmental inspections in India,” *Econometrica*, 2018, 86 (6), 2123–2160.
- GAO, US**, “LARGE TRUCK SAFETY: Federal Enforcement Efforts Have Been Stronger Since 2000, but Oversight of State Grants Needs Improvement,” *GAO-06-156*, 2005.
- Gonzalez-Lira, Andres and Ahmed Mushfiq Mobarak**, “Slippery Fish: Enforcing Regulation when Agents Learn and Adapt,” Technical Report, National Bureau of Economic Research 2021.
- Gosnell, Greer K, John A List, and Robert D Metcalfe**, “The impact of management practices on employee productivity: A field experiment with airline captains,” *Journal of Political Economy*, 2020, 128 (4), 000–000.
- Graham, Jove, Jennifer Irving, Xiaoqin Tang, Stephen Sellers, Joshua Crisp, Daniel Horwitz,**

- Lucija Muehlenbachs, Alan Krupnick, and David Carey**, “Increased traffic accident rates associated with shale gas drilling in Pennsylvania,” *Accident Analysis & Prevention*, 2015, 74, 203–209.
- Gray, Wayne B and Jay P Shimshack**, “The effectiveness of environmental monitoring and enforcement: A review of the empirical evidence,” *Review of Environmental Economics and Policy*, 2011, 5 (1), 3–24.
- Johnson, Matthew S, David I Levine, and Michael W Toffel**, “Improving Regulatory Effectiveness Through Better Targeting: Evidence from OSHA,” *Harvard Business School Technology & Operations Mgt. Unit Working Paper No. 20-019*, 2020.
- Keeler, Theodore E**, “Highway safety, economic behavior, and driving environment,” *The American Economic Review*, 1994, 84 (3), 684–693.
- Kwigizile, Valerian, Andrew H Ceifetz, Jun-Seok Oh, Joyce L Yassin, and Jason Firman**, “Benefit-Cost Analysis of Fixed Weigh Stations: The Michigan Case,” *Transportation Research Board 95th Annual Meeting*, 2016, pp. 16–4434.
- Li, Shanjun**, “Traffic safety and vehicle choice: quantifying the effects of the ‘arms race’ on American roads,” *Journal of Applied Econometrics*, 2012, 27 (1), 34–62.
- Loeb, Peter D and Benjamin Gilad**, “The efficacy and cost-effectiveness of vehicle inspection: a state specific analysis using time series data,” *Journal of Transport Economics and Policy*, 1984, pp. 145–164.
- Lv, Jinpeng, Dominique Lord, Yunlong Zhang, and Zhi Chen**, “Investigating Peltzman effects in adopting mandatory seat belt laws in the US: Evidence from non-occupant fatalities,” *Transport Policy*, 2015, 44, 58–64.
- Muehlenbachs, Lucija, Stefan Staubli, and Mark A Cohen**, “The impact of team inspections on enforcement and deterrence,” *Journal of the Association of Environmental and Resource Economists*, 2016, 3 (1), 159–204.
- , —, and **Ziyan Chu**, “The accident externality from trucking: Evidence from shale gas development,” *Regional Science and Urban Economics*, 2021, 88, 103630.
- Okat, Deniz**, “Deterring fraud by looking away,” *The RAND Journal of Economics*, 2016, 47 (3), 734–747.
- Oliva, Paulina**, “Environmental regulations and corruption: Automobile emissions in Mexico City,” *Journal of Political Economy*, 2015, 123 (3), 686–724.
- Peltzman, Sam**, “The effects of automobile safety regulation,” *Journal of political Economy*, 1975,

83 (4), 677–725.

Sun, Liyang and Sarah Abraham, “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 2020.

Viscusi, W Kip, “The lulling effect: the impact of child-resistant packaging on aspirin and analgesic ingestions,” *The American Economic Review*, 1984, 74 (2), 324–327.

—, “The impact of occupational safety and health regulation, 1973-1983,” *The RAND Journal of Economics*, 1986, pp. 567–580.

Zou, Eric, “Unwatched pollution: The effect of intermittent monitoring on air quality,” *URL: <https://static1.squarespace.com/static/56034c20e4b047f1e0c1bfca>*, 2018, 5, 1535772291252.

Appendices

A Inspection outcomes

Different violations are given for problems found during the inspections based on criteria created by the Commercial Vehicle Safety Alliance (CVSA). Multiple violations could be issued if multiple issues were found. Among all inspections conducted at weigh stations, 41% result in at least one violation but not in a violation so severe that it leads the vehicle to be taken out of service. For those that result in violations, 40% are assigned to drivers, 75% are assigned to vehicles, and 1% for hazmat carriers.³² In general, driver violations can incur a higher fine than vehicle violations. Driver violations can result in the suspension of drivers' licenses – which likely serves as more of a deterrent for violating regulations than vehicle violations. During inspections, trucks with violations that do not result in the vehicle to be placed out-of-service are still allowed to continue operating, but any violations or defects noted at the time of inspection must be corrected within 15 days of receiving the violation.

In all, 20% of inspections result in out-of-service violations, indicating that the vehicle or driver presents an imminent hazards to the public. The effect is immediate; the vehicle may not be driven until all necessary repairs are made, and all the violations are corrected.

Another 38% of inspections do not find any violations. If no critical violations are assigned to the vehicle during a thorough vehicle inspection, then the vehicle could be issued a CVSA decal. Among all inspected vehicles, 16.6% receive a decal. Generally, a vehicle displaying a valid CVSA decal will not be reinspected during the three-month time frame in which the decal is valid. Note that the CVSA decal is only valid for the vehicle, so the driver could still be subject to an inspection.

³²An inspection could result in both driver and vehicle violations.

