

DISCUSSION PAPER SERIES

DP17317

Laboratory Safety and Research Productivity

Alberto Galasso, Hong Luo and Brooklynn Zhu

ORGANIZATIONAL ECONOMICS

CEPR

Laboratory Safety and Research Productivity

Alberto Galasso, Hong Luo and Brooklynn Zhu

Discussion Paper DP17317

Published 20 May 2022

Submitted 19 May 2022

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Organizational Economics

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Alberto Galasso, Hong Luo and Brooklynn Zhu

Laboratory Safety and Research Productivity

Abstract

Are laboratory safety practices a tax on scientific productivity? We examine this question by exploiting the substantial increase in safety regulations at the University of California following the shocking accidental death of a research assistant in 2008. Difference-in-differences analyses show that relative to `dry lab' scientists who use theoretical and computational methods, the publication rates of `wet lab' scientists who conduct experiments on chemical and biological substances did not change significantly after the shock. At the same time, we find that the shock induced the wet laboratories that more frequently used dangerous substances to reduce their reliance on flammable materials and unfamiliar hazardous compounds. Our findings suggest that laboratory safety may shape the production of science, but they do not support the claim that safety practices impose a significant tax on research productivity.

JEL Classification: K13, J24, O31

Keywords: economics of science, risk perception, Safety regulation

Alberto Galasso - alberto.galasso@rotman.utoronto.ca
University of Toronto and CEPR

Hong Luo - hluo@hbs.edu
Harvard Business School

Brooklynn Zhu - brooklynn.zhu@bi.no
BI Norwegian Business School

Laboratory Safety and Research Productivity

Alberto Galasso
University of Toronto, CEPR and NBER

Hong Luo
Harvard Business School

Brooklynn Zhu
BI Norwegian Business School*

May 18, 2022

Abstract

Are laboratory safety practices a tax on scientific productivity? We examine this question by exploiting the substantial increase in safety regulations at the University of California following the shocking accidental death of a research assistant in 2008. Difference-in-differences analyses show that relative to ‘dry lab’ scientists who use theoretical and computational methods, the publication rates of ‘wet lab’ scientists who conduct experiments on chemical and biological substances did not change significantly after the shock. At the same time, we find that the shock induced the wet laboratories that more frequently used dangerous substances to reduce their reliance on flammable materials and unfamiliar hazardous compounds. Our findings suggest that laboratory safety may shape the production of science, but they do not support the claim that safety practices impose a significant tax on research productivity.

Keywords: economics of science, risk perception, safety regulation

JEL Codes: K13, J24, O31

“If you want to become a chemist, you have to ruin your health. Who does not ruin his health by his studies, nowadays will not get anywhere in Chemistry” — Justus Liebig

1 Introduction

Workplace accidents are one the leading causes of death and disabilities.¹ An important example of hazardous working environments is academic research laboratories. In fact, during the past decades

*Weida Li, Shenye Liu, Yash Karani, Ehsen Tayyabi, and Bijan Miriabi provided excellent research assistance. Alberto Galasso is grateful for financial support from the Social Sciences and Humanities Research Council of Canada.

¹A recent joint study by the World Health Organization and International Labour Organization estimates that almost 2 million people die from work-related causes each year (WHO-ILO, 2021).

growing concerns have been voiced in both academic and policy debates about the safety culture and protocols in place across universities in the United States, Europe, and China (Silver, 2021). Several proposals to introduce stricter safety regulations have been put forward, including the use of safety metrics in journal publication, faculty promotion, and grant award decisions (National Research Council, 2014).

Implementing stricter safety measures is challenging in part because of the widely-held belief that they may reduce research productivity. Time spent on safety training, risk assessment, and documentation is time not spent on research. Research funds spent on personal protection equipment (PPE) may be diverted from other uses. Supervision by experienced researchers and co-workers also constrains resources.² Indeed, survey evidence indicates that safety precautions are seen by many academic scientists as a tax on scientific productivity and as an infringement on academic freedom (e.g., Van Noorden, 2013; National Research Council, 2014). In contrast, proponents of stricter regulations argue that a safe working environment could increase research productivity. In hazardous environments, adequate safety precautions allow scientists to focus on their research rather than worry about emergencies. Accidents not only threaten researchers' health and safety, but ultimately may result in research projects being shelved or substantially delayed. Serious injuries and death may also affect a laboratory's reputation, making it difficult to attract new researchers.

These contrasting views highlight the need for empirical research. In this paper we provide a first set of evidence on the relationship between lab safety and research productivity. We focus on academic chemistry laboratories. Hazardous chemicals impose health and safety risks on researchers, and they are not as stringently regulated as radioactive or biological materials (Van Noorden, 2011). Our analysis exploits a substantial and quasi-exogenous surge in academic lab safety regulations and examines its impact on researchers' productivity.

Specifically, on December 29, 2008, Sheharbano Sangji, a research assistant in an organic chemistry laboratory at the University of California Los Angeles (UCLA), spilled a highly flammable compound that ignited and severely burned over half of her body. Sangji died two weeks later. Investigations

²Surveys by the American Chemical Society in 2010 suggest that 70.5% of faculty and 52.1% of graduate students often or occasionally work alone in laboratories, which is forbidden in industry (Van Noorden, 2013).

determined that inadequate safety training and not wearing PPE were among the causes of this tragic incident. This led to the first-ever criminal case against a university professor in a lab safety accident. The unprecedented event received extensive press coverage and was described as “*a pivotal development*” for academic lab safety regulations that “*had more impact on lab safety than anything else that’s happened in the last 20 years*” (Basken, 2012; Trager, 2014). The University of California (UC) system responded to these investigations by implementing major changes to their safety programs. These included more frequent lab inspections, more stringent protocols for dangerous chemicals, and more frequent safety training classes for laboratory scientists (Gibson, et. al. 2014).

From the perspective of principal investigators (PIs), who are the heads of the labs, these events have led to two major changes. The first is a significant increase in the perceived risk (and liability) associated with laboratory accidents, and the second is a significantly more stringent regulatory environment. To guide our empirical analysis of the impact of these changes on research productivity, we develop a simple model describing the relationship between risk/liability perception, safety levels, and research output. The model builds on the multi-tasking literature (Holmstrom and Milgrom, 1991) and assumes that researchers direct their efforts towards either conducting research or improving the safety of their labs.

The model first clarifies that even in the absence of stricter safety regulations mandated by the university, an increase in risk perception in itself incentivizes scientists to increase efforts related to lab safety. Ultimately, whether an increase in safety leads to an increase or a decrease in research output depends on the relationship between research and safety efforts. Intuitively, when safety and research efforts are substitutes (that is, safety practices are a tax on research productivity), we will observe a decrease in research output. Conversely, if these two efforts are complements, we will observe an increase in research output.

Our empirical analysis relies on a sample of 609 chemistry labs affiliated with the University of California system between 2004 and 2017. For each lab, we retrieve its complete record of publications from the Web of Science database. Importantly, we hire a team of chemistry PhD students who examine the information available on a lab’s webpage and the CV of the PI and classify each lab into one of two groups. The first (treatment group) includes ‘wet labs,’ in which scientists conduct experiments on chemical and biological substances, and the second (control group) includes ‘dry labs’ that specialize in

computational and theoretical research. Our identification assumption is that while the UCLA accident and subsequent events are likely to have a prominent impact on wet labs, they will have limited impact on dry labs.

Our main finding is that despite the implementation of more stringent safety regulations, we do not see a significant decline in research productivity. Specifically, difference-in-differences estimations show that, relative to dry labs, wet labs at UC experienced a reduction of about 3 percent in their yearly publications after 2008, but we *cannot* statistically distinguish this effect from zero. Event study analysis shows no pre-trend differences between these two groups. The small and insignificant estimate is robust across a variety of robustness checks, including different econometric models, alternative metrics to measure research output, and a different method to distinguish between wet and dry labs.

While we do not find a significant change in the level of research output, utilizing complementary datasets—which allow us to observe all the compounds involved in the research presented by a journal article and, importantly, the compounds’ safety hazard information—we find that wet labs that tended to use dangerous chemicals more frequently before the accident have reduced their usage of dangerous chemicals after the accident. This is especially the case for dangerous chemicals that are flammable and those that are relatively unfamiliar to the researchers. These results provide evidence of changes in scientists’ behavior that increases lab safety levels. They also illustrate ways in which lab safety may shape the production function of science.

Taken together, our results do not support the claim that lab safety imposes a significant tax on research productivity. Interpreted through the lens of our model, the null finding is consistent with the idea that safety and research efforts can be linked by multiple channels that simultaneously generate complementarity and substitutability effects. While increasing investment in safety may crowd out time and resources for research, it may also help to mitigate potential declines in researchers’ productivity due to concerns about safety, especially when the perceived risk associated with working in a chemistry lab has substantially increased. The data suggest that these contrasting effects tend to compensate each other, leading to limited changes in research output.

The paper is organized as follows. Section 2 discusses related literature. Section 3 describes the UCLA accident and what happens to chemistry lab safety regulations in the wake of this accident.

