'Dead Man Working': A Place-based Approach to Occupational Safety and Health

Augusto Cerqua[†], Costanza Giannantoni[*], Marco Letta[‡], Gabriele Pinto[§]

Abstract

Despite increasingly stringent regulations, there has been a concerning stagnation in reducing workplace fatalities recently. Can place-based targeting help? By coupling machine learning techniques with comprehensive data from Italy, we develop a place-based approach to workplace fatalities. Harnessing accurate forecasts, we construct a granular risk map and compare it to the allocation of on-site inspections and public subsidies for occupational safety, uncovering limited overlap. Counterfactual estimates reveal that current measures are effective only in areas flagged as high-risk by ex-ante machine predictions. AI-powered territorial targeting can reduce the incidence of this chronic issue while lowering the costs of policy implementation.

JEL codes: H51, I18, J28, R50.

Keywords: occupational safety and health; workplace deaths; algorithmic decision-making; machine learning; place-based policies.

Acknowledgments: We thank Chiara Goretti, Melisso Boschi, and Laura Maragnani for their valuable advice on an earlier version of this work. We are grateful to Antonella Milieni of INL for support with the inspection data and to Adelina Brusco and Alessandra Filottrano for assistance with the INAIL data. We express our thanks to Nicolò Danieli and Filippo Nardoni for their support in the data collection process. We also thank participants to the 2025 EPCS Annual Conference for useful comments and suggestions. Augusto Cerqua and Marco Letta gratefully acknowledge financial support from the Italian PRIN 2022 grant for the project "Strengthening TARgeting and Guidance with Actionable and Timely Evidence (STARGATE) for the implementation of the Italian National Recovery and Resilience Plan", CUP: F53D23003220006.

† Department of Social Sciences and Economics, Sapienza University of Rome, Italy. Email: augusto.cerqua@uniroma1.it
* Department of Social Sciences and Economics, Sapienza University of Rome, Italy. Email: costanza.giannantoni@uniroma1.it
‡ Department of Social Sciences and Economics, Sapienza University of Rome, Italy. Email: marco.letta@uniroma1.it
§ Department of Social Sciences and Economics, Sapienza University of Rome, Italy. Email: gabriele.pinto@uniroma1.it

1. Introduction

Workplace fatalities, a primary target of occupational safety and health regulation, remain a significant problem worldwide, even in the most developed countries. According to the International Labour Organization (ILO), nearly three million people died in 2019 due to work-related accidents and diseases, representing an increase of more than 5% compared to 2015 (ILO, 2023). This lack of progress has made the issue a policy priority across many countries. The European Union (EU), for example, set out an ambitious *Strategic Framework on Health and Safety at Work 2021-2027*, based on the *Vision Zero* approach, which aims to eliminate work-related deaths entirely (EASHW, 2023).

However, despite its relevance in political debates and continued prominence in the media, this issue has received limited attention from researchers in public policy and economics. Perhaps even more surprisingly—considering both the heterogeneity of occupational fatalities across jobs and sectors and the high degree of clustering and specialization of modern local economies—there is a notable lack of studies examining the issue from a territorial perspective. In this work, we develop a place-based approach to identify hotspot areas at the highest risk of workplace deaths in advance, pinpoint the main territorial predictors of this phenomenon, and target and support policy interventions aimed at enhancing occupational safety and health.

The economic literature on these issues is relatively scarce and thus far has focused only on specific aspects and relationships. One body of work has investigated linkages with labor market characteristics, such as the relationship between worker sorting and the risk of on-the-job fatalities (DeLeire and Levy, 2004); the impact of immigration on natives' workplace conditions (Dillender and McInerney, 2020); the nexus between worker deaths and firms' demand for incumbent workers and new hires (Jäger and Heining, 2022); and the unintended effects of temporary contracts and minimum wages on injury severity (Picchio and van Ours, 2017, Liu et al., 2024). Other studies have explored the role of economic and market conditions, specifically examining the consequences of positive price shocks on workplace injury rates (Charles et al., 2022), the effects of import competition on workers' health (Mcmanus and Schaur, 2016), and the relationship between recessions and workplace safety (Boone and van Ours, 2006). Other analyzed aspects include the role of safety regulations, such as the deterrence effects of publicizing violations of workplace safety (Johnson, 2020), the use of randomized safety inspections to estimate the value of a statistical life (Lee and Taylor, 2019), and the impact of environmental conditions on workplace accidents (Drescher and Janzen, 2025). Finally, more closely related to our paper, a recent study by Johnson et al. (2023) used micro-level establishment data on workplace injuries and combined randomization and causal forests to assess the impacts of inspections and evaluate alternative targeting strategies to random inspection allocation. They found that the use of data-driven allocation criteria for targeting inspections could have prevented as many as twice as many injuries and that targeting based on predicted injuries is more effective than targeting based on estimates of the treatment effects of inspections.

None of these microeconomic studies adopt a territorial approach to explore occupational safety and health outcomes. However, given the spatial heterogeneity of modern local labor markets, in terms of the composition of jobs, workers, sectors, etc., geographic patterns and territorial data may hold intrinsic predictive power and provide valuable insights concerning these phenomena. Similarly, the ex-post, retrospective setting of all these studies—with the only exception of Johnson et al. (2023)—is not particularly informative about *whom to treat* and the related development and deployment of preventive and deterrent measures, but only about the effects, or lack thereof, of the realized policies.