B Data

Carrier census data. The carrier company census file from FMCSA provides a snapshot of all active operating carrier companies in the US at the time of a request. The company census file contains information on the address, registration, type of cargo transported, number of vehicles owned, and number of drivers employed. This data file is useful in particular to analyze the heterogeneity of drivers' responses to the inspection program among different carrier companies. In Section 6, I compare the impact of an inspection on accidents for large and small companies defined by power units or the number of drivers, and for companies carrying different cargo types, and whether they conduct interstate or intrastate business. Table 1 panel D shows that there were 1,669,661 active companies registered as of October 2018. On average, each carrier company employs 5 drivers and owns 21 trucks. Note that a median-size carrier company only have one driver and one truck. So, half of the carrier companies are very small, but there are also a small number of giant carrier companies in the industry.

Other data sets. Supplementary data sets assembled for the analysis include the traffic monitoring and traffic volume data from Federal Highway Administration (FHWA), the daily weather records from Global Historical Climate Network Daily (GHCN-Daily). Traffic monitoring data contain the traffic volume records by vehicle class at hourly frequencies for the participating states from 2012 to 2018. This data set allows me to directly observe truck traffic volume, and to use traffic volume of other types of vehicles as controls for the analysis. GHCN-Daily weather data provide weather records including maximum and minimum temperatures, and total daily precipitation, snowfall, and snow depth. I use this data set to test whether adverse weather conditions lead to more accidents after inspections.

Analysis sample construction. The inspection and crash files are the two main data sets combined to create an truck-inspection-crash daily panel that traces all accidents that happened to the same truck inspected within the time frame of interest. More specifically, since there are 11 million trucks inspected at the fixed weigh stations, it would involve too much computational burden to create a balanced panel for each truck at the daily level for 23 years. The number of observations would be larger than 50 billion. Instead, I construct a 28-day event window around each inspection for any given truck so that I can compare events in the 14 days before to the 14 days after the inspection. I then identify all accidents for the same truck that receives an inspection within the 28-day event window using the VIN in both files.³³

³³The sample period of this paper is from 1996 to 2018, some trucks may change license plate number in between, so VIN is the best identifier for the same vehicle. However, VIN is missing for many inspection records from 1996 to 2009, so license plate numbers are used to identify a truck if VIN is missing. After 2010, 90% of the inspections have VIN records.

C Empirical analysis: identification


C.1 Weather and traffic as potential confounders

In this appendix, I show how I rule out weather and traffic conditions as confounders to the identification. I cannot directly control for weather or traffic conditions in the regressions because I do not observe the exact location of the trucks at any times other than the time of inspections and the time of any accident that occurs. I use the following ways to rule out the two confounders:

I rule out that weather conditions are confounders by showing that the inspections are not chosen at times of worse weather conditions (such as rain or snow). On the contrary, I find that the inspections take place on days with better weather conditions. Therefore, even if adverse weather conditions do lead to higher probabilities of accidents, the observed *increase* in accidents after an inspection is not caused by adverse weather conditions. For example, imagine that the inspectors operating the weigh station in Tompkins County in New York state decide which day to conduct inspections according to the weather forecast for the entire following week.³⁴ I test whether the inspections are chosen at better or worse weather conditions by looking at an event study of inspections in relation to weather conditions in the county where the weigh station is located, for instance, a rain indicator for Tompkins County. Figure A.1 shows the relationship between inspections and the probability of raining. There is a small but quite significant 1.5% drop in the probability of raining on the day of the inspection ($t = 0$) comparing to days before or after. Because it is certain that inspections do not cause rain, the relationship has to be the other way around; that is, inspections are chosen at days when rain is less likely to occur. The slightly positive coefficients on days before inspection also suggest that inspectors choose to do inspections on a sunny day following rainy days. The results are similar when looking at snowy days. In addition, I do not find that temperatures affect the inspection schedule.

³⁴In this exercise, I pre-assume that weather conditions are going to affect the inspection schedules. If the assumption does not hold, which means that inspection schedules are set as stone so weather conditions are not correlated with inspection schedules, then the weather confounders can be ruled out as well.

Figure A.1: Ruling out weather condition as a confounder



Note: This figure shows the relationship between inspections and a rain indicator in the county where the inspections take place. The coefficients plotted in this figure are estimated using an event study setting where I regress a rain indicator of the county where the inspection takes place on a set of inspection indicators from 14 days before an inspection to 14 days after inspection, controlling for individual truck fixed effects, year, month, and day-of-week fixed effects. The shaded area is the 95% confidence interval for the estimates. The effect of an inspection on crash accidents happening in the two-day bin $(-2,-1)$ is normalized to 0.

I can also rule out higher traffic volume as a potential cause of the increase in accidents after an inspection. I do this in the following ways: I control for the traffic conditions in the inspection county at the inspection hour as shown in Table B.1 in the appendix. The first column prints the baseline estimation using equation 2 for the sample period (2012-2018).³⁵ To measure traffic conditions more comprehensively, I look at both the truck counts (in column 2 and 3) and the total traffic volume of all motor vehicles (in column 4). The post-inspection coefficient in Column 2 is almost the same as those in Column 1, indicating that the effect of inspections on accidents stays the same after controlling for traffic volume. Then, in Columns 3 and 4, I interact traffic conditions with the post-inspection indicator to see whether the effect size depends on traffic volume. When using the number of truck counts to measure traffic condition as in Column 3, the coefficient on interaction term ($\text{post_insp} \times \text{truck_count}$) is not significant at 5% level. The effect size is compara-

³⁵ Because the traffic volume data are available only from 2012 onwards, the event study period is from 2012 to 2018 in this case.

ble to Columns 1 and 2. When using the total traffic volume as another measure of traffic condition in Column 4, the coefficient on the interaction term ($\text{post_insp} \times \text{total_traffic}$) is significantly negative.³⁶ It suggests that if a truck gets inspected when the total traffic volume is high, the increase in accidents is smaller after the inspection. Therefore, combining all evidences, I can rule out that heavy traffic volume is contributing to the increase of accidents after inspections.