Section 4 presents the theoretical model. Section 5 describes the data and econometric methods. Section 6 describes the impact on the level of publications, and Section 7 examines the impact on the use of dangerous chemicals. Concluding remarks summarize and discuss our main findings.

2 Related Literature

Our paper connects and contributes to the following three literatures. The first comprises studies on the link between worker safety and productivity.³ Theoretically, this literature generally posits that there is a tradeoff between productivity and worker safety (Oi, 1974; Thaler and Rosen, 1976; Gowrisankaran et. al. 2015). Empirical studies also largely find a negative relationship between safety and productivity, with much of the evidence coming from the mining and oil extraction industries (Gray, 1987; Gowrisankaran et. al., 2015; Boomhower, 2019). An exception is Levine et. al. (2012), who find that randomized Occupational Health and Safety Health Administration inspections lower injury rates with no loss in sales or employment. Our contributions to this literature are two-fold. First, we provide novel evidence on academic labs and scientific productivity, which is an important sector that fuels R&D and, ultimately, economic growth. Second, by applying a multi-tasking model, we allow for the theoretical possibility that safety and productivity may also assume a complementary relationship.

Second, this paper joins a growing number of empirical studies on how risk perception and liability costs may shape innovation and R&D strategies.⁴ Despite the dominant policy view that greater liability risk chills innovation (Porter, 1990), existing large-sample evidence mostly does not support this view. Viscusi and Moore (1993) and Galasso and Luo (2017) both find that on average, higher liability risk induces higher R&D spending and more patenting. Relatedly, Galasso and Luo (2021) show that shocks that increase consumers' perceived risk of product use can induce firms to develop risk-mitigating technologies and increase overall R&D. The exception is Galasso and Luo (2022), which shows a signif-

³More broadly, our paper is also related to studies on the economic costs of regulation. For example, Coffey et. al. (2020) find that regulatory restrictions have had a net effect of dampening economic growth by approximately 0.8 percent per year since 1980. Greenstone et. al. (2012) find that stricter air quality regulations are associated with a roughly 2.6 percent decline in productivity.

⁴Theoretical models in law and economics suggest that the impact of liability risk on innovation is ambiguous (e.g., Daughety and Reinganum, 2013; Dawid and Muehlheusser, 2022). On the one hand, higher liability may reduce innovation incentives by raising the costs of or chilling the demand for new technologies associated with greater risk. On the other hand, it may also increase the profitability of and the demand for risk-mitigating technologies and safer product designs that reduce the likelihood of injuries.

icant chilling effect through a specific mechanism: a sudden increase in product liability risk faced by suppliers may disrupt supply chains and negatively impact downstream innovation investments. Our paper contributes to this literature by examining the relationship between research productivity and workplace safety, rather than studying the safety of the final products as in previous studies.

Finally, our paper contributes to the literature on the drivers of productivity in academic research. Stephan (2010) provides a comprehensive survey of the topic. Several studies in this literature exploit premature death of scientists (from illness or non-work-related accidents) to estimate human capital spillover effects on peer researchers (Azoulay et. al. 2010; Oettl, 2012; Jaravel et. al. 2018) and the evolution of research fields (Azoulay et. al. 2019). Baruffaldi and Gaessler (2021) examine the effect of adverse events (such as fires and floods that do not involve human casualties) and show that losing specialized physical capital can lead to substantial and persistent declines in scientists' research productivity. We contribute to this literature by providing new evidence on lab safety regulations, which is an important, contentious, but understudied issue.

3 The UCLA accident and subsequent events

On December 29, 2008, research assistant Sheharbano (Sheri) Sangji was conducting an experiment in Professor Patrick Harran's organic chemistry lab at UCLA. When she attempted to transfer a tert-butyl-lithium solution, a highly flammable compound, from a bottle to a flask, the syringe plunger came out of the barrel and the chemical burst into flames. Sangji was not wearing a protective lab coat at the time, and her synthetic sweater caught fire. Sangji was rushed to the hospital with over half of her body severely burned, and she died from her injuries a few days later on January 16, 2009 (Kemsley, 2009).

In December 2009, the investigation by the State of California's Division of Occupational Safety and Health (OSHA) concluded that Sangji had not received adequate training for work with hazardous chemicals as required by the State of California. The report noted that UCLA's environmental health and safety (EH&S) department "*was well aware that research staff within virtually all laboratories at the University routinely did not wear lab coats and other personal protective equipment while working in the labs.... The practice was so well known by EH&S that it was simply regarded 'as part of the culture.'*"

(Baudendistel, 2009).

In December 2011, almost three years after the accident, the Los Angeles District Attorney filed criminal charges against the Regents of the University of California and Patrick Harran for willful violation of safety regulations. The case against the university was settled in July 2012. The terms of the settlement required UC to accept responsibility for the event, establish a scholarship in honor of Sangji, pay the OSHA litigation costs, and implement a number of specific laboratory safety practices. These practices encompassed lab safety manuals, hygiene plans, training of staff and PIs, and following standard operating procedures for hundreds of chemical substances (Merlic, 2013).

Harran, if convicted, could have faced up to four and a half years in prison. In June 2014, Harran reached a settlement under which he was required to: (i) teach organic chemistry to inner city high school graduates for five years; (ii) complete 800 hours of non-teaching community service, (iii) speak to UCLA students about the importance of lab safety and (iv) pay a \$10,000 fine to the regional burn center where Sangji was treated (Trager, 2014). On September 6, 2018, having determined that Harran had now met the terms of the agreement, a Los Angeles County Superior Court judge dismissed the criminal charges against Harran (Maxmen, 2018).

The UCLA accident and what happened subsequently have received wide media coverage and close attention by the academic scientific community. News outlets including the Los Angeles Times, Chemical & Engineering News, and Chemistry World reported the accident shortly after Sheri Sangji's death and followed up on subsequent events.⁵ This accident and laboratory safety became a major topic of discussion at the American Chemical Society meeting in March 2009 (Benderly, 2009). Numerous editorials and blog posts in Science, Nature, Scientific American, and other publications (including blogs maintained by chemists and Reddit) discussed the details of this accident and called for actions to improve laboratory safety and prevent future accidents.

The events presented above led to two major changes to the institutional environment in which

⁵The initial reports at these outlets are: "Researcher Dies After Lab Fire UCLA research assistant burned in incident with tert-butyl lithium," by Jyllian Kemsley, January 22, 2009, C&EN (Source: <https://cen.acs.org/articles/87/web/2009/01/Researcher-Dies-Lab-Fire.html>); "UCLA lab assistant dies," by Rebecca Trager, 23 January 2009, Chemistry World (Source: www.chemistryworld.com/news/ucla-lab-assistant-dies/3004085.article); "Deadly UCLA lab fire leaves haunting questions," by Kim Christensen, March 1, 2009, Los Angeles Times (Source: <https://www.latimes.com/local/lanow/la-xpm-2009-mar-01-me-uclaburn1-story.html>).

scientists operated. The first was a significant increase in the perceived risk of working in academic chemistry laboratories involving dangerous substances. This happened for several reasons. First, people became more aware of the hazards of working with dangerous chemicals and their potentially severe consequences. In 2009, after reviewing the reports on the UCLA and other accidents, Dr. Neal Langerman, former chair of the Division of Chemical Health and Safety of the American Chemical Society (ACS), said that: *“I have come to the disheartening conclusion that most academic laboratories are unsafe venues for work or study”* (Langerman, 2009). A 2013 survey conducted by Nature and UCLA shows that 30 percent of the 2,400 responding scientists reported having witnessed a lab injury that was severe enough to warrant attention from a medical professional (Van Noorden, 2013).

From the perspective of the PIs, the criminal case against Patrick Harran, the first ever of its kind, also clarified their personal responsibilities for laboratory safety and heightened their expected liability costs. Media described the unprecedented case as *“a pivotal development”* for academic lab safety regulations (Basken, 2012). In a 2014 interview, Russ Phifer, the executive director of the National Registry of Certified Chemists stated that the case against Harran *“had more impact on lab safety than anything else that’s happened in the last 20 years”* (Trager, 2014).

Apart from legal liability, laboratory safety issues may also have imposed greater costs on PIs’ reputations, statuses, and careers. Various proposals were made to take accident reports, laboratory investigations, and safety policy compliance into account for promotion, tenure and the allocation of grants and departmental resources. In December 2015, the American Association for the Advancement of Science (AAAS) decided to withhold a recognition of Patrick Harran. In an interview with Nature, Langerman stated that *“this action is huge, and impacts every scientist who aspires to be named for national recognition or international recognition. . . . If I were a young chemist, and I set a career goal to win a Priestley Medal [the highest honour conferred by the American Chemical Society] this says that if my lab has a serious incident, I may never achieve my goal”* (Hayden, 2015).