We take a different route and develop a prospective approach to study workers' health and safety through an urban and public economics lens, which complements these micro-level studies by exploring the problems of vulnerability detection and public resources allocation from an ex-ante and place-based, rather than ex-post and firm-based, perspective.[1] To this end, we combine state-of-the-art machine learning (ML) methods with granular territorial data to predict and target areas most at risk of workplace fatalities.

Our exclusive focus on deaths—rather than injuries—as the outcome variable to study occupational safety and health is uncommon[2] and particularly noteworthy, as nonfatal work accidents are generally subject to vast and hard-to-detect underreporting issues, especially in contexts characterized by sizable shadow economies where the incentive not to report is particularly strong (OECD/ILO, 2019). Mismeasurement in outcome data is always a threat in predictive problems, as it can lead ML to automate, reproduce, and even amplify errors in the recording of imperfect data (Mullainathan & Obermeyer, 2017). On the other hand, underreporting bias is likely minimal for deaths, given the considerable challenges associated with not reporting fatalities. This implies that our analysis relies on more accurate data than most previous studies focused on non-fatal injuries, thereby minimizing the challenges associated with predictive tasks based on inaccurate training outcome data (Cannings et al., 2020).

Using detailed workplace fatality data from Italian local labor markets (LLMs) for the period of 2017-2023, the ML algorithms reveal significant spatial heterogeneity in the risk of fatalities on the job and accurately forecast the local number of workers' deaths in the held-out years. We leverage these accurate machine predictions to construct a granular risk map, which can inform place-based policy interventions aimed at enhancing occupational safety and health. By comparing the ML risk map with the actual distribution of public policies, we demonstrate that the allocation of on-site inspections and public subsidies for occupational safety is currently not more prevalent where it is most needed.

[1] To our knowledge, the only study that adopted an ex-ante approach to the issue of occupational accidents is a forecasting exercise by Melchior et al. (2021), which relies on parametric time series methods to forecast work-related mortality rates in Brazil, without any focus on public policy analysis or place-based aspects.

[2] Among the aforementioned works, only Lee and Taylor (2019) employed fatality data as the outcome variable.

We then explore the implications of this mismatch by i) conducting a descriptive analysis to highlight the main differences in terms of demographic characteristics and economic structure between high-risk hotspots and other areas and ii) undertaking an ex-post evaluation of the impact of on-site inspections and public subsidies on the number of workplace deaths via double/debiased machine learning (Chernozhukov et al., 2018). Our findings indicate that, on average, an increase in the number of on-site inspections or in the amount of public subsidies has a null and insignificant impact. However, the effect becomes negative and statistically significant for 'red flag' areas independently classified as having the highest risk by the ex-ante ML forecasts, suggesting a high-risk/high-benefit linkage regarding the effects of these public policies.

Overall, these results suggest that the current policy framework for occupational safety and health might be improved. Implementing an ML-powered place-based policy allocation rule would simultaneously increase the effectiveness of interventions—significantly enhancing workplace safety by better directing on-site inspections and public subsidy assignments—and lower the costs of policy implementation. A back-of-the-envelope calculation based on the most conservative estimates suggests that replacing current allocation rules with machine forecasts would prevent approximately 86 fatal accidents per year, or approximately 10% of the annual on-the-job deaths registered in the period we study. Such a reduction could enhance Italian policymakers' ability to consistently achieve their stated goal of reducing the total number of workplace deaths and making progress toward the achievement of the ambitious EU *Vision Zero* target (see EASHW, 2023), while at the same time decreasing the overall cost of the policy by concentrating efforts in vulnerable hotspots. The Italian case is relevant because the country ranked 8^{th} among the EU-27 countries in 2021 in terms of deaths at work, reporting 3.3 workplace fatalities per day, equivalent to 2.7 deaths per 100,000 people employed, a figure 50% higher than the European Union average.[3] For this reason, the issue is currently highly salient in the media and prominent in policy debates.[4] Italy is also a country where underreporting of workplace injuries is notably high due to the prevalence of informal work and the shadow economy, further motivating the use of fatalities over injuries.[5]

In terms of contribution, our work can be placed at the intersection of three related but previously unintegrated strands of literature: the aforementioned occupational health and safety literature; the urban economics and public policy literature on place-based policies (see, among many, Blesse and Diegmann, 2022; Cerqua and Letta, 2022; Cerqua and Pellegrini, 2022; Criscuolo et al., 2019; Faggio 2019; Gallé et al., 2024; Freedman, 2015; Lu et al., 2019; Mayer et al., 2017; Schweiger et al., 2022); and the body of works investigating the relationship between artificial intelligence (AI) and occupational and workplace outcomes (e.g. Autor, 2015; Brynjolfsson and Mitchel, 2017; Brynjolfsson et al., 2018; Felten et al., 2018). While the literature on place-based policies has largely concentrated on ex-post evaluation, mixed evidence concerning their effects has drawn attention to their targeting mechanisms, following the notion that, for place-based policies to succeed, they must target the right areas (Corinth et al., 2025). We show that ML can be leveraged for such targeting, shifting the paradigm from ex-post evaluation to ex-ante analysis of targeting mechanisms in place-based policymaking. Concerning the occupational and workplace impact of AI, we complement this literature, which mainly explores the workplace consequences of AI-related disruptions and the impact of automation on productivity and workers' tasks and skills, to show that AI can contribute to the design and implementation

[3] The full data from Eurostat is available at https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Accidents_at_work_statistics.

[4] As an example, in his 2024 end-of-the-year speech, the Head of State Sergio Mattarella referred to workplace fatalities as a policy priority and commented on a recent tragedy, in which five workers died in a depot explosion in Tuscany, with the following statement: "Words of outrage are no longer enough: action is needed, with responsibility and severity. All fatal accidents can and must be prevented" (our translation).