Table B.1: The impact of an inspection on crashes, controlling for the traffic volume

Dependent var: number of crashes (per 100,000 trucks)	(1)	(2)	(3)	(4)
<i>post_insp</i>	3.20*** (0.193)	3.20*** (0.193)	3.27*** (0.243)	3.54*** (0.248)
<i>truck_count</i>		-.00005 (.00018)	.0001 (.00019)	
<i>post_insp</i> × <i>truck_count</i>			-.0001 (.0002)	
<i>total_traffic</i>				.00002 (.00002)
<i>post_insp</i> × <i>total_traffic</i>				-.00004** (.00002)
Average crash rate	9.08	9.08	9.08	9.08
Effect size	35.24%	35.24%	36.01%	37.40%
Observations	96,841,680	96,841,680	96,841,680	96,841,680

Note: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The sample period in this exercise is from 2012-2018 since the traffic volume data is available from 2012 onwards. All columns use the same estimation framework following equation 2, which looks at the impact of an inspection on crashes by comparing 14 days before to 14 days after inspection. All regressions include dummies for inspection happened outside of the time window of interest, as well as individual truck, year, month, and day-of-week fixed effects. Standard errors are clustered at the truck level. *post_insp* is a post-inspection indicator, it equals to 0 before an inspection happen on truck *i*, and it equals to 1 on and after the inspection. *truck_count* is the number of trucks passed by in the same hour and county as when the inspection happened. *total_traffic* is the total traffic volume of all motor vehicles in the same hour and county as the inspection. The median of total traffic volume is 3,422, with standard deviation equals to 10,236.

³⁶The median of total traffic volume is 3,422, with standard deviation equals to 10236. When the total traffic volume increases by 1 standard deviation from the median, the effect size decreases by 4.5%.


C.2 Eliminate reverse causality

For the identification assumption to hold, inspections induced by accidents should be excluded in the estimation sample. The issue warrants attention in this particular case because it is not uncommon for a police officer to call a nearby inspector to check a truck that has been involved in an accident. In this case, there is reverse causality because the accident itself invites the inspection. To eliminate that concern, I drop all trucks that are inspected within 18 hours after accidents.

In Figure A.2, the slight increase in accident rate on day (-2, -1) without dropping any inspections reveals the fact that an accident could invite an inspection, which causes a reverse causality problem. The rationale of solving the problem by dropping trucks inspected right after accidents is the following. In principle, an inspection happening on day 0 should not have any differential effect on accidents one day or two days before. Even if the truck drivers anticipate that an inspection is going to happen in the near term, they would not be able to know for sure whether there will be an inspection one or two or three days later. Thus, their behaviors in the days approaching the inspection should be similar. After varying the length of the time interval during which the inspections are dropped after accidents to see what the proper time frame would be, I decide to use 18-hour as the time frame; I select this because the coefficient on the (-2,-1) bin is consistent with that on (-4,-3) bin. Figure A.2 shows that no other coefficients are affected when varying the length of the time interval - except for the coefficients on days (-2, -1) and (0, 1), which drop as the length gets longer.³⁷ Therefore, in any case, the choice of the time interval would not affect the impact of the inspection on accidents when comparing the accident rates further back in the pre-period with those in the post-period.

³⁷In Figure A.2, I normalize the coefficient on days (-14,-13) to be 0 in order to compare across samples

Figure A.2: Eliminating the reverse causality



Note: This figure shows that by varying the length of the time interval during which the inspections are dropped after crash accidents, the coefficients on two-day bin $(-2, -1)$ and $(0, 1)$ drop as the length gets longer, but it does not affect any other coefficients. The first line marked as "all" is the full sample without dropping an inspections. Then I look at the choice of dropping inspections within 3, 6, 12, 18, and 24 hours of crash. I normalize the effect of an inspection on crashes on day $(-14, -13)$ to be 0 in order to compare across samples. The coefficients plotted in this figure are estimated using equation 1 where I regress the number of crash accidents for the same truck that receives the inspection on a set of inspection indicators from 14 days before an inspection to 14 days after inspection, controlling for individual truck fixed effects, year, month, and day-of-week fixed effects.


C.3 Treatment effect heterogeneity

To address potential treatment effect heterogeneity across different trucks inspected at various timing (De Chaisemartin and d'Haultfoeuille (2020)), I compute the weights underlying the linear combination of treatment effects following the framework developed in Sun and Abraham (2020). I find that there is little treatment effect heterogeneity, and the parallel pretrend assumption holds

in this context.

I randomly selected 1 percent trucks out of all trucks, and then group them into 23 cohorts according to their year of inspection respectively. In order to save computation time, I further restricts the event window to $(-3,3)$ around the inspection date for each truck because the parallel trend test falls on the time indicator on -2 of the inspection date. I follow [Sun and Abraham \(2020\)](#) in performing the regression that assigns weights on each truck cohort average treatment effect that help gauge the amount of treatment effect heterogeneity. Figure A.3 replicates Figure 2 in their paper with a few modifications according to the structure of my dataset. The figure below shows that the weights are almost zero for lags of treatments $\tau = 1, 2, 3$, which suggests that the pretrend, as indicated in β_{-2} , is not sensitive to the effect of inspections, and there is little treatment effects heterogeneity in $\tau = -3, 0, 1, 2, 3$ from different cohorts.

Figure A.3: Estimating cohort-specific weights for event study following [Sun and Abraham \(2020\)](#)



Note: Estimated weights underlying β_{-2} . Each curve represents the weights on a specific truck cohort average treatment effect.

D Calculate the probability of reinspection

This appendix explains how the probability of reinspection plotted in Figure 4 is constructed. The probability of reinspection represents the probability of getting another inspection within N weeks after the most recent inspection. I first generate a list of inspection history for each one of the trucks inspected in the sample. I then calculate the time interval between the current inspection and the closest next inspection in the future for a given truck. As for each truck's last inspection in the sample, since there is no future inspections, I treat the time interval to be longer than 1 year. Next, I calculate the percentage of reinspections that were conducted within 1,2,3 to 52 weeks out of all inspections received for a given truck; then take the average of all trucks across the 52 weeks. Note that this is the same as calculating the percentage of all trucks that were reinspected within 1 to 52 weeks.