The second major change to the institutional environment was the introduction of stricter laboratory safety rules. UCLA, specifically, responded to the OSHA report by substantially increasing the laboratory safety protocols implemented by its EH&S department. These included more frequent lab inspections and more stringent rules for the handling of dangerous chemicals. Lab safety training classes

were made mandatory for all laboratory personnel and made available both online and in person on a weekly basis, rather than quarterly as before (Kemsley, 2009). Gibson et al. (2014) report a marked increase in the number of lab safety class participants at UCLA in 2009 (about 13,000) compared to 2008 (about 6,000), and the number of class participants increased to almost 22,000 in 2012. Similarly, the number of safety inspections at UCLA increased from about 1,100 in 2008 to about 2,000 in 2009, and to about 4,500 in 2012. Beyond UCLA, the University of California also implemented UC-system wide changes by establishing a Center for Laboratory Safety (CLS) in March 2011, with the goal of supporting research in laboratory safety as well as of developing and diffusing best practices (National Research Council, 2014).

The increased focus on laboratory safety went beyond the University of California and spread to chemistry departments across universities in the United States. Immediately after the UCLA accident, Russell Phifer, chair of ACS’s chemical health and safety division, said in an interview with Science: “*I know for a fact that many universities immediately reviewed their protocols for dealing with pyrophorics and many of them looked at their documentation of safety training*” (Benderly, 2009). In 2012, the American Chemical Society (ACS) issued the report “Creating Safety Cultures in Academic Institutions” that described best practices and provided recommendations to university departments. This was followed by an ACS presidential commission recommending the adoption of the best safety practices as a key requirement for the advancement of graduate education in chemical sciences (ACS, 2012a; 2012b).⁶

4 Theoretical considerations

In this section, we develop a simple theoretical model that helps illustrate the potential impact of the UCLA accident and the events that followed on research output. Section 3 suggests that these events led to two major changes from a PI’s perspective: i) a significant increase in the perceived risk (and liability) associated with laboratory safety; and ii) stricter safety regulations implemented by the university. Research shows that changes in risk perceptions may lead to behavior changes or innovations that help mitigate risk, especially when markets fail to provide full insurance against uncertainty (Arrow,

⁶Recommendations in the ACS report include: (i) the development of written protocols and training, (ii) appointing a university official who has the authority to oversee research laboratories safety, and (iii) documenting and communicating all laboratory near-misses and accidents (Merlic, 2013).

1970; Galasso and Luo, 2021). Thus, even in the absence of stricter safety regulations mandated by the university, an increase in risk perception itself may change how labs operate—e.g., greater use of PPEs, more rigorous safety training, and extensive documentation of lessons learned from prior accidents and near misses—that could, in turn, affect research outcomes.

To illustrate this point, we use a multi-tasking model in the spirit of Holmstrom and Milgrom (1991), in which a PI decides to allocate time and resources to two types of activities: $r \in (0, 1)$ is the research effort directed at conducting experiments and publishing new results, and $s \in (0, 1)$ is the effort aimed at reducing the risk of accidents and their consequences.

With efforts (r, s) , the PI enjoys a benefit of

$$B(r, s) = r - \delta(1 - s)$$

where r is the monetary and reputational rewards from publications, $1 - s$ is the risk level of the lab, and δ captures the perceived costs associated with lab accidents. For principal investigators, $\delta > 0$ includes both the legal liability and reputational loss due to accidents, as well as their internalization of harms to researchers working in the lab. We also assume that $\delta < 1$ to reflect the notion that from the PI’s perspective, the marginal benefit of safety is typically less than that of research output. This is an issue prominently highlighted in the 2014 National Research Council report discussing how academic reputation and decisions about advancement, salary, and space tend to focus on research productivity.

We model the cost of these efforts as

$$C(s, r) = \frac{s^2}{2} + \frac{r^2}{2} + \rho sr.$$

This quadratic-cost formulation is standard in the multi-tasking literature (inter alia see Fryer and Holden, 2013; Benabou and Tirole, 2016; and De Philippis, 2021). The parameter $\rho \in (-1, 1)$ reflects the notion that these two efforts may affect each other. Intuitively, because resources are fixed, allocating more time and budget to implementing safety protocols means less time and budget for research; that is, these two efforts may be substitutes. This substitutability relationship is consistent with the recurring theme that an important barrier to improving laboratory safety is the perceived conflict between safety and research productivity. Harry J. Elston, editor of the *Journal of Chemical Health and Safety*, writes

that the Sangji case is “*a harbinger of things to come*” unless scientists devoted to accident prevention are willing to “*stand in the gap between worker’s safety and [scientific] productivity.*”⁷ This view also emerged from the UCLA-Nature survey, with one-fifth of the respondents indicating that lab safety rules had negatively impacted their research productivity (Van Noorden, 2013). This perception was also highlighted by a 2014 National Research Council report on safe science, which states that “*one of the most recalcitrant problems in many chemistry laboratories is the attitude, unfortunately often reinforced by principal investigators, that safety practices are time-wasting inhibitions to research productivity*” (National Research Council, 2014).

While there seems to be a strong prior that research and safety efforts are substitutes, the two efforts may also have a complementary relationship if a safer work environment facilitates research. Safer workspaces, by reducing the risk of accidents, enable researchers to “*focus on their tasks rather than worrying about emergencies,*” because “*in addition to putting people at the risk of harm, these incidents ultimately decrease productivity, as they hinder the researchers’ ability to work*” (Hersh, 2017). Labs with bad safety records may also find it harder to attract post-docs and PhD students. Moreover, laboratory safety may also affect a lab’s ability to obtain research funding, given the increasing calls to include safety records as part of promotion decisions and allocation of departmental resources, grants, and prizes (National Research Council, 2014).

These mechanisms are *not* mutually exclusive, and ρ reflects the net effect. If substitutability dominates complementarity, we will have $\rho > 0$; if complementarity dominates substitutability, $\rho < 0$; and if these two effects are roughly equal, we will have $\rho = 0$.

We map the two institutional changes that followed the UCLA accident to the theoretical model as follows. First, we capture the increase in risk perception with an increase in the parameter δ . Specifically, we assume that $\delta = \underline{\delta}$ before the accident, and $\delta = \bar{\delta}$ after the shock, with $\Delta\delta = \bar{\delta} - \underline{\delta} > 0$. Second, we model the implementation of stricter safety protocols as an increase in minimum safety requirements. For simplicity, we normalize the minimum required safety level to zero before the accident, and indicate the requirement after the accident as $\underline{s} > 0$.

⁷Journal of Chemical Health and Safety, in a lead editorial entitled "Recipe for disaster," posted to the Internet on 31 March.

We solve this simple model in the Appendix and examine how the joint change in risk perception and mandatory safety levels affects the research output of the lab. Specifically, we compute the optimal research output produced by the lab before the shock, which is indicated by $r(\underline{\delta}, 0)$ and is a function of the pre-shock risk perception and minimum safety requirement, and the optimal research output after the shock, $r(\bar{\delta}, \underline{s})$. The effect of the shock on the research level is thus $\Delta r = r(\bar{\delta}, \underline{s}) - r(\underline{\delta}, 0)$.

Despite its simplicity, the model delivers several implications that help guide our empirical analysis. First, whether the shock increases or decreases research output (that is, the sign of Δr) depends on one specific parameter of the model: the net relationship between the two types of efforts, ρ . Intuitively, the increase in risk perception leads scientists to increase their investment in lab safety (that is, the optimal level of s increases as δ increases from $\underline{\delta}$ to $\bar{\delta}$). This translates into a reduction in research output when the two efforts are substitutes, because an increase in safety investment increases the marginal cost of research. However, the accident leads to an increase in research output when the two efforts are complements; that is, when an increase in safety investments lowers the marginal cost of research.

The fact that a change in the economic environment translates into an increase or decrease in effort depending on the substitutability/complementarity parameter is a common result in multi-tasking models. The slight complication in our setting is that we consider two simultaneous changes $(\Delta\delta, \underline{s})$, rather than focusing on a single change as in most of the previous studies. We show in the Appendix that, despite this issue, the sign of Δr is uniquely determined by the sign of ρ .

Second, the model shows that the magnitude of the effect depends on some combination of ρ , $\Delta\delta$ and \underline{s} . This implies that, in principle, empirical estimates of a small magnitude could be driven by a low level of interaction between the two types of efforts (ρ close to zero) or by small changes in perception and safety regulations ($\Delta\delta$ and \underline{s} close to zero), or both. The evidence presented in section 3 suggests that in our empirical setting, the change in risk perception and the university's response were substantial. Thus, an empirical estimate of small magnitude is likely to reflect a relatively small ρ . As discussed above, research efforts and safety investments can be linked by multiple channels involving complementarity and substitutability, and when these effects tend to compensate each other, ρ will be small.

Finally, the model illustrates how safety levels mandated by the university may or may not be binding. The PIs and lab scientists may want to invest in safety beyond the mandated level when safety

effort complements research, or if their perceived risk is sufficiently high when safety and research efforts are substitutes. Otherwise, the PIs would find the mandated safety level too stringent. This may be the case if δ is relatively low; that is, if the PI does not fully internalize the liability risk faced by the university (which may capture more fully the potential harms to lab researchers).