[5] The underreporting of injuries in Southern Italy is evident when comparing the geographical patterns of on-the-job deaths versus all injuries (including non-fatal ones), which are reported in Figure A.1 of Supplemental Appendix A. In addition, Antonelli et al. (2024) reveal the existence of a relevant under-notification phenomenon of accidents at work with respect to moderate accidents, which is higher especially for the southern regions of Italy.

of workplace policies aimed at increasing on-the-job safety for workers.

Methodologically, our work expands the growing literature in economics that harnesses the predictive power of ML techniques to improve decision-making (Kleinberg et al., 2015), informing program design and resource allocation across diverse sectors and domains, such as corruption (Ash et al., 2024), taxation (Andini et al., 2018; Battaglini et al., 2024), local public finance (Antulov-Fantulin et al., 2021), education (Athey et al., 2025), energy efficiency (Christensen et al., 2024), and health (Carrieri et al., 2021; Glaeser et al., 2016). We build upon these studies by focusing on a still unexplored issue and go beyond them by proposing a ML-based approach explicitly suited for panel data forecasting owing to a rigorous panel cross-validation strategy.

More generally, the problem of *whom to target* arises in many settings (Athey et al., 2025), some of which are inherently complex and characterized by a low signal-to-noise ratio, making it challenging for traditional approaches to recover policy-relevant signals. In this respect, our work introduces an original approach based on machine-guided territorial policies that can also be adopted to identify vulnerability hotspots and inform public decision making to address societal issues in other domains that are rarely examined from a place-based perspective. For example, complex phenomena such as gender violence, suicides, and gambling addiction might also exhibit systematic territorial patterns and benefit from a targeting analysis coupling AI tools with granular territorial data.

The rest of this paper is arranged as follows. Section 2 presents the institutional framework. Section 3 describes the data and methods, and Section 4 presents the results. Section 5 concludes.

2. Institutional framework

The Italian system for ensuring workplace safety involves multiple institutional actors, primarily from the Ministry of Health and the Ministry of Labor. Three main activities can be identified: workplace supervision and inspections, insurance against workplace injuries, and public incentives for firms' investments in occupational safety. Prior to 2015, workplace inspections were shared between the Ministry of Labor, the Italian National Institute for Insurance against Accidents at Work (INAIL), and the Italian National Social Security Institute (INPS). Then, Law no. 149/2015 established the National Inspectorate for Labor (*Ispettorato Nazionale del Lavoro* – INL), which subsequently centralized these inspection responsibilities.[6] The INL exercises and coordinates surveillance across the national territory, covering work, contributions, mandatory insurance, social legislation, and health and safety protection in workplaces.[7] Inspection activities are complemented by local health units (*Aziende Sanitarie Locali*), whose activity depends directly on regional governments. As of December 31, 2022, the INL coordinated 3,983 inspectors, comprising *i)* 2,412 civilian INL inspectors, *ii)* 884 INPS inspectors, *iii)* 210 INAIL inspectors, and *iv)* 477 military personnel from the Arma, partially assigned to judicial police functions.[8] Inspections may originate from reported concerns (with limited anonymous reporting) or through proactive "intelligence activities" initiated by the INL to identify potential high-risk cases. These investigations can proceed in the absence of the employer and culminate in a detailed report. In 2022, the INL carried out 82,183 inspections, revealing 314,069 violations of worker protection regulations, i.e., 3.8 violations per inspection, including 19,932 cases of undeclared labor.

In addition to workplace inspections, the workplace safety system includes insurance coverage for occupational injuries. This function is overseen by INAIL, which manages the mandatory insurance scheme for workplace injuries and

---

[6] The aim of the law was to overcome the fragmentation of inspection activities among the different institutions that monitor compliance with labor contracts (social security, black practices, etc.), and health and safety conditions (compulsory insurances, regular work environments, etc.), concentrating them in a single agency (Colombo et al., 2019).

[7] The INL, which operates under the direction of the Italian Ministry of Labor, covers the whole Italian territory with the exception of three special-status regions: Aosta Valley, Sicily and Trentino-South Tyrol, where the functions of the INL are carried out by regional and provincial bodies through the Local Health Units.

[8] See here (in Italian). Moreover, in July 2024, the INL announced the recruitment of a total of 750 non-executive personnel on a permanent basis, to be classified in the functional area of technical health and safety inspection officers.

occupational diseases while also promoting precautionary measures to enhance worker protection and support workers' "rehabilitation and reintegration to social life and work" (Campo et al., 2020). To further enhance workers' protection, in 2010, INAIL launched a state-aid scheme (named "ISI calls") to support firms' (especially SMEs) investments to improve occupational safety and health performance. With a total funding allocation exceeding 3 billion euros from 2010, the ISI calls represent the most significant incentive measure dedicated to the protection of health and safety at the European level (Castaldo et al., 2023). The rationale for this policy is that SMEs can be constrained in investments for workplace safety, making them underequipped (Barile et al., 2024). This is particularly relevant for the Italian context, where SMEs represent over 75% of the industrial sector. Eligible proposals may include investment projects, training programs, and the adoption of organizational models and social responsibility frameworks. The recognized incentive consists of a nonrepayable capital grant covering 50% to 75% of the project costs, with grants ranging from a minimum of €5,000 to a maximum of €100,000. The assignment process is conducted at the regional level through calls for tenders, which are based on eligibility criteria that do not consider the risk of workplace accidents.[9] By not taking accident risk into account, the geographical allocation of funds is unlikely to optimize the reduction of workplace injuries and fatalities.