E Crash external conditions

In section 5.1, I find a larger percentage increase in single-vehicle accidents following an inspection comparing to multi-vehicle accidents. It indicates that the reason for the increase of accident on average is attributable to the truck driver's behaviors following an inspection. Furthermore, I find that the increase in single-vehicle accidents is even higher when the external conditions are worse. This suggests that the increase is caused by drivers paying less attention to driving conditions.

Crash external conditions include road surface conditions, weather conditions and light conditions. Table B.2 compares the impact of inspections on crashes under normal conditions versus the impact under adverse conditions. I categorize all single-vehicle crash accidents into crashes under normal conditions and crashes under adverse conditions. Crashes under normal conditions are crashes happening on dry road surface conditions, with normal weather conditions, and under daylight. They represent 39.7% of all single-vehicle crashes. Crashes under adverse conditions are crashes happening when either of the three conditions are not as good as previous listed normal conditions. They represent 55.7% of all single-vehicle crashes. Column 2 shows that crashes under normal conditions increase by 60.9% following an inspection, and column 3 shows that crashes under adverse conditions increase by 83.7% following an inspection. Since when external conditions are bad, drivers need to pay more attention to the traffic than usual in order to prevent an accident. The larger increase in single-vehicle accidents under adverse conditions reveals that drivers drive less carefully after they undergo an inspection, which then leads to even more accidents.

F Other potential mechanisms

In this appendix, I address some alternative mechanisms that could also affect drivers' behaviors after inspections. First, I discuss the possibility of drivers feeling safer after passing inspections that causes them to reduce caution or seek risky behaviors. Second, I look at the incentives for speeding to make up for the time lost due to the weigh station inspections.

F.1 Peltzman effect

There is an alternative mechanism for the increase in accidents after inspections related to the Peltzman effect (Peltzman (1975)). After passing inspections, the drivers perceive that the result indicates that their trucks are safe. They feel safer driving the trucks, so they might reduce caution towards the traffic or the vehicle conditions, or push the trucks to the limit. It could lead to an increase of accident after inspections. It is analogous to the situation with the compulsory seat belt law described in Sam Peltzman's framework. He first points out the offsetting effect of drivers' responses to the seat belt law in reducing the highway death rate. Drivers feel safer wearing seat

Table B.2: The impact of an inspection on crashes under different crash external conditions

Dependent var: number of crashes (per 100,000 trucks)	(1) All single-vehicle crashes	(2) Normal conditions	(3) Adverse conditions
post_insp	1.64*** (0.06)	0.53*** (0.02)	1.03*** (0.03)
Average crash rate	2.19	0.87	1.23
Effect size	74.89%	60.92%	83.74%
Observations	670,402,180	670,402,180	670,402,180

Note: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All three columns use the same estimation framework following equation 2, which looks at the impact of an inspection on crashes by comparing 14-days before to 14 days after inspection. *post_insp* is a post-inspection indicator, it equals to 0 before an inspection happen on truck *i*, and it equals to 1 on and after the inspection. All regressions include dummies for inspection happened outside of the time window of interest, as well as individual truck, year, month, and day-of-week fixed effects. Standard errors are clustered at the truck level. Column 1 prints the baseline estimation using all single-vehicle crashes. Column 2 selects single-vehicle crashes under normal external conditions (all three external conditions are normal). Column 3 selects single-vehicle crashes under adverse external conditions (any one of the three external conditions is worse).

belts so they tend to drive more recklessly on the road, which then create safety problems.

In this paper, the empirical evidence I find also supports the Peltzman effect explanation. The data does not empower me to clearly tease out whether drivers are responding to the reduced risk of accidents because inspections provide information about the safety of the trucks, or drivers are responding to the reduced risk of detection because the reinspection probability is very low in the near future. From the existing evidences I find, I think both mechanisms are possible in affecting drivers' decisions. From the regulator's point of view, we can change regulatory designs so that drivers would not expect the reinspection probability to drop sharply after a recent inspection. In Section 7, I leverage the differences in regulatory designs across states to show that drivers indeed respond to variations in the reinspection probability.

F.2 Making up for time lost

The truck drivers may have an incentive to make up for the time lost during inspections in the trip afterwards. An inspection at weigh station could take from 5 minutes to 1 hour depending on the scrutiny and procedures performed. If there is a line waiting to enter the station, then the total time spent at the weigh station could be even longer. Because drivers are mostly paid per mileage, and carrier companies incur a relatively significant cost for delays in delivery. The time lost during inspections might cause the drivers to speed up in the trip following the inspections, which leads to more accidents.

Empirically, I do not find a significant evidence that shows that drivers' speeding behaviors after inspections are caused by the time lost during inspections. The increase in accident rate lasts for at least 14 days after the inspection. Even if drivers have incentive to make up for the time lost during inspections, it would not be affecting their driving behaviors in such a long period after inspections. It suggests that the major factors that influence drivers' driving behaviors is the low reinspection probability that persists in a long time after inspections.

If there is any effect from speeding to make up for the time lost, I expect it to show up in the first few days after the inspections. To see how the time lost at weight stations affects accident rates after inspection in a short period of time, I conduct an hourly event study that looks at the accident rate in the 48 hours before to 48 hours after the inspection at 4-hour interval. I focus on single-vehicle crashes since they are more representative of the truck drivers' behaviors. I use the time lapse between the inspection start time and end time to calculate the inspection duration for each inspection. I distinguish between short and long inspection duration by the median duration in Table B.3. Comparing the two groups, first, it shows that the average accident rate for trucks that have long inspection durations is higher than trucks that have short inspection durations. Second, the percentage increase in accident rate is comparable between the two groups. The result does not support the argument that there are more accidents following inspections for trucks that have

a longer inspection duration. However, we know that the inspection durations are not distributed randomly across trucks, so trucks with a longer duration might have more serious safety problems, which shows up as a higher average accident rate for trucks in this group. Therefore, the comparison of the two groups cannot be served as a causal evidence. In summary, I cannot completely rule out that drivers have incentives to make up for time lost during inspections, but it is not likely to be the major contributing factor to the increase in accident rate following inspections.