5 Data

Our analysis relies on a sample of chemistry labs that were affiliated with the University of California system and active around the time of the UCLA accident. The data source used to identify these labs and their research output is the Web of Science (WoS) Journal Citation Reports database.

To generate this sample, we first identify journals that are natural publishing outlets for chemistry researchers. These include: i) journals in the top decile of the impact factor in each of the nine chemistry subfields, as provided by WoS; and ii) the ten multidisciplinary scientific journals with the highest impact factors such as *Science*, *Nature*, and the *Proceedings of the National Academy of Sciences*. We download all the articles published in these journals between 1998 and 2017 from WoS. This step gives us 698,094 articles published in 105 journals.

We then use the author and affiliation information provided by WoS to identify individual labs. In chemistry, typically, the author listed last is the PI of the lab that hosted most of the research, and the corresponding author comes from this lab (Venkatraman, 2010). We rely on this convention to identify the primary PI of each article and the institution that this PI is affiliated with.⁸ We refer to each unique PI-institution combination as a lab.

Many of these labs are not very active; some of them publish only one article, and others are active for only a year or two. Rather than meaningful research units, these low-activity labs are likely to capture errors in the PI names and affiliations or organizations with one-off publication projects. To

⁸Specifically, we used the following procedure. First, based on the author information provided by WoS, we identified the last name and the first initial of each author of a given article, as well as the institution (university, firm, or government agency) with which an author is affiliated. Second, we examined whether the last name and the first initial of the re-print (corresponding) author—which is provided by WoS—matched one of the last three authors listed in the article. If there was a match, we classified the re-print author as the PI. If there was not a match, we examined whether the affiliation of any of the last three authors coincides with the affiliation of the reprint author. If this is the case, we classified the author (among the last three authors) with the matched affiliation as the PI. Otherwise, we classified the re-print author as the PI.

address this issue, we drop labs with an active life span of fewer than three years. Among the remaining labs, we further drop those with a yearly publication rate below the median (0.6 articles per year in the 105 journals described above). This step gives us 6,704 relatively active chemistry labs based in United States.

Among labs identified from the above step, we create the UC Sample as the 609 labs affiliated with the University of California system, and this is the main sample for our empirical investigation. We focus on UC labs for two reasons. First, while the accident had a nation-wide impact, UC labs are most directly affected. Second, as is explained below, with the smaller sample, we can manually collect information that is critical for our identification strategy. In Section 6.1, we examine the external validity of our main finding among non-UC universities.

Leveraging information from sources such as lab websites, department websites, and news releases, we manually confirm that the labs in the UC Sample were indeed run by UC-affiliated PIs. We also collect information on the year in which the PIs joined and left their respective UC institutions, which we use to control for the PI's tenure and to construct a balanced panel for robustness analyses. Finally, we retrieve from WoS all the journal articles (not just in the 105 journals described above) published by these labs between 1998 and 2017 and the citations received by these articles until 2020. The 609 UC labs have published a total of 65,458 journal articles.

UC Berkeley accounts for about 25% of the labs in the UC Sample, UCLA and UC San Diego for about 15% each, and UC Davis for 10%. The remaining labs are affiliated with UC Irvine, UC Riverside, UC San Francisco, UC Santa Barbara or UC Santa Cruz. The PIs have an average tenure of 19 years at UC. On average, the labs publish 7.365 articles each year during our sample period, which, by 2020, have received about 688 citations.

To identify the effect of the shock on research productivity, we distinguish between dry and wet labs. Wet labs refer to labs in which scientists conduct experiments on chemicals and biological substances. Scientists in dry labs, instead, specialize in computational and theoretical analysis and do not handle chemical or biological substances. The assumption underlying our identification assumption is that the shock has a bigger impact on the day-to-day operations of wet-lab scientists (treatment group) than on dry-lab scientists (control group).

An example of a dry lab is the one run by Professor Anastassia Alexandrova, a member of the UCLA chemistry department. Professor Alexandrova’s research focuses on computational and theoretical design and multi-scale description of new materials. Her work relies on quantum and statistical methods, including artificial intelligence and machine learning algorithms. An example of a wet lab is the one run by Professor Ohyun Kwon, also at UCLA. Her research focuses on the transformation, catalyst, and synthesis of natural compounds. Her research group uses various research equipment including fume hoods, solvent stills, and glove boxes to manipulate hazardous materials.

We hire a team of chemistry PhD students to examine the lab webpage, CV, and publication records of each PI in our UC Sample. A lab is classified as a wet lab if the available information indicates that it is equipped to handle biological specimens, chemicals, drugs, and other materials used in experiments. Otherwise, the lab is coded as a dry lab. Of the 609 labs in the UC Sample, 589 (87%) are wet labs, and 81 (13%) are classified as dry labs.

5.1 Econometric method

Our empirical strategy relies on difference-in-differences estimations in which the treatment group includes wet labs and the control group includes dry labs. The pre-treatment period is 2004-2008, and the treatment period is 2009-2017. The unit of observation is a lab-year. Specifically, we estimate

$$Y_{l,t} = \alpha + \beta WetLab_l \times AfterAccident_t + \theta X_{l,t} + \delta_t + f_l + \varepsilon_{l,t}, \quad (1)$$

where the dependent variable, $Y_{l,t}$, captures the publication level of lab l in year t . The treatment variable, $WetLab_l$, is equal to one for wet labs. The dummy, $AfterAccident_t$, is equal to one for the years after 2008. The term $X_{l,t}$ captures time-varying controls at the lab level. The terms δ_t and f_l are year and lab fixed effects. The coefficient β is a difference-in-differences estimator for the effect of the changes in risk perception and safety regulations following the UCLA accident on the research output of wet labs relative to dry labs.

The baseline results are estimated by OLS regressions, with the standard errors clustered at the lab level. Section 6 confirms the robustness of our findings using alternative specifications that account for both the count nature and other features of our data. Table 1 provides summary statistics for the key

empirical variables used in our empirical analysis.

6 Impact on publication levels

Table 2 presents the first set of estimates of the effect of the shock on the publication levels of wet labs relative to dry labs at UC universities. The analysis uses the difference-in-differences model described in equation (1). To account for experience and career concerns of the PIs, we control for the lab’s past productivity (measured as the total number of publications in the past three years) and for the PI’s tenure at UC (in log).

Column 1 shows that relative to UC dry labs, UC wet labs reduce their publications by about 0.28 papers per year, on average, after 2008 compared to before. The estimate is *not* statistically significant at the ten percent level. Assuming the same difference between dry and wet labs before and after 2008, the hypothetical average for wet labs would have been 8.88 publications per year after 2008. This implies that the average decline in publication is about 3 percent for wet labs after 2008.⁹ This is small relative to the estimates of other drivers of scientists’ productivity in the literature. For example, Oettl (2012) estimates a 20 percent decrease in performance associated with an unexpected loss of a highly productive and helpful co-author, whereas Baruffaldi and Gaessler (2021) show that unexpected loss of lab equipment leads to a publication decline of about 15 percent.

Column 2 focuses on the most impactful publications; specifically, instead of all publications, we count only publications with citation counts in the top decile of our sample, and the estimated effect of the shock remains small and statistically insignificant.

In column 3, we again use the total number of publications as the dependent variable but restrict the sample to only labs at UCLA, and column 4 further drops the lab of Patrick Harran, where the accident took place. The estimated effects are even smaller than those estimated for the full UC sample. Finally, column 6 re-estimates the effect of the shock in the UCLA sample using the number of articles published per lab member as the dependent variable. It is difficult to recover historical data on lab members. As a proxy, we use the number of unique researchers who publish with a PI in a given year

⁹The average number of papers for dry-labs after 2008 is 8.88, and the pre-2008 difference between wet and dry labs is -0.31 papers per year.

and who are also affiliated with UCLA.¹⁰ We estimate a small, positive, and statistically insignificant coefficient, which indicates that the shock had a limited impact not only on a lab’s total research output but also on the productivity per lab member.¹¹

The analyses above suggest that relative to dry labs, wet labs do not appear to experience a significant decline in research output after the UCLA accident, despite the significantly more stringent safety regulations. Our model suggests that this is consistent with the idea that the complementarity and the substitutability relationships of safety and research efforts roughly cancel each other (that is, ρ is close to zero). On the one hand, increasing investment in safety may crowd out time and resources for research output. On the other hand, it also helps mitigate potential decline in research productivity due to concerns about safety issues, especially when the risk associated with working in a chemistry lab has substantially increased.

6.1 Robustness and extensions

Below, we show that our main finding is robust to alternative specifications and extensions.

Alternative econometric models. We first confirm the results of Table 2 using alternative econometric models. Specifically, to address the skewed and count nature of our dependent variable, column 1 of Table A1 uses a Poisson quasi maximum-likelihood estimation. Using a similar Poisson model, column 2 estimates the change in publications using a citation-weighted measure as the dependent variable. In both cases, the estimated effects are negative, small, and statistically insignificant.

Another concern is that our estimates may be affected by the entry and exit of scientists who are not active at UC for the entire sample period. To address this issue, we replicate our analysis in column 3 of Table A1, using the sub-sample of labs run by scientists who remain at UC during the entire sample period. This balanced sample includes 326 labs (40 dry labs and 286 wet labs). The estimated effect is essentially the same as for our baseline result.