3. Data and Methodology

3.1 Data

3.1.1 Data on workplace fatalities

For the empirical analysis, we use data on occupational fatalities provided by

[9] To submit an application, businesses must meet the eligibility requirements. These include that the firm must not be undergoing voluntary liquidation or subject to any insolvency proceedings, must comply with insurance and social security contribution obligations, must not have requested or received other public funding for the proposed project, nor have been granted funding approval under the three previous calls for proposals.

INAIL, the authority responsible for managing insurance and compensation for work-related injuries and fatalities. The workforce covered by INAIL is quite extensive, encompassing approximately 80% of total employment; however, it does not include certain special categories of workers, such as firefighters, police officers, military personnel, and journalists, as they are covered by other insurers (Filomena and Picchio, 2024). INAIL provides data at the municipal level on the number of workers' fatalities for the period 2017-2023, split by gender, macroeconomic sector (agriculture, public sector, industry and services) and place of occurrence ("in itinere" or "on-the-job"). Occupational fatalities include both deaths occurring at the work site ("on-the-job") and those occurring during the commute to work (so-called "in itinere" fatalities). The latter also encompass incidents between the workplace and a lunch location or between two different workplaces. In Italy, the majority of occupational fatalities occur at the workplace itself: over the period 2017-2023, 76.9% of these deaths fall into this category.

Table 1 reports the number of occupational deaths by sector of activity in Italy for the period 2017-2023, split by "on-the-job" and "in itinere". The table clearly shows the sharp increase in the number of "on-the-job" deaths in 2020 and 2021 (Panel B) due to the COVID-19 pandemic. In the same years, it was possible to observe a decrease in the number of "in itinere" deaths (Panel C) due to the reduction in mobility and economic activities caused by the pandemic.[10] Nevertheless, the impact of the COVID-19 pandemic on occupational deaths was limited to 2020 and 2021. According to the INAIL Annual Report (2022), the number of deaths directly or indirectly attributable to COVID-19 in 2022 was only 8. The similarity in the number of occupational deaths between 2017-2019 and 2022-2023 suggests that the COVID-19 shock can be considered a temporary

[10] During the COVID-19 pandemic, Italy implemented stringent measures to curb the spread of the virus, including the closure of non-essential industries. In March 2020, the Italian government issued a decree ordering the shutdown of all non-essential factories and businesses, significantly affecting the country's workforce. This decision was part of a broader national lockdown, one of the first and most severe in Europe, which also included restrictions on movement and public gatherings. The closure impacted millions of workers across various sectors, from manufacturing to retail, as the government aimed to protect public health and prevent the healthcare system from being overwhelmed.

shock, after which conditions returned to “business as usual”, at least concerning worker safety.

Table 1: Yearly number of occupational fatalities in Italy

Panel A – Total number of occupational fatalities

	2017	2018	2019	2020	2021	2022	2023
Industry and Services	981	1,122	1,044	1,503	1,228	1,073	978
Agriculture	163	152	171	138	148	137	133
Public Sector	34	20	24	82	59	37	36
Total	1,178	1,294	1,239	1,723	1,435	1,247	1,147

Panel B – “On-the-job” workplace fatalities

	2017	2018	2019	2020	2021	2022	2023
Industry and Services	705	796	748	1,298	979	775	740
Agriculture	139	118	144	120	127	116	121
Public Sector	17	7	12	74	49	20	21
Total	861	921	904	1,492	1,155	911	882

Panel C – “In itinere” fatalities

	2017	2018	2019	2020	2021	2022	2023
Industry and Services	276	326	296	205	249	298	238
Agriculture	24	34	27	18	21	21	12
Public Sector	17	13	12	8	10	17	15
Total	317	373	335	231	280	336	265

Notes: The table reports the number of occupational fatalities in Italy over time by sector of activity. The data are provided by INAIL. Table A.1 in Supplemental Appendix A presents the number of deaths by year and sector before and after the imputation for the years 2020 and 2021.

The primary goal of the targeting analysis is to identify the most significant territorial predictors of “on-the-job” (workplace) fatalities in the near future and to create a risk map for policy implementation during “normal times.” However, workplace deaths caused by COVID-19 pose a challenge, as they skew the number of workplace deaths observed in “normal times,” particularly in certain areas (see Cerqua et al., 2021). Failing to account for these COVID-19-related workplace deaths would result in ML algorithms learning and forecasting non-systematic patterns associated with a temporary shock and possibly erroneous selection of the most relevant territorial predictors. A risk map based on such estimates would be biased and could undermine rather than support policy efforts (see Supplemental Appendix B for a more detailed discussion on the consequences of not accounting for COVID-19-related workplace deaths). Consequently, before conducting the analysis, we impute the number of "on-the-job" deaths for the areas most impacted by the pandemic in 2020 and 2021. This modeling choice follows the approach developed by the World Health Organization (WHO) in imputing all-cause death data in the no-pandemic counterfactual scenario to provide official estimates of the global excess death toll associated with the COVID-19 pandemic (see the two publications by the WHO team, Knutson et al., 2023; Msemburi et al., 2023). We emphasize that this imputation is conducted solely to enable the models to learn the true data-generating process during ordinary times (we report more details on the imputation process in Supplemental Appendix B). As detailed below, the out-of-sample forecasting ability of our models, our primary focus, will be evaluated exclusively with real-world, non-imputed data.