Table B.3: The impact of an inspection on hourly crashes

	(1) Avg duration = 13min Single-vehicle crashes (per 100,000 trucks)	(2) Avg duration = 36min Single-vehicle crashes (per 100,000 trucks)
<i>post_insp</i>	0.39*** (0.02)	0.53*** (0.03)
Average crash rate	0.38	0.46
Effect size	102%	115%
Observations	289,458,950	306,280,275

Note: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Both columns use the same estimation framework following equation 2, which looks at the impact of an inspection on crashes by comparing 48 hours before to 48 hours after inspection at each 4-hour interval. *post_insp* is a post-inspection indicator, it equals to 0 before an inspection happen on truck i , and it equals to 1 on and after the inspection. So the coefficient of *post_insp* should be interpreted as the effect of inspection on crash in every 4 hours per 100,000 trucks. The average crash rate is the number of crashes in 4 hours per 100,000 trucks inspected. All regressions include dummies for inspection happened outside of the time window of interest, as well as individual truck, year, month, time of day and day-of-week fixed effects. Standard errors are clustered at the truck level. The median duration of an inspection is 22 minutes. Column 1 looks at single-vehicle crashes for trucks that have an average inspection duration of 13 minutes. Column 2 looks at single-vehicle crashes for trucks that have an average inspection duration of 36 minutes.

G Heterogeneity

G.1 Effects by inspection outcomes

Table B.4: The impact of an inspection on crashes by inspection outcomes

Dependent var: number of crashes (per 100,000 trucks)	(1) No violation	(2) Any violation	(3) Driver violation	(4) Vehicle violation
post_insp	2.78*** (0.09)	2.90*** (0.08)	4.04*** (0.15)	2.42*** (0.09)
Average crash rate	6.52	6.23	7.64	5.65
Effect size	42.64%	46.55%	52.88%	42.83%
Observations	323,231,972	347,170,208	138,706,904	261,024,568

Note: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All columns use the same estimation framework following equation 2, which looks at the impact of an inspection on crashes by comparing 14-days before and after. *post_insp* is a post-inspection indicator, it equals to 0 before an inspection happen on truck i , and it equals to 1 on and after the inspection. All regressions include dummies for inspection happened outside of the time window of interest, as well as individual truck, year, month, and day-of-week fixed effects. All regressions dropped inspections that happen within 18 hours after crashes. Standard errors are clustered at the truck level.

Column 1 looks at the impact of an inspection on crashes for trucks that do not receive any violation from the inspection, which accounts for 38% of all inspections. Column 2 looks at trucks that do receive violations but not out-of-service violations (41%). Among trucks that receive violations, column 3 and 4 further break down to driver-related violations and vehicle-related violations. Note that there are trucks that receive both driver and vehicle violations in one inspection.

G.2 Effects by firm characteristics

Firm size: large versus small I look at whether drivers at large and small firms respond differently after inspections. I measure firm sizes according to their inventories in two ways: the number of power units and the number of drivers. Panel A in Table B.5 compares the effect of an inspection on crashes for trucks belong to firms with large or small numbers of power units.³⁸ It shows that the increase in accidents from smaller firms is slightly smaller than large firms: the effect size for

³⁸I define a large firm as one having more than the median number of power units (or drivers) among all inspected firms; the small firms are those in the rest of the sample. So, I can split the full sample into two equally sized sub-samples. A large firm has more than 48 power units (or more than 47 drivers). This number is larger than the number of power units for a median firm in the carrier census file. This is because large firms with more power units are more likely to receive inspections.

large firms is 44.52%, for small firms is 39.26%. The average accident rate for the two types of firms is almost the same. Panel B shows that defining firm sizes using the number of drivers instead of power units gives a similar result.

However, it does not suggest that large and small firms have the same incentives in doing so. There are as many reasons why large firms exhibit compensating behaviors as small firms. For examples, for 1-truck/1-driver carriers, the driver's earnings are tied with how much work he/she chooses to finish. So these drivers have more incentive to drive longer once their perceived inspection probability is lower. On the other hand, a large carrier, such as FedEx, could potentially have a principal-agent problem (Baker and Hubbard (2004)), which suggests that the drivers would not drive in the best manner to preserve the truck's value. Thus, drivers in large and small firms are living in quite different worlds with different profit structures and liabilities, which in turn affect their risk preferences when driving. In this paper, my ability to tease out the different constraints that firms face is limited by data availability. That does not suggest, however, that the dynamics are the same when facing inspections.

Firm business: interstate versus intrastate I look at how carriers respond differently to inspections depending on their major business types – that is, whether they mainly operate on interstate or intrastate routes. This distinction is of particular interest because, depending on the route, drivers face different time constraints. Imagine a driver who encounters an inspection, which could take from 5 minutes to 1 hour, plus the time waiting in the queue entering the weigh station. An intrastate driver is likely on a tight delivery schedule within the day because the driving distance is likely shorter. Thus, the delay caused by the inspection could possibly result in the driver speeding in order to be on time for delivery. On the other hand, a long-haul driver has less incentive in doing so. Panel C in Table B.5 shows that the effect for interstate-only firms is 43.15%, for intrastate-only firms is 35.69%, and for firms that offer both business is 37.84%. Both the effect size and average accident rate for intrastate-only firms are smaller probably because those drivers are familiar with local routes where they operate. But the effect sizes for both interstate and intrastate firms are considerable, which also suggests that the incentive to make up for time lost after inspections is not the major driving force behind the increase in accidents after inspections.