¹⁰Specifically, we first identify all the unique co-authors of the articles published by a focal PI in a given year. We then keep those who are also affiliated with UCLA. Finally, we drop names that match PIs in our UCLA sample and names with a relatively long tenure at UCLA—specifically, above the 95th percentile of the tenure distribution among all the coauthors of UCLA PIs—to exclude potential faculty members in other departments at UCLA.

¹¹Unreported results also confirm that relative to dry labs at UCLA, there is no significant change in the lab size for wet labs after the shock.

In column 4 of Table A1 we show that results are statistically similar in a weighted OLS model, in which observations are weighted by the pre-2008 publication level of the lab. This suggests that our baseline finding is not driven by a few large labs.¹²

We also show that our baseline estimate is robust to controlling for time-varying university characteristics such as the market value of the university endowment, the number of chemistry PhD students enrolled at the university, and federal science and engineering grants received by the institution.¹³ Results are similar if we replace these time-varying controls with institution-year fixed effects.

Pre-treatment trend and time-specific treatment effects. Our empirical model assumes that the publications of wet labs follow a similar trend as that for dry labs. To provide support for this assumption, we extend our baseline model to estimate the time-specific differences between treatment and control labs. Figure 1 provides a graphical illustration of the estimated coefficients and their 90-percent confidence intervals. The dependent variable is the number of yearly publications. There are no statistically significant differences in the publication level between the two groups of labs before 2008, which supports the common-trends assumption.

Interestingly, despite being statistically insignificant, the post-2008 estimates show a lower level of publications in 2010 relative to 2009 and 2011, and a slow decline in publications during the period 2012-15, which is followed by a recovery in later years. On the one hand, this is consistent with the idea that the 2008 accident may have led to an initial chilling effect on wet lab research, which is manifested in fewer publications in 2010, and that the introduction of stricter safety protocols following the OSHA settlement in 2012 may also have temporarily slowed down some research projects. On the other hand, the magnitude of the differences between the two groups remain relatively small after 2008, and the coefficients are not statistically significant. This suggests that the chilling effect, even if present, is very limited and does not have a long-lasting impact on a lab’s research output.

¹²This empirical exercise is closely related to the econometric analysis in Greenstone et al. (2010). Intuitively, one can interpret our baseline unweighted regression as estimating the impact of the shock on the average lab, whereas the output-weighted regression as an estimate of the impact on the average output level. When there is a substantial difference in the estimates using these two approaches, one can infer that the effect is not homogenous across labs with different output levels.

¹³Endowment data are collected from the National Association of College and University Business Officers, the grants and PhD student data are from surveys by the National Science Foundation. The information is missing for a few institutions-years which explains the lower number of observations in this regression.

Journal-based lab classification. We also use an alternative method to distinguish between dry and wet labs. Instead of manually classifying each lab based on information available from its webpage, we ask the team of chemistry PhD students to identify the subset of journals (among the 105 described in Section 5) that specialize in theoretical and computational work. We classify a lab as a dry lab if the fraction of its pre-2008 publications in these theoretical journals is in the top decile of the sample. Table A2 shows no significant difference in the publication level between the treatment and control groups using this alternative definition of dry versus wet labs in the UC sample.

Impact on non-UC universities. The above journal-based classification method allows us to extend our analysis to non-UC universities in the much larger US sample. Table A2 also shows no significantly differential change in the publication level between dry and wet labs for non-UC universities. The coefficient, a precisely estimated zero, provides support for the external validity of our main finding. Unreported results also show similar null results across the university size distribution (measured by their total number of labs in the US sample). This suggests that the shock also did not translate into a large tax on research productivity for smaller institutions that may not have sufficient endowments or access to government funds to withstand potential disruptions of more stringent safety regulations.

7 Impact on the use of dangerous chemicals

Our theoretical model suggests that after the UCLA accident, investment in safety should have increased, either mandated by the university or due to voluntary adjustments by the labs themselves. In section 3, we provide evidence for a significant increase in lab safety practices based on historical reports (e.g., Gibson et. al. 2014). In this section, we examine the impact of the shock on a lab’s propensity to use dangerous chemicals. Documenting a shift away from dangerous chemicals not only provides additional evidence for an increase in lab safety, it also illustrates channels through which laboratory safety may shape the production function of science.

For this analysis, we complement the data used in the previous section with two additional datasets. The first comes from Laboratory Chemical Safety Summary (LCSS), which is publicly accessible via PubChem. LCSS provides a comprehensive list of hazardous (e.g., acute toxic, flammable, or explosive)

chemical compounds. The second complementary dataset comes from Scifinder, which is a proprietary chemistry database that documents the list of chemical compounds used by the research presented in a journal article. Because Scifinder restricts the number of entries that can be downloaded in total and each time, we limit this analysis to the 83 UCLA labs. Using the digital object identification (doi) information, we are able to find about 80 percent of the WoS articles in the UCLA sample in Scifinder.¹⁴ In Appendix Table A2, we replicate our analysis of the level of publications using the Scifinder instead of the WoS data, and the estimates are in line with those presented in Section 6.¹⁵

Merging these two datasets using the CAS registry numbers, which are unique identifiers for chemical substances, we are able to determine whether each article published by the UCLA researchers references dangerous chemicals. Our data indicate that both treatment and control labs reference dangerous substances in their publications. Specifically, UCLA dry labs publish 2.65 articles a year referencing dangerous compounds, whereas wet labs publish 3.30 articles per year referencing dangerous compounds.

Manual examination of a few publications supports the idea that most of the references made by dry labs capture the study of the theoretical properties of these compounds through mathematical modelling. For example, in 2011 the UCLA dry lab run by PI Anastassia Alexandrova published “On the mechanism and rate of spontaneous decomposition of amino acids” in the *Journal of Physical Chemistry B*. The study relies on Monte Carlo simulations and quantum mechanics to examine the properties of methylamine (CAS no 74-89-5), which is classified as a dangerous compound in the LCC database. The substance is labeled as extremely flammable (may cause severe skin burns and eye damage) and irritant (may cause skin and eye irritation). For wet labs, many of the compounds referenced by the wet labs typically capture the use of these chemicals in experiments. For example, in 2014 the UCLA wet lab run by PI Ohyun Kwon published a paper in the *Asian Journal of Organic Chemistry* titled “Phosphine-Initiated General-Base-Catalyzed Quinolone Synthesis.” The method section of this article describes an experimental procedure in which triphenylphosphine (CAS no 603-35-0) is mixed in a flask with other compounds and stirred. The compound is flagged as a dangerous substance in the LCC

¹⁴The fraction of unmatched papers appears fairly constant across sample years and is not driven by specific labs.

¹⁵There is a decline in publications after 2008 for wet lab scientists relative to dry lab scientists, but the estimated coefficients are small and statistically insignificant both in OLS and Poisson regressions. If anything, the magnitude of the effect is even smaller in this sample relative to the WoS sample.

database, which labels it as an irritant (may cause an allergic skin reaction) and health hazardous (may cause cancer).

We use this information in two ways. First, it allows us to construct dependent variables that are proxies for whether a wet lab uses dangerous chemicals in a given year. Second, this information also allows us to differentiate wet labs that handle substantial amounts of dangerous chemicals from wet labs that only work with these chemicals in a limited fashion. It seems reasonable to expect the effect of the shock to differ for labs that deal with dangerous chemicals with different intensities.

7.1 Dry versus wet labs

In Table 3, we present a series of regressions examining whether UCLA wet labs changed their propensity to publish research referencing dangerous compounds relative to dry labs after 2008. In these regressions, publications of dry labs that reference dangerous chemicals are used as a control for general trends of research interest in these compounds. In column 1, the dependent variable is the number of publications involving any kind of dangerous compounds. In the remaining columns, we examine the number of publications that involve specific types of compounds: acute toxic (column 2), explosive (column 3), and flammable (column 4). Across the specifications, we find that on average, wet labs do not reduce research referencing dangerous chemicals relative to dry labs. The coefficients of the interaction term are all statistically insignificant, and they are mostly small in magnitude. In unreported regressions, we confirm this finding using other hazard categories reported in the LCSS data, which include corrosive or irritant substances, and compressed gas.

7.2 Wet labs using dangerous chemicals more versus less frequently

The findings above show that, on average, UCLA wet labs are not significantly less likely to refer to dangerous chemicals after 2008 relative to dry labs. In this section, we contrast wet labs whose need for dangerous chemicals is more intense—which is defined by the fraction of Scifinder publications referring to dangerous chemicals during the pre-shock period 2004-08 being among the top 20% of the sample—to wet labs for which such need is less intense. In performing this analysis, we focus on labs run by PIs affiliated with UCLA for the entire sample period because we need a relatively long pre-period to define

a lab’s intensity of using dangerous chemicals. This leaves us with a sample of 42 wet labs, 8 of which are classified as heavy users of dangerous chemicals.