We have chosen the LLM level as the most appropriate unit of spatial analysis.[11] Each LLM is an aggregation of two or more neighboring municipalities (13 on average), defined by the Italian National Institute of Statistics (Istat) on the basis of daily commuting flows from the place of residence to the place of work. Opting for LLMs rather than municipalities as the unit of spatial analysis achieves a compromise between data granularity and the relative rarity of workplace deaths. As a result, we use yearly LLM data spanning the period from 2017 to 2023. We cover all 610 Italian LLMs and use the number of on-the-job deaths as the main outcome variable.

3.1.2 Set of predictors

The initial set of predictors comprises 160 variables collected at the LLM level. The set includes a detailed description of the industrial structure, macro-regional dummies, labor market characteristics, socio-economic and demographic data, features of the housing market, local politics, and detailed information on several

[11] The criteria used to determine Italian LLMs are similar to those used to define Metropolitan Statistical Areas in the US or Travel to Work Areas in the UK.

aspects of workplace deaths. In particular, we use as predictors the lagged values of the number of "in itinere" deaths (overall, for each sector and by gender) and the number of "on-the-job" deaths (overall, for each sector and by gender) (see Table A.2 in Supplemental Appendix A for a detailed description of the variables and their sources and Table A.3 for the descriptive statistics). In this set of covariates, we included the first lag of all time-varying predictors (including the dependent variable), along with their first lagged differences, to incorporate temporal dynamics and improve the forecasting model. This implies that we collapsed the original 2017–2023 dataset into a dataset covering the period 2019–2023.

3.1.3 Inspections and public subsidies

Finally, we collected data on the number of inspections from the INL and the amount of state aid received annually through the ISI calls by each territory to support firm investments in occupational safety and health. Owing to privacy concerns, these variables are only available at the provincial level. We use data on the number of inspections and the amount of public subsidies to assess whether public efforts are currently deployed in the most effective way to reduce workplace fatalities. See Table A.4 in Supplemental Appendix A for a detailed description of the variables used in the ex-post analysis and their sources and Table A.5 for the descriptive statistics.

An important caveat lies in the partial nature of the measures analyzed: inspections and subsidies do not cover all policy interventions implemented in Italy to address occupational safety and health. These specific dimensions were chosen because they are two of the most relevant policies, and we could not access data on other policies. It is therefore crucial, both when interpreting the results and in assessing external validity, to recognize that while highly relevant, they are not the only measures enacted to reduce workplace fatalities by Italian policymakers.

Regarding the ISI public subsidy data, additional caution is warranted due to the multifaceted nature of subsidies available to firms in Italy. Some of these subsidies, although not explicitly linked to occupational safety and health, may still contribute to improving workplace safety for employees. As a result, the distribution of public subsidies for occupational safety and health examined later in this study complements other forms of subsidies that are excluded from our analysis. This should be carefully considered when evaluating the findings of both the targeting and counterfactual analyses.

3.2 Methodology

The primary objective of this empirical analysis is to develop an ML pipeline to forecast the number of on-the-job fatalities in a given year for all LLMs. If these forecasts prove reliable, policymakers could leverage them to focus resources on the areas most at risk. To achieve this aim, it is essential to treat available data in a way that can resemble a real-world situation. In this regard, applying a standard ML pipeline to panel data is problematic owing to their peculiar nature, as they are inherently characterized by substantial cross-sectional and temporal autocorrelation, compared with the type of data for which ML techniques were originally developed. To address the challenges involving the application of ML routines to panel data, such as cross-validation (Arkhangelsky and Imbens, 2024), our approach is based on the following key departures from the standard ML routine:

- We use only lagged values of predictors in the ML pipeline to maintain the integrity of forecasts, as forecasting future data points using exclusively past information is key.[12] To forecast future values of a

[12] To illustrate this point, consider the following example: the number of deaths on the job is mechanically correlated with the number of workers. Imagine a massive employment shock occurs, causing many people to lose their jobs in a given year. Including the contemporaneous value of such a predictor would lead the algorithm to forecast a drop in the number of deaths due to the mechanical correlation learned from past data (i.e., if the algorithm sees a drop in the key predictor, which is highly and positively correlated with the outcome, it predicts an analogous drop in the outcome). However, assuming the shock could not be

variable, only information available at the time the forecast is made can be used; i.e., the forecasting ability of a model must be evaluated by generating forecasts over some past period only using data known at each forecast origin (Petropoulos et al., 2022).

- We diverge from the conventional ML approach, which typically uses a random division of the sample into training and testing sets. Instead, we implement a non-random, time-based splitting criterion to divide our sample into two disjoint sets (Cerqua, Letta, and Menchetti, 2024). This methodological choice is driven by the need to preserve the temporal sequence of the dataset and prevent time-dependent data leakage. Such leakage would occur if a random split allowed future observations to be included in the training set, thereby inappropriately influencing predictions of past events in the testing set (Cerqua, Letta and Pinto, 2024). Owing to these adjustments to the ML forecasting pipeline, we can obtain unbiased forecasts that are not overoptimistic with respect to the true out-of-sample performance on future data.

- For the same reason related to data leakage issues, we replace the standard random k-fold cross-validation (CV) routine typically applied in the ML literature (Hastie et al., 2009) with a panel CV strategy, which is adapted from the time series forecasting literature and fully described below.

Our starting point is the balanced panel dataset at the LLM-level from 2017 to 2023 described in Section 3.1. We employ five ML algorithms: two linear—LASSO and partial least squares (PLS)—two tree-based—random forest and stochastic gradient boosting—and the Long Short-Term Memory neural networks (LSTMs).[13] In the main analysis, we train these ML algorithms using only data

forecasted based on past data, the ability to forecast is an illusion due to the inclusion of a predictor contemporaneous to the outcome.