I look at how carrier firms respond differently to inspections depending on whether their main vehicles are trucks or buses. Since trucks transport cargo while buses transport people, the protocols that bind each type of carriers is different. For examples the hours-of-service regulation for bus drivers is stricter than that for truck drivers. Most bus drivers do not work on long-distance route across the US; by contrast, long-distance route are very common for interstate truck drivers. Panel D in Table B.5 compares between these two types of firms. The result shows that the effect of an inspection on accidents for truck-only firms is 41.55%, while for bus-only firms is 42.76%. Both

Table B.5: The impact of an inspection on crashes by firm characteristics

Number of crashes (per 100,000 trucks)	(1)	(2)	(3)
Panel A: Firms size, measured by the number of vehicles			
	Large	Small	1-truck-1-driver
post_insp	3.45*** (0.13)	2.85*** (0.12)	3.01*** (0.30)
Average crash rate	7.75	7.26	6.76
Effect size	44.52%	39.26%	44.53%
Observations	192,570,028	195,500,088	30,159,248
Panel B: Firms size, measured by the number of drivers			
	Large	Small	
post_insp	3.48*** (0.13)	2.82*** (0.12)	
Average crash rate	7.76	7.26	
Effect size	44.85%	38.84%	
Observations	193,200,840	194,869,332	
Panel C: Firms business: interstate versus intrastate			
	Interstate only	Intrastate only	Both inter- & intrastate
post_insp	3.40*** (0.11)	2.12*** (0.29)	2.49*** (0.19)
Average crash rate	7.88	5.94	6.58
Effect size	43.15%	35.69%	37.84%
Observations	290,099,572	26,502,364	70,620,844
Panel D: Firms type: trucks or buses			
	Truck only	Bus only	
post_insp	3.10*** (0.09)	2.54*** (0.65)	
Average crash rate	7.46	5.94	
Effect size	41.55%	42.76%	
Observations	365,358,840	5,755,428	

Note: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All four panels use the same estimation framework following equation 2, which looks at the impact of an inspection on crashes by comparing 14-days before and after. *post_insp* is a post-inspection indicator, it equals to 0 before an inspection happen on truck *i*, and it equals to 1 on and after the inspection. All regressions include dummies for inspection happened outside of the time window of interest, as well as individual truck, year, month, and day-of-week fixed effects. Standard errors are clustered at the truck level. Panel A and B both compare large firms with small firms. I measure company sizes according to their inventories in two ways: the number of vehicles and the number of drivers. I define a large company as one having more than the median number of power units (or drivers) among all companies in the inspection sample, and small company are the rest of the sample. The median is 48 power units or 47 drivers. I also look at firms with 1 truck and 1 driver only. Panel C compares firms that hire drivers to drive interstate, or intrastate, or both routes. Panel D compares firms that own trucks or buses as their main inventory.

type of firms respond similarly to inspections except that bus-only firms have a smaller average accident rate.

Cargo carried I look at how carrier firms respond differently to inspections depending on the type of cargo they transport. I summarize the 30 types of cargo into general freight, chemicals, food and beverage, paper products, building materials, metal sheet, heavy duty commodities, and passengers. Such categorization makes sure that there are enough observations within each category. The results are shown in Table B.6. In general, the effects of an inspection on accidents is very comparable among trucks that transport different types of cargo.

Table B.6: The impact of an inspection on crashes by carrier's cargo types

Type of Cargo	post_insp (per 100,000 trucks)	S.E.	Avg crash rate	Effect size	No. of Obs
General freight	3.34***	(0.11)	7.73	43.60%	275,121,476
Chemicals	3.69***	(0.20)	7.81	47.25%	81,415,936
Food and beverage	3.48***	(0.14)	7.84	44.39%	153,407,212
Paper products	3.44***	(0.16)	7.85	43.82%	126,093,324
Building materials	3.11***	(0.16)	7.54	41.25%	116,545,100
Metal: sheets, coils, rolls	3.19***	(0.19)	7.88	40.48%	88,770,444
Heavy duty commodities	3.00***	(0.12)	7.28	41.21%	190,061,592
Passengers	3.82***	(0.43)	8.53	44.78%	18,884,460

Note: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All columns use the same estimation framework following equation 2, which looks at the impact of an inspection on crashes by comparing 14-days before and after. *post_insp* is a post-inspection indicator, it equals to 0 before an inspection happen on truck *i*, and it equals to 1 on and after the inspection. All regressions include dummies for inspection happened outside of the time window of interest, as well as individual truck, year, month, and day-of-week fixed effects. Standard errors are clustered at the truck level.

This table looks at the impact of inspections on crashes of different carrier companies depending on the type of cargo they transport. I summarized the 30 types of cargo that the carrier/shipper companies transport into general freight, chemicals, food and beverage, paper product, building materials, metal sheet, heavy duty commodities, and passengers. Such categorization makes sure that there are enough observations within each category.

H Predictability index

I develop a predictability index to measure the variation in the day-of-week inspection schedule for each state in the following way. I first demean the number of inspections conducted ($Insp_{ct}$) in a county c at a given day t using interactive fixed effects estimators that control for variations within in county \times day-of-week, county \times year, day-of-week \times year, and county \times day-of-week \times year as shown in the following equation.

$$Insp_{ct} = \alpha_c + \gamma_{dow} + \eta_{year} + \theta_{c \times dow} + \mu_{c \times yr} + \lambda_{dow \times yr} + \omega_{c \times dow \times yr} + \varepsilon_{ct}, \quad (6)$$

where α_c , γ_{dow} , η_{year} are the county, day-of-week, and year fixed effects, $\theta_{c \times dow}$, $\mu_{c \times yr}$, $\lambda_{dow \times yr}$ are the corresponding three-way fixed effects, and $\omega_{c \times dow \times yr}$ is the fully interactive fixed effect. I then define the predictability index $pred_{s, yr}$ as the variance of the residual, ε_{ct} , at the state \times year level.

$$pred_{s, yr} = var(\varepsilon_{ct}) \quad (7)$$

The lower the $pred_index$, the easier the drivers could predict an inspection. The state \times year predictability index can also be interpreted as the average for all county indices in the given state. Figure 12 shows the yearly average of the state \times year predictability index across all states.