Table 4 presents several regressions exploring the differences between UCLA wet labs with more versus less intense need for dangerous chemicals. Consistent with our baseline finding, Column 1 shows that the publication level is not statistically significant between these two types of labs.¹⁶ However, column 2 shows a significant decline (at the 0.05 level) in the number of articles referring to dangerous chemicals by about 1.17 articles per year after the shock. Columns 3 to 5 examine different types of hazards. These regressions indicate that the decline found in column 2 is primarily driven by flammable substances. Wet labs with the most intensive need for dangerous chemicals, on average, experienced a decline of about 0.96 articles per year after 2008 relative to the other wet labs (significant at the 0.01 level).¹⁷

In Figure 2, we provide a graphical representation of the dynamic evolution in the differential use of flammable substances between UCLA wet labs with the most intensive need for dangerous chemicals and wet labs with less intensive need. There is no evidence of pre-trend differences between the two groups. After 2008, the coefficients become negative and significant (or marginally significant) in most years until 2014. The reduction in the reliance on flammable materials could be driven either by the increase in risk perception, which has led to a voluntary shift away from the use of such materials, or by an increasingly more stringent safety rules at the university, which may have made the use of such chemicals logistically more costly. The magnitudes are especially large in 2013 and 2014. Perception of personal liability for PIs may be especially high due to the criminal case against Harran by the Los Angeles District Attorney in 2011. UCLA, as a result of its own settlement, also substantially strengthened its safety programs and increased the number of mandated training classes and inspections in 2012 (Gibson et. al., 2014). The effects appear to be smaller and statistically insignificant after 2015, suggesting that the use of flammable substances recovered gradually after the implementation of safer practices.

We also examine whether the 2008 accident is associated with a change in the propensity to use

¹⁶This is measured with Scifinder articles. A decline of similar magnitude is observed when we measure total publications using the WoS data.

¹⁷Unreported regressions also confirm that flammable substances appear to be the main driver of the decline estimated in Column 2. The estimates for the other hazard types (corrosive substances, compressed gas, and irritant substances) are much smaller in magnitude and statistically insignificant.

compounds that are relatively unfamiliar to the lab. In particular, we generate a dummy variable capturing cases in which a lab used a compound that had not previously been referenced by any of the UCLA labs. Assuming that labs are more likely to be familiar with the properties of compounds already handled by themselves or by their local colleagues, working with these chemicals is likely to imply greater risk.¹⁸

The regressions in Table 5 show no evidence of a decline in the use of unfamiliar compounds, on average. However, separating first references to safe compounds (column 2) from dangerous compounds (column 3) shows that while there is no significant change in the use of unfamiliar safe compounds, there is a strong negative effect for unfamiliar dangerous compounds. The magnitude of the decline is about 40 percent of the mean level of the dependent variable. Column 4 confirms this result using the number of papers published by the lab in a given year that refer to unfamiliar dangerous compounds as the dependent variable. Also in this case, the estimate indicates a significant decline in the references to dangerous chemicals not previously used at UCLA.¹⁹

Taken together, the results in this section suggest that while the shock has had no broad impact on the use of dangerous chemicals among wet labs, we do see a significant shift away from these hazardous materials for the small set of labs that tend to use them most intensely. This is especially the case for flammable compounds and dangerous chemicals that are relatively unfamiliar to the researchers. These results provide evidence for an increase in safety induced by the shock and suggest that lab safety does have some effect on the scientific research process. That said, such an effect is localized to a small subsets of labs, rather than broadly affecting the entire institution.

8 Conclusion

In this paper, we study the relationship between lab safety and research productivity by examining the impact of a substantial increase in safety regulations at the University of California following the

¹⁸We classify a publication as the first use of a compound at UCLA if we cannot find a reference to that compound in previous publications (since 1998) of all the UCLA labs in our data.

¹⁹We also attempt to examine whether the shock changed the propensity of wet labs to discover new compounds (e.g., whether the publication is the first ever referencing a compound in the academic literature). This is a challenging exercise as publicly available data sources that provide comprehensive historical bibliographical data for compounds (such as PubChem) can only be matched to a small subset of the compounds in our data (roughly 20 percent). The available data indicate an extremely low propensity that wet labs in our sample discover new compounds for the first time.

shocking death of a research assistant in a UCLA lab fire in 2008. There are two key empirical findings. First, the 2008 shock did not lead to any statistically significant differences in yearly publications for wet labs, which are more affected relative to dry labs, which are less affected. Second, we find evidence for research practices that increase lab safety. In particular, wet labs that used dangerous chemicals more frequently before the accident reduced their usage of dangerous chemicals—especially for flammable materials and dangerous chemicals that were unfamiliar to the researchers—after the accident.

A common view in academic and policy debates is that “safety is a tax on research productivity.” Our results do not support this claim. Instead, they suggest that safety can shape the production of science through multiple and potentially off-setting mechanisms. In highly hazardous work environments, adequate safety practices are not only necessary for workers’ health and safety, they also appear to be important in enabling researchers to focus on research instead of worrying about accidents. This effect potentially counteracts any negative impacts of stronger safety regulations on research productivity. We acknowledge that, rather than a random policy change, the increase in safety regulations in our setting follows a high-profile accident, which is likely to have made such complementary effect between safety and research productivity especially salient. That said, accidents, for better or worse, are typically the impetuses for the strengthening of safety practices that we observe in practice. Finally, because the magnitude of the increase in safety regulations is substantial in our setting, our findings suggest that large drops in research productivity are unlikely in cases of less extensive, more common revisions to safety rules.

References

- [1] ACS (2012a), Creating safety cultures in academic institutions: A report of the safety culture task force of the ACS committee on chemical safety, report available at www.acs.org/content/acs/en/education/students/graduate/creating-safety-cultures-in-academic-institutions.html
- [2] ACS (2012b), Advancing graduate education in the chemical sciences, report available at www.acs.org/content/acs/en/about/governance/acs-presidential-commission-on-graduation-education-in-the-chemical-sciences.html
- [3] Arrow, Kenneth (1970) *Essays in the theory of risk bearing*, North Holland, Amsterdam
- [4] Azoulay, Pierre, Christian Fons-Rosen, and Joshua Graff-Zivin (2019) “Does science advance one funeral at a time?” *American Economic Review* 109: 2889-2920
- [5] Azoulay, Pierre, Joshua Graff-Zivin, and Jialan Wang (2010) “Superstar extinction,” *Quarterly Journal of Economics* 125: 549-589
- [6] Baruffaldi, Stefano and Fabian Gaessler (2021) “The returns to physical capital in knowledge production: Evidence from lab disasters,” Max Planck Institute for Innovation & Competition Research Paper 21-19
- [7] Basken, Paul (2012) “UCLA criminal charges focus universities’ attention on laboratory safety,” *Chronicles of Higher Education*, April 15
- [8] Baudendistel, Brian (2009) “Investigation report University of California, Los Angeles,” Case No. S1110-003-09, (Department of Industrial Relations, Division of Occupational Safety and Health, Los Angeles)
- [9] Bénabou, Roland, and Jean Tirole (2016) “Bonus culture: Competitive pay, screening, and multi-tasking,” *Journal of Political Economy* 124: 305-370

- [10] Benderly, Beryl (2009) “The burning question of laboratory safety,” *Science* online article available at www.science.org/content/article/burning-question-laboratory-safety
- [11] Boomhower, Judson (2019) “Drilling like there’s no tomorrow: Bankruptcy, insurance, and environmental risk,” *American Economic Review* 109: 391-426.
- [12] Coffey, Bentley, Patrick McLaughlin and Pietro Peretto (2020) “The cumulative cost of regulations,” *Review of Economic Dynamics* 38: 1-21
- [13] Daughety, Andrew and Jennifer Reinganum (2013) “Economic analysis of products liability: Theory,” In *Research handbook on the economics of torts*, Edward Elgar Publishing
- [14] Dawid, Herbert, and Gerd Muehlheusser (2022) “Smart products: Liability, investments in product safety, and the timing of market introduction,” *Journal of Economic Dynamics and Control* 134: 1-20
- [15] De Philippis, Marta (2021) “Multi-task agents and incentives: The case of teaching and research for university professors,” *Economic Journal* 131: 1643-1681
- [16] Fryer, Roland and Richard Holden (2013) “Multitasking, dynamic complementarities, and incentives: A cautionary tale,” Harvard University working paper
- [17] Galasso, Alberto and Hong Luo (2017) “Tort reform and innovation,” *Journal of Law and Economics* 60: 385-412.
- [18] Galasso, Alberto and Hong Luo (2021) “Risk-mitigating technologies: The case of radiation diagnostic devices,” *Management Science* 67: 3022-3040
- [19] Galasso, Alberto and Hong Luo (2022) “When does product liability risk chill innovation? Evidence from medical implants,” *American Economic Journal: Economic Policy* 14: 366–401
- [20] Gibson, James, Imke Schröder and Nancy Wayne (2014) “A research university’s rapid response to a fatal chemistry accident: Safety changes and outcomes,” *Journal of Chemical Health and Safety* 21: 18-26