[13] LASSO is a regression analysis linear method that performs both variable selection and regularization. It works by imposing a constraint on the sum of the absolute values of the model coefficients, effectively

up to 2022 and use the value of the dependent variable in 2023 only in the held-out testing set to assess the final accuracy of the forecasts.[14] During the training phase, we employ a panel CV based on an expanding forecasting origin (Hyndman & Athanasopoulos, 2021), in which we iteratively train ML algorithms on expanding training sets from earlier periods and test their performance on sliding testing sets from later periods to select hyperparameters that optimize their performance in forecasting future outcome data points. In particular, as shown in Figure A.2 in Supplemental Appendix A, we use the first fold (2019) as the training set and the second fold (2020) as the test set; then the first two folds as the training set (2019 and 2020) and the third fold (2021) as the test set; and finally the first three folds (2019, 2020, and 2021) as the training set and the fourth fold (2022) as the test set. We run the ML algorithms—using a grid search for hyperparameter configuration—across all three training sets, test their performances on the corresponding test sets, and select the hyperparameters of each algorithm that result in the best average performance. The tuned models are then applied to the full 2019-2022 dataset to forecast 2023 outcomes out of sample, and their performance is evaluated on these data to determine the winner of the horse race. Since 2023 might just be a 'lucky year' in terms of forecasting performance, in a separate analysis, we repeat the same approach but exclude data from 2023, using only predictors in 2020 and 2021 to forecast 2022 outcomes on the testing set. This additional check ensures that the ML forecasting pipeline

shrinking some coefficients to zero. PLS reduces dimensionality by projecting predictors and responses onto orthogonal components that maximize covariance, thus improving predictive power. Random forest is an ensemble learning technique that constructs many decision trees during training and aggregates their predictions to increase out-of-sample accuracy by reducing overfitting risk. Each tree is built using only a random subset of the training data and of the predictors at each candidate split. The final prediction is determined by averaging the outputs of all the trees. Stochastic gradient boosting is also a tree-based ensemble technique, but it works in a sequential manner, where each new tree works on the residuals, i.e., the errors of the previous ones, and improves on gradient boosting. The final output is a weighted sum of all tree predictions. For a more detailed description of these algorithms, see Hastie et al. (2009). LSTMs are a specialized type of recurrent neural networks (RNNs) adept at learning order dependence in sequence prediction problems using textual or time series data.

[14] It is worth noting that in situations with only a few time periods available, time-series techniques like those adopted by Melchior et al. (2021) are not feasible, as they require at least several dozen time periods for effective implementation.

produces accurate forecasts for any held-out year.[15] The detailed implementation process is reported in Table 2.

Table 2: ML forecasting pipeline

Preliminary step

Data splitting. We split the full dataset into a training set with data on Y from 2019 to 2022 and a testing set with data on Y in 2023.

Training phase – Use only the training set

1. Algorithm selection. We select five supervised ML algorithms: 1) LASSO; 2) Partial Least Squares; 3) stochastic gradient boosting; 4) random forest; 5) LSTMs. As we are agnostic about the functional form of the underlying data-generating process, we opt for a mix of non-linear and linear models.

2. Principled input selection. We build an initial LLM dataset with 160 predictors on the basis of literature insights and data availability (see the Data section). From this dataset, we then keep only the most important predictors selected by a preliminary random forest run on the training data (see step below), following the approach proposed by Athey and Wager (2019). However, for each algorithm, we also report the estimates obtained using all 160 predictors. Figure A.3 in Supplemental Appendix A presents the relative importance of the 20 variables with the highest scores assigned by the preliminary random forest model. It is interesting to note that the lagged value of the dependent variable ranks only 14th in the ranking of most predictive covariates.

3. Panel cross-validation (CV). For each algorithm, we tune hyperparameters via panel CV, involving iterative estimation: we use covariates in 2017 and 2018 to forecast $Y_{i,2019}$, covariates in 2018 and 2019 to forecast $Y_{i,2020}$, covariates in 2019 and 2020 to forecast $Y_{i,2021}$ and covariates in 2020 and 2021 to forecast $Y_{i,2022}$. For each algorithm, we consider the specific hyperparameter(s),[16] and we also treat the number of covariates as an additional hyperparameter (by using the most predictive 10, 20, or 30 covariates selected via the preliminary random forest analysis—step 2 above—or all 160 covariates.).

4. We retrain all the algorithms on the full 2019-2022 sample using the hyperparameters cross-validated in the previous step to forecast $Y_{i,2023}$.

Testing phase

1. We assess the performance for the four algorithms by comparing average forecasted

[15] To tune the model used for forecasting outcomes in 2022 out of sample, we exclude the last training-testing pair from the last row of Figure A.2 to avoid data leakage.

[16] The hyperparameters we select via panel CV are the following: for LASSO, the parameter λ which controls the shrinkage penalty; for the random forest, the parameter m, i.e., the number of features randomly sampled as candidates at each split (for the number of trees to grow, instead, we use the default value of 1,000); for boosting, the shrinkage parameter representing the learning rate, the number of trees to fit and the minimum number of observations in the terminal nodes of the tree; and for LSTMs the number of hidden layers and the number of neurons. Panel CV is used by running different models with several candidate values (or combinations of values, in the case of boosting) for all these parameters.

vs. actual outcomes on the 2023 held-out test data. In particular, we estimate the mean squared forecasting error (MSFE).

2. Final model selection. On the basis of the comparative performance assessment, we pick the best-performing algorithm (random forest) to produce the final risk map.