The predictability index defined above captures the variation left in the inspection schedule after accounting for fixed patterns in the schedule at the county-by-day-of-week-by-year level. So, the predictability index is not driven by different inspection frequencies, traffic and road conditions across states. For example, in appendix table B.7, I test that the effects from different predictability in schedules on post-inspection accidents are not affected much after controlling for the total number of inspections a state does in a year. In the two-state example in Figure 11, State A and B both have 58 inspections per week in total. Using the predictability index defined above, State A has a predictability index of 0, State B has a predictability index of 32.


Table B.7: Alternative regulatory designs: control for the total number of inspections

Panel A: Predictability in schedule	
post_insp	2.83*** (0.06)
std. $pred_{s,yr}$	-0.04 (0.11)
post_insp \times std. $pred_{s,yr}$	-0.30*** (0.10)
std. $\#insp_{s,yr}$	-0.07 (0.11)
post_insp \times std. $\#insp_{s,yr}$	-0.37*** (0.10)
Panel B: Randomness in selection	
post_insp	2.83*** (0.06)
std. $Dtime_{s,yr}$	-0.10 (0.08)
post_insp \times std. $Dtime_{s,yr}$	0.53*** (0.07)
std. $\#insp_{s,yr}$	-0.15* (0.08)
post_insp \times std. $\#insp_{s,yr}$	-0.37*** (0.07)
Average crash rate	6.36
Observations	669,801,440

Note: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. $post_insp$ is a post-inspection indicator, it equals to 0 before an inspection happen on truck i , and it equals to 1 on and after the inspection. $std.pred_{s,yr}$ is the predictability index of inspection schedule for state s in year yr , standardized to mean 0 and standard deviation 1. $std.Dtime_{s,yr}$ is the average reinspection time interval calculated for state s in year yr , standardized to mean 0 and standard deviation 1. $std.\#insp_{s,yr}$ is the total number of inspections conducted in state s in year yr , standardized to mean 0 and standard deviation 1. All regressions include dummies for inspection happened outside of the time window of interest, as well as individual truck, year, month, and day-of-week fixed effects. Standard errors are clustered at the truck level.

I Appendix Figures & Tables

Figure A.4: The number of CMV crashes vs. VMT over time



Note: The left y-axis is the number of truck-related accidents in thousands, which is aggregated using the FMCSA crash data for all commercial vehicles. It is the same data used in the main analysis. The right y-axis is the annual vehicle miles traveled (VMT) by commercial vehicles, which is obtained from the Bureau of Transportation Statistics.

Figure A.5: An inspection at a weigh station



(a) step1: enter the weigh station



(b) step2: get weighted



(c) step3: the inspector chooses



(d) step4: closer inspection

Figure A.6: The number of trucks and buses over time

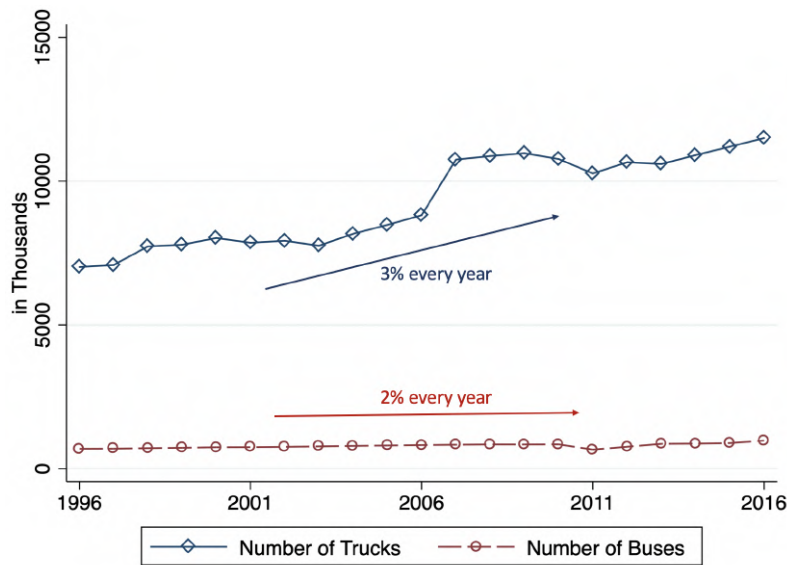



Figure A.7: The number of inspections over time



Note: In Figure A.6, the number of trucks and buses are collected from the Bureau of Transportation Statistics. Data for 2007-17 were calculated using a new methodology developed by FHWA. Data for these years are based on new categories and are not comparable to previous years. So the rate of increase are calculated separately for years before 2007 and after 2007 and then taken the mean. They average at 3% per year over the period.

Figure A.7 is the author's calculation based on all inspection records, including fixed weigh station inspections and roadside inspections. The y-axis on the left shows the total number of inspections in the US per year, while the y-axis on the right shows the number of inspections per truck. The number of inspections per truck declines over time despite the increase in the total number of inspections since the number of vehicles increase much faster.