- [21] Gray, Wayne (1987) “The cost of regulation: OSHA, EPA and the productivity slowdown,” *American Economic Review* 77: 998–1006
- [22] Greenstone, Michael, Richard Hornbeck, and Enrico Moretti (2010) “Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings,” *Journal of Political Economy* 118: 536-598
- [23] Greenstone, Michael, John List and Chad Syverson (2012) “The effects of environmental regulation on the competitiveness of US manufacturing,” National Bureau of Economic Research working paper 18392
- [24] Gowrisankaran, Gautam, Charles He, Eric Lutz, and Jefferey Burgess (2015) “Productivity, safety, and regulation in underground coal mining: Evidence from disasters and fatalities,” National Bureau of Economic Research working paper 21129
- [25] Hayden, Erika (215) “Scientific society withholds honour over lab death,” *Nature News*, December
- [26] Hersh, Erica (2017) “Five ways effective laboratory design impacts health, safety, and productivity,” electronic article available at www.hsph.harvard.edu/ecpe/five-ways-effective-laboratory-design-impacts-health-safety-and-productivity/
- [27] Holmstrom, Bengt and Paul Milgrom (1991) “Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design,” *Journal of Law Economics and Organizations*: 24-52
- [28] Jaravel, Xavier, Neviana Petkova, and Alex Bell (2018) “Team-specific capital and innovation,” *American Economic Review* 108: 1034-73
- [29] Kemsley, Jyllian (2009) “Learning from UCLA,” *Chemical & Engineering News* 87: 29-31
- [30] Langerman, Neal (2009) “Laboratory safety?” *Journal of Chemical Health and Safety*, 16: 49-50
- [31] Levine, David, Michael Toffel, and Matthew Johnson (2012) “Randomized government safety inspections reduce worker injuries with no detectable job loss,” *Science* 336: 907-911

- [32] Maxmen, Amy (2018) “Judge dismisses criminal charges against US university chemist over lab death,” *Nature* online article available at www.nature.com/articles/d41586-018-06726-1
- [33] Merlic, Craig (2013) “Enhancing laboratory safety - Academic laboratory safety in a post 12/29 World,” unpublished presentation deck - April 11
- [34] National Research Council (2014) *Safe science: Promoting a culture of safety in academic chemical research*, National Academies Press.
- [35] Oettl, Alexander (2012) “Reconceptualizing stars: Scientist helpfulness and peer performance,” *Management Science* 58: 1122-1140
- [36] Oi, Walter (1974) “On the economics of industrial safety,” *Law and Contemporary Problems* 38: 669-699
- [37] Porter, Michael (1990) *The competitive advantage of nations*, New York: Free Press
- [38] Silver, Andrew (2021) “Fatal lab explosion in China highlights wider safety fears,” *Nature* electronic publication available at www.nature.com/articles/d41586-021-03589-x
- [39] Stephan, Paula (2010) “The economics of science,” in *Handbook of the Economics of Innovation*, Vol 1, North-Holland
- [40] Thaler, Richard and Sherwin Rosen (1976) “The value of saving a life: Evidence from the labor market,” in *Household production and consumption*, National Bureau of Economic Research
- [41] Trager, Rebecca (2014) “UCLA chemist avoids prison time for lethal lab accident,” *Chemistry World*, 26 June
- [42] Van Noorden, Richard (2011) “A death in the lab,” *Nature* 472: 270-271
- [43] Van Noorden, Richard (2013) “Safety survey reveals lab risks,” *Nature* 493: 9-10
- [44] Venkatraman, Vijaysree (2010) “Conventions of scientific authorship,” *Science* 12: 243-244

- [45] Viscusi, Kip and Michael Moore (1993) “Product liability, research and development, and innovation,” *Journal of Political Economy* 101: 161-184
- [46] World Health Organization and International Labour Organization (2021) WHO/ILO Joint Estimates of the Work-related Burden of Disease and Injury, 2000–2016, technical report available at www.who.int/publications/i/item/9789240034945

Appendix: Solving the model

The maximization problem for the scientist before the shock is

$$\begin{aligned} \max_{s,r} r + \underline{\delta}s - \underline{\delta} - \frac{s^2}{2} - \frac{r^2}{2} - \rho sr \\ \text{s.t. } s \geq 0 \text{ and } r \geq 0. \end{aligned}$$

After the shock, this problem becomes

$$\begin{aligned} \max_{s,r} r + \bar{\delta}s - \bar{\delta} - \frac{s^2}{2} - \frac{r^2}{2} - \rho sr \\ \text{s.t. } s \geq \underline{s} \text{ and } r \geq 0. \end{aligned}$$

In this setup, the optimal research output before the shock is

$$r^*(\underline{\delta}, 0) = \begin{cases} \frac{1 - \underline{\delta}\rho}{1 - \rho^2} & \text{if } \underline{\delta} - \rho > 0 \\ 1 & \text{if } \underline{\delta} - \rho \leq 0 \end{cases}.$$

After the shock, the optimal research output becomes

$$r^*(\bar{\delta}, \underline{s}) = \begin{cases} \frac{1 - \bar{\delta}\rho}{1 - \rho^2} & \text{if } \frac{\bar{\delta} - \rho}{1 - \rho^2} > \underline{s} \\ 1 - \rho\underline{s} & \text{if } \frac{\bar{\delta} - \rho}{1 - \rho^2} < \underline{s} \end{cases}.$$

The above solutions imply that $\Delta r = r(\bar{\delta}, \underline{s}) - r(\underline{\delta}, 0)$ are:

$$\Delta r = \begin{cases} -\rho \frac{\Delta\delta}{1 - \rho^2} & \text{if } \underline{\delta} - \rho > 0 \quad \& \quad \frac{\bar{\delta} - \rho}{1 - \rho^2} > \underline{s} \\ \frac{1 - \bar{\delta}\rho}{1 - \rho^2} - 1 & \text{if } \underline{\delta} - \rho \leq 0 \quad \& \quad \frac{\bar{\delta} - \rho}{1 - \rho^2} > \underline{s} \\ -\rho\underline{s} & \text{if } \underline{\delta} - \rho \leq 0 \quad \& \quad \frac{\bar{\delta} - \rho}{1 - \rho^2} < \underline{s} \\ 1 - \rho\underline{s} - \frac{1 - \underline{\delta}\rho}{1 - \rho^2} & \text{if } \underline{\delta} - \rho > 0 \quad \& \quad \frac{\bar{\delta} - \rho}{1 - \rho^2} < \underline{s} \end{cases}.$$

Inspections of the first and the third cases immediately imply that $\Delta r < 0$ if $\rho > 0$ and $\Delta r > 0$ if $\rho < 0$. For the second case, notice that $\Delta r = \frac{1 - \bar{\delta}\rho}{1 - \rho^2} - 1 = \frac{(\rho - \bar{\delta})\rho}{1 - \rho^2}$ and that this happens when


if $\frac{\bar{\delta} - \rho}{1 - \rho^2} > \underline{s}$. The latter condition implies that $\rho - \bar{\delta} < 0$. Thus, we again have $\Delta r < 0$ if and only if $\rho > 0$. For the last case notice that this occurs when $\underline{s} \geq \frac{\bar{\delta} - \rho}{1 - \rho^2}$. This implies that if $\rho > 0$

$$\Delta r = 1 - \rho \underline{s} - \frac{1 - \underline{\delta} \rho}{1 - \rho^2} \leq 1 - \rho \left(\frac{\bar{\delta} - \rho}{1 - \rho^2} \right) - \frac{1 - \underline{\delta} \rho}{1 - \rho^2} = -\rho \frac{\Delta \delta}{1 - \rho^2} < 0.$$

Conversely, when $\rho < 0$, we have that


$$\Delta r = 1 - \rho \underline{s} - \frac{1 - \underline{\delta} \rho}{1 - \rho^2} \geq 1 - \rho \left(\frac{\bar{\delta} - \rho}{1 - \rho^2} \right) - \frac{1 - \underline{\delta} \rho}{1 - \rho^2} = -\rho \frac{\Delta \delta}{1 - \rho^2} > 0.$$

Figure 1: Annual treatment effects – total publications



Note: The figure plots the coefficients (and the 90% confidence intervals) of the interaction terms between year dummies and the wet lab dummy, which equals one if the lab conducts experiments using chemical and biological substances.

Figure 2: Annual treatment effects – publications referring to flammable chemicals



Note: The figure plots the coefficients (and the 90% confidence intervals) of the interaction terms between year dummies and a dummy capturing wet laboratories with high use of dangerous substances (i.e., wet labs above the 80th percentile in the use of dangerous chemicals). The control group includes wet labs below the 80th percentile in the use of dangerous chemicals.

Table 1 - Summary Statistics

Panel A: UC sample	obs.	mean	sd	min	max
Articles	6837	7.365	6.819	0	79
Wet Lab	6837	0.868	0.339	0	1
Year	6837	2010.827	3.931	2004	2017

Panel B: UCLA sample	obs.	mean	sd	min	max
Articles	986	7.677	7.278	0	57
Wet Lab	986	0.826	0.378	0	1
Year	986	2010.843	3.914	2004	2017

NOTES: Unit of observation is a lab-year. Panels A and B report summary statistics for UC and UCLA samples, respectively. Articles = the number of articles published by the lab in year t. Wet Lab = 1 if the lab conducts experiments using chemical and biological substances.