The general estimation model underlying the forecasting pipeline of Table 2 is the one reported in Equation 1:

$$Y_{i,t} = f(Y_{i,t-1}, Y_{i,t-2}, \boldsymbol{X}_{i,t-1}, \boldsymbol{X}_{i,t-2}) + \varepsilon_{i,t} \tag{1}$$

This forecasting model aims to predict the number of workplace fatalities Y in LLM i and year t, where the outcome is assumed to be a function of highly predictive Y outcome lags and lagged covariates $\boldsymbol{X}$ from the previous two years.[17] We remain agnostic regarding the data-generating process and flexibly allow for possibly arbitrary complexity of the function f to be addressed via a range of ML algorithms characterized by varying degrees of complexity.

3.2.1 Ex-post impact assessment

For the ex-post analysis of the impact of current measures (on-site work inspections and public subsidies for occupational safety) on the number of workplace deaths, we employ the double/debiased ML method developed by Chernozhukov et al. (2018). The method is based on double orthogonalization and leverages the Frisch–Waugh–Lovell theorem to assess the impact of a treatment on the outcome of interest under a partially linear model, where the treatment effect is assumed to be additive, whereas the data-generating process of the untreated potential outcome can be arbitrarily complex. This involves estimating two predictive models, one for the treatment and another for the outcome, to remove bias from confounding factors. The idea is to use ML models to create an orthogonal score for the target parameter, which helps in removing the bias introduced by regularization. In other words, the method uses flexible

[17] Where the set of predictors includes only those selected at the end of step 2 in Table 2.

ML techniques to select the most important confounders and flexibly adjusts for them. Once the treatment and outcome have been debiased and denoised using ML learners, the average treatment effect is estimated via OLS regression of the residualized outcome on the residualized treatment. The main assumption is that of unconfoundedness, i.e., that the treatment is orthogonal to the outcome given the covariates. Furthermore, double/debiased ML relies on cross-fitting to prevent overfitting. We use several different ML learners and then rely on the coefficient estimates of the best-performing learner in terms of the root mean squared error of both the outcome and treatment regressions. Specifically, we employ a set of linear (LASSO, ridge, elastic net) and non-linear (random forest, boosting, artificial neural networks) ML techniques, as well as a simple OLS regression model for comparability.[18] Originally developed for cross-sectional settings, the method is now also routinely applied to panel data settings (e.g., Chernozhukov et al., 2022; Girma & Paton, 2024). Compared with traditional estimators, the key advantage of double ML is that it does not make any a priori assumptions or subjective judgments regarding the functional form of the data-generating process.

We apply this technique for two separate impact evaluation analyses: one on the effect of on-site inspections and the other on the impact of public subsidies for occupational safety. For these analyses, we employ a dataset at the province level, which offers a reasonably high level of resolution and the only administrative level for which inspection and subsidy data are available. Provinces in Italy are broadly comparable to counties in the United States (Barone et al., 2024) and have been employed by many studies interested in the detection of local phenomena (e.g., Le Moglie and Sorrenti, 2022). To overcome the problem represented by the outlier years marked by the COVID-19 pandemic, we run the analysis on a first-differenced dataset where the outcome is the province-level change in the number of workplace deaths per 100,000 inhabitants between 2019 and 2023.

[18] See Hastie et al. (2009) for a detailed description of the ML techniques.

We explore the effects of two policies: on-site inspections and public subsidies to firms. Hence, the treatment is a binary variable taking a value of 1 if there has been an increase in the number of inspections (and, alternatively, in the amount of public subsidies) per 100,000 inhabitants between the years 2018-2019 and 2022-2023 and 0 otherwise, and the confounders are changes in the lagged values of a set of province-level covariates aimed at capturing a variety of demographic, socio-economic, and geographic characteristics of Italian provinces, including all the variables that we later identify as structurally different between high-risk areas and other areas in the descriptive statistics provided at the end of the forecasting analysis (see Table A.4 for their description and A.5 for summary statistics). Considering that we have data for 95 provinces and include 27 predictors, the ratio of predictors to the number of observations approaches a high-dimensional setting, where ML algorithms have an edge over traditional methods, particularly in terms of handling complex interactions and non-linearities and avoiding overfitting.

Since our treatment variables are both highly endogenous and there are multiple confounding factors, we follow Britto et al. (2022) in using first-differencing for causal ML techniques based on the unconfoundedness assumption, ensuring that double orthogonalization is plausible in our context and that time-invariant unobserved heterogeneity is properly accounted for. First-differencing has recently been assessed as the most appropriate method for combining double machine learning with panel data (Clarke and Polselli, 2025). This difference-in-differences framework, which uses only 2019 and 2023 as the start and endpoints for the death data, not only prevents data quality issues for the pandemic years that led us to impute data for those years in the forecasting exercise but also allows us to filter out the impact of the COVID-19 pandemic, as it was a systemic shock common to both treated and untreated areas. For the timing of inspections, we follow Johnson et al. (2023) in that we assume that the impact of inspections can either be contemporaneous or lagged with respect to the year in which the number of deaths is measured. We assume the same for our alternative policy of interest, the amount of public subsidies for occupational safety and health.