Figure A.8: The days of inspection per year by commuting zone


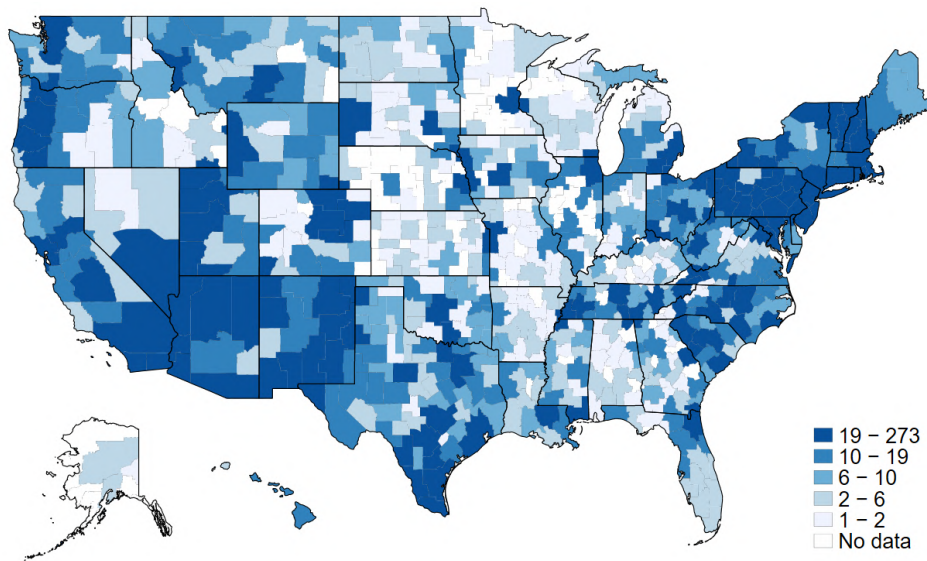



Figure A.9: The number of weigh stations by commuting zone




Note: From author's calculation based on the inspections conducted at the fixed weigh stations. Data are aggregated to the 709 commuting zones in the US.

Figure A.10: Raw data patterns: number of crashes by inspection-crash time lapse




Note: I count the number of crashes that happened in 1 to 12 months following the most recent inspection of the same trucks that had the accidents. The left figure adds up all accidents, while the right figure separately plots accidents by state.

Figure A.11: Comparing effect sizes in different crash categories




Note: This figure complements Figure 5 by further breaking down the single-vehicle crash group into detailed crash categories. First 5 categories in the figure are all single-vehicle crashes: truck run-off road, truck rollover, cargo loss, equipment failure, and collision with fixed objects. Multi-vehicle crashes include all collisions with moving motor vehicles. The vertical axis is the percentage change relative to each of the average crash rate. The figure shows that, for most single-vehicle crash categories, the percentage increase in crashes are all larger than the multi-vehicle crashes. The coefficients plotted in this figure are estimated using equation 1 where I regress the number of crash accidents for the same truck that receives the inspection on a set of inspection indicators from 14 days before an inspection to 14 days after inspection, controlling for individual truck fixed effects, year, month, and day-of-week fixed effects. The effect of an inspection on crash accidents happening in the two-day bin (-2,-1) is normalized to 0.

Figure A.12: Additional evidence: the impact of an inspection on accidents in Texas




Note: This figure uses crash data obtained from the Texas Department of Transportation (TxDOT) from 2010 to 2018. This set of crash files contain all types of crashes, including crashes that result in property damage only, injury, and fatalities. This figure uses the exact same econometric framework in equation 1 for all inspections and crashes in Texas. There are 1,204,497 inspections in this analysis. It shows that there is a 39.7% increase in truck crashes following inspections in Texas, which is almost the same as the average effect size using the data from the whole US.

Figure A.13: Placebo test: t-statistics for post-inspection coefficients



Note: The histogram shows the distribution of the t-statistics of the post-inspection coefficients from the 500 placebo tests. This figure complements Figure 2 for the placebo test results. The kernel density and the normal density estimate are both generated from the t-statistics. Both density plots are close to the Student's t-distribution.

Figure A.14: Distribution of Rest Areas and Weigh Stations



Note: I collect geo-coded location information on all the safety rest areas from Texas Department of Transportation. While the exact locations of weigh station are difficult to pin down, I could locate them at the county-level. So darker colors indicate that the county has a weigh station, and the points show where the rest areas are. This figure shows that there are a abundant number of counties that have weigh stations present but not rest areas. There are on average 1.2 (s.d.=1) rest areas per county, or 2 (s.d.=1) per county where there is at least one rest area.

Table B.8: The impact of an inspection on no-injury truck crashes

	(1) Baseline (all crashes) (per 100,000 trucks)	(2) No-injury & nonfatal crashes (per 100,000 trucks)
post_insp	2.84*** (0.06)	1.61*** (0.05)
Average crash rate	6.37	3.53
Effect size	44.58%	45.61%
Observations	670,402,180	669,950,988

Note: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Both columns use the same estimation framework following equation 2, which looks at the impact of an inspection on crashes by comparing 14-days before to 14 days after inspection. *post_insp* is a post-inspection indicator, it equals to 0 before an inspection happen on truck *i*, and it equals to 1 on and after the inspection. All regressions include dummies for inspection happened outside of the time window of interest, as well as individual truck, year, month, and day-of-week fixed effects. Standard errors are clustered at the truck level. All regressions dropped inspections that happen within 18 hours after crashes. Column 1 is the baseline estimation using all kinds of crashes. Column 2 only looks at no-injury and nonfatal crashes, other crashes are dropped in the estimation. The average crash rate is calculated using the number of crashes per 100,000 trucks inspected.

Table B.9: The impact of an inspection on accidents by rest area distributions, Texas

	(1) Baseline	(2) Rest area distribution	(3) Rest area distribution
post_insp	1.75*** (0.06)	2.11*** (0.06)	2.11*** (0.22)
# rest areas		0.36 (0.28)	
post_insp \times # rest areas		-0.32*** (0.11)	-0.31*** (0.11)
Additional FEs			county-by-year FEs
Average crash rate	3.98	3.98	3.98
Effect Size	44.0%	53.0% (no rest area)	53.0% (no rest area)

Note: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The Texas sub-sample contain 77,265,048 observations. All three columns use the same estimation framework following equation 2, which looks at the impact of an inspection on crashes by comparing 14 days before to 14 days after inspection. *post_insp* is a post-inspection indicator, it equals to 0 before an inspection happen on truck *i*, and it equals to 1 on and after the inspection. # rest areas is the number of rest areas in the county that the inspections took place. All regressions include dummies for inspection happened outside of the time window of interest, as well as individual truck, year, month, and day-of-week fixed effects. Column 3 adds county-by-year fixed effects. Standard errors are clustered at the truck level. The coefficients estimated and average crash rates are multiplied by 100,000.