Table 2: Introduction of stricter lab safety protocols is not associated with changes in wet labs' publication levels relative to dry labs

	(1)	(2)	(3)	(4)	(5)
Dep. Var.	Articles	Highly cited articles	Articles	Articles	Articles/lab members
Wet Lab × After Accident	-0.280 (0.488)	0.003 (0.122)	-0.009 (1.042)	-0.002 (1.048)	0.200 (0.401)
Sample	Full	Full	UCLA labs	UCLA labs without Patrick Harran	UCLA labs
Year effects	YES	YES	YES	YES	YES
Lab effects	YES	YES	YES	YES	YES
Observations	6837	6837	986	976	986

NOTES: OLS regressions. Articles = the number of articles published by the lab in year t. Highly cited articles = the number of articles published by the lab in year t in the top decile of citations. Wet Lab = 1 if the lab conducts experiments using chemical and biological substances. After Accident = 1 if after year 2008. All regressions control for the total publications by the lab in the previous three years and the logarithm of the lab's tenure. Lab members = the number of unique local non-PI and non-faculty researchers listed as coauthors in papers published by the lab in year t. Robust standard errors clustered at the lab level in parentheses. * p < 0.10, ** p < 0.05, *** p < 0.01

Table 3: Introduction of stricter lab safety protocols is not associated with changes in references to dangerous chemicals in UCLA wet lab publications relative to UCLA dry labs

	(1)	(2)	(3)	(4)
Dep. Var.	Dangerous substances	Acute toxic substances	Explosive substances	Flammable substances
Wet Lab × After Accident	0.307 (0.914)	-0.115 (0.730)	-0.085 (0.102)	-0.120 (0.881)
Sample	UCLA labs	UCLA labs	UCLA labs	UCLA labs
Year effects	YES	YES	YES	YES
Lab effects	YES	YES	YES	YES
Observations	986	986	986	986

NOTES: OLS regressions. Acute toxicity substances = the number of articles referring to acute toxicity substances published by the lab's in year t. Explosive substances = the number of articles referring to explosive substances published by the lab in year t. Flammable substances = the number of articles referring to flammable substances published by the lab in year t. After Accident = 1 if after year 2008. All regressions control for the total publications by the lab in the previous three years and the logarithm of the labtenure. Robust standard errors clustered at the lab level in parentheses. * p < 0.10, ** p < 0.05, *** p < 0.01.

Table 4: The introduction of stricter lab safety protocols is associated with a reduction in the use of flammable substances in UCLA wet labs with high use of dangerous substances relative to other UCLA wet labs

	(1)	(2)	(3)	(4)	(5)
Dep. Var.	SciFinder articles	Dangerous substances	Acute toxic substances	Explosive substances	Flammable substances
High Use × After Accident	-0.307 (0.782)	-1.171*** (0.492)	-0.303 (0.390)	-0.030 (0.045)	-0.967** (0.477)
Sample	UCLA wet labs active between 2004 and 2017	UCLA wet labs active between 2004 and 2017	UCLA wet labs active between 2004 and 2017	UCLA wet labs active between 2004 and 2017	UCLA wet labs active between 2004 and 2017
Year effects	YES	YES	YES	YES	YES
Lab effects	YES	YES	YES	YES	YES
Observations	588	588	588	588	588

NOTES: OLS regressions. SciFinder articles = the number of articles in the the SciFinder database published by the lab in year t. Dangerous substances = the number of articles referring to dangerous substances published by the lab in year t. Acute toxicity substances = the number of articles referring to acute toxicity substances published by the lab in year t. Explosive substances = the number of articles referring to explosive substances published by the lab in year t. Flammable substances = the number of articles referring to flammable substances published by the lab in year t. High Use = 1 if the lab is in the top quintile in terms of publications referencing dangerous substances. After Accident = 1 if after year 2008. All regressions control for the total publications by the lab in the previous three years and the logarithm of the lab's tenure. Robust standard errors clustered at the lab level in parentheses. * p < 0.10, ** p < 0.05, *** p < 0.01.

Table 5: Introduction of stricter lab safety protocols is associated with a reduction in the use of relatively unfamiliar dangerous compounds

	(1)	(2)	(3)	(4)
Dep. Var.	Use of compounds new to UCLA	Use of safe compounds new to UCLA	Use of dangerous compounds new to UCLA	Articles referring to new dangerous compounds
High Use × After Accident	0.012 (0.053)	0.054 (0.065)	-0.174** (0.078)	-0.633** (0.247)
Sample	UCLA wet labs active between 2004 and 2017	UCLA wet labs active between 2004 and 2017	UCLA wet labs active between 2004 and 2017	UCLA wet labs active between 2004 and 2017
Year effects	YES	YES	YES	YES
Lab effects	YES	YES	YES	YES
Observations	588	588	588	588

NOTES: OLS regressions. Use of compounds new to UCLA = 1 if at least one of the compounds referenced in the lab publications in year t was never used before at UCLA. Use of safe compounds new to UCLA = 1 if at least one of the safe compounds referenced in the lab publications in year t was never used before at UCLA. Use of dangerous compounds new to UCLA = 1 if at least one of the dangerous compounds referenced in the lab publications in year t was never used before at UCLA. Articles referring to new dangerous compounds = the number of articles with dangerous substances first used at UCLA published by the lab in year t. High Use = 1 if lab in top quintile in terms of articles published using dangerous substances. After Accident = 1 if after year 2008. All regressions control for the number of compounds referenced by the lab in year t, total publications by the lab in the past three years, and the logarithm of the lab's tenure. Robust standard errors clustered at the lab level in parentheses. * p < 0.10, ** p < 0.05, *** p < 0.01.

Table A1: Laboratory safety and publication levels -- Robustness to alternative econometric models

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep. Var.	Articles	Citation weighted articles	Articles	Articles	Articles	Articles	Articles	Articles
Model	Poisson	Poisson	OLS	Weighted-OLS	OLS	OLS	OLS	OLS
Wet Lab × After Accident	-0.053 (0.061)	-0.176 (0.208)	-0.382 (0.619)	-0.464 (0.685)	-0.211 (0.486)	-0.246 (0.486)		
Endowment					0.001** (0.001)			
Chem PhDs					0.008 (0.009)			
Science & Engineering Grants					0.001 (0.001)			
Journal Based Wet Lab × After Accident							0.130 (0.540)	-0.043 (0.087)
Sample	FULL	FULL	UC labs active between 2004 and 2017	FULL	Institutions with available data	FULL	FULL	Non-UC US academic labs
Year effects	YES	YES	YES	YES	YES	YES	YES	YES
Lab effects	YES	YES	YES	YES	YES	YES	YES	YES
Institution-Year Effects	NO	NO	NO	NO	NO	YES	NO	NO
Observations	6828	6828	4564	6244	5563	6837	6837	38763

NOTES: Articles = the number of articles published by the lab in year t. Citation weighted papers = the number of articles weighted by citations received as of 2020 published by the lab in year t. Wet Lab = 1 if the lab conducts experiments using chemical and biological substances. After Accident = 1 if after year 2008. Endowment = total market value of the endowment held by the institution in year t (in 1,000s USD). Chem PhDs = the number of individuals receiving a research doctorate in the field of chemistry in the institution in year t. Science & Engineering Grants = the amount of federal science and engineering (S&E) funding received by the institution in year t (in 1,000s USD). Journal Based Wet Lab defines wet versus dry labs using classification of journals and publications between 2004 and 2008. All regressions control for the total publications by the lab in the previous three years and the logarithm of the lab's tenure. In Column 8, the DV is constructed using the subset of WoS journals specialized in chemistry and ten multidisciplinary science journals with the highest impact factors. Robust standard errors clustered at the lab level in parentheses. * p < 0.10, ** p < 0.05, *** p < 0.01.

Table A2: Laboratory safety and publication levels -- Robustness using SciFinder data

	(1)	(2)	(3)
Dep. Var.	Articles	SciFinder articles	SciFinder articles
Model	OLS	OLS	Poisson
Wet Lab × After Accident	-0.009 (1.042)	-0.041 (1.165)	-0.036 (0.087)
Sample	UCLA labs	UCLA labs	UCLA labs
Year effects	YES	YES	YES
Lab effects	YES	YES	YES
Observations	986	986	986

NOTES: Articles = the number of articles published by the lab in year t. SciFinder articles = the number of articles recorded by the SciFinder database published by the lab in year t. Wet Lab = 1 if the lab conducts experiments using chemical and biological substances. After Accident = 1 if after year 2008. All regressions control for the total publications by the lab in the past three years and the logarithm of the number of years that the lab has existed as of year t. Robust standard errors clustered at the lab level in parentheses. * p < 0.10, ** p < 0.05, *** p < 0.01.