Consequently, the partially linear model we employ for this counterfactual analysis is as follows:

$$\Delta Y_{p,2023-2019} = \beta\left(\Delta D_{p,2022/23\,-\,2018/2019}\right) + g\left(\Delta \mathbf{Z}_{p,2022-2018}\right) + \Delta\varepsilon_{p,2023-2019} \quad (2)$$

where on the left-hand side, we have the change in the number of workplace fatalities ΔY in province p between 2023 and 2019. On the right-hand side, we have the change in the treatment variable D, which takes value 1 if number of inspections (and, alternatively, the amount of public subsidies for occupational safety and health) in province p has increased between 2018-19 and 2022-2023 (allowing for both contemporaneous and lagged impacts) and 0 otherwise, plus a vector of covariates composed of changes in the lagged values of a vector of confounders $\mathbf{Z}$, without any *a priori* assumption regarding the function g that maps these covariates to the outcome. The coefficient β captures the impact of an increase in inspections (or in public subsidies) on the change in workplace deaths. We estimate it via a final residual-on-residual regression of the residualized outcome on the residualized treatment after flexibly debiasing and denoising them from confounders with double orthogonalization and preventing overfitting with 5-fold cross-fitting. Our interest in this ex-post analysis is twofold: i) to assess the overall effect of increasing workplace inspections and public subsidies on the risk of death on the job and ii) to connect this counterfactual analysis to the previous risk forecasting analysis to verify whether the impacts are stronger in areas forecasted to be at the highest risk by the independently run ML forecasting analysis. Should this prove to be the case, it would imply that replacing current allocation rules with machine predictions would boost the effectiveness of on-site inspections and/or public subsides by targeting red flag areas that are more responsive to this policy intervention. For this reason, after running the main analysis where we consider 'treated' all areas experiencing an increase in inspections (or the amount of subsidies received) between 2018-2019 and 2022-2023, we conduct an additional heterogeneity analysis where we consider 'treated' only a subsample of areas experiencing such an increase, namely, those that belong to the highest risk decile for 2023 according to the ex-ante forecasts of the best-performing ML technique, and check whether the estimates differ with respect to the effect retrieved in the main analysis. In this way, we effectively simulate a scenario in which a policymaker increases the number of inspections and subsidies in certain high-risk areas.

4. Results

4.1 Main estimates

The last column of Table 3 below reports the performance—in terms of MSFE—of the five selected ML algorithms in forecasting the number of on-the-job deaths in 2023 for each LLM. It also reports the performance for different numbers of covariates, showing that, at least for this application, using only a small subset of the most predictive covariates results in more accurate forecasts. The best performing algorithm is PLS with 9 principal components. This algorithm produces an MSFE of 1.694, which is approximately 5% smaller than those of LASSO and random forest, 10% smaller than the LSTMs and 20% smaller than the stochastic gradient boosting. However, an external benchmark is needed to assess whether this ML performance is truly accurate. To this end, we compare the performance of the ML algorithms with that of the naïve estimator, i.e., a basic method that is widely used in the forecasting literature and used as a forecast for the i^{th} LLM at time t+1 ($Y_{i,t+1}$), which is the last observed value of Y at time t ($Y_{i,t}$) (Hyndman and Athanasopoulos, 2021). In other words, the naïve estimator uses the last available data point as a forecast for the subsequent one; it is intuitive enough that it could be used as a simple risk assessment by a policymaker. The MSFE error of the naïve estimator is 3.031, which is much larger than that of all the ML algorithms and 79% larger than that of the PLS, the best performing technique.

In Supplemental Appendix A, we also test whether the ML algorithms provide accurate forecasts for every population size category (Table A.6) and for an alternative held-out year (Table A.7). These tables demonstrate that the ML algorithms provide reliable forecasts and outperform the naïve estimator for each population size. Additionally, all the models consistently overtake the naïve estimator even for 2022. Notably, the LSTMs outperform all the other models in the smallest municipalities, whereas LASSO is the best-performing algorithm when 2022 is used as the held-out year.

Table 3: Performance by number of predictors for 2023

	10 variables	20 variables	30 variables	ALL (160) variables	Smallest MSFE
LASSO	1.734	1.729	1.763	1.968	1.729
PLS	1.727	1.694	1.855	2.432	1.694
Random Forest	1.810	1.784	1.785	1.839	1.784
Stochastic gradient boosting	2.190	2.741	2.555	2.184	2.184
Neural network	1.893	5.421	2.180	2.449	1.893
Naïve			3.031		

Notes: We report the parameters selected via the panel CV for the best performing case of each ML algorithm. For boosting, 1,500 trees were selected, 2 as the maximum depth of each tree, 10 as the minimum number of observations in terminal nodes and 0.01 for the learning rate. For LASSO, penalty parameters equal to 0.9, 9 as the number of components (PLS), and 5 for the number of variables randomly sampled as candidates at each split (random forest) are used. The neural network comprises two LSTM layers followed by a dense layer. The first LSTM layer has 500 units and feeds into a second LSTM layer with 50 units. A dense layer with a single unit produces the final output, using ReLU activation throughout.

Figure 1 presents two diagnostic tests for the best-performing ML algorithm, comparing the actual number of on-the-job deaths in each LLM in 2023 with the forecasts obtained via PLS. The scatterplot (Panel A) shows that the forecasts are quite accurate, as the data points are closely aligned along the 45-degree diagonal. Notably, the correlation between the observed and forecasted outcomes is extremely high, i.e., 0.929, demonstrating the ability of the selected ML algorithm to accurately forecast the number of on-the-job deaths in each LLM. Panel B of Figure 1 illustrates that the distribution of the forecasting errors is approximately normal and centered around zero, suggesting that there is no bias in the forecasts. We also run the Holden-Peel test (1990) to provide a formal test in which there is no bias in the forecast errors. The p-value of the test is 0.749, suggesting that there is no systematic forecast error. This also implies that in this context, there are no substantial distributional shifts over time or that ML algorithms in this context are capable of capturing them.

Figure 1: Forecasted and actual number of on-the-job deaths in 2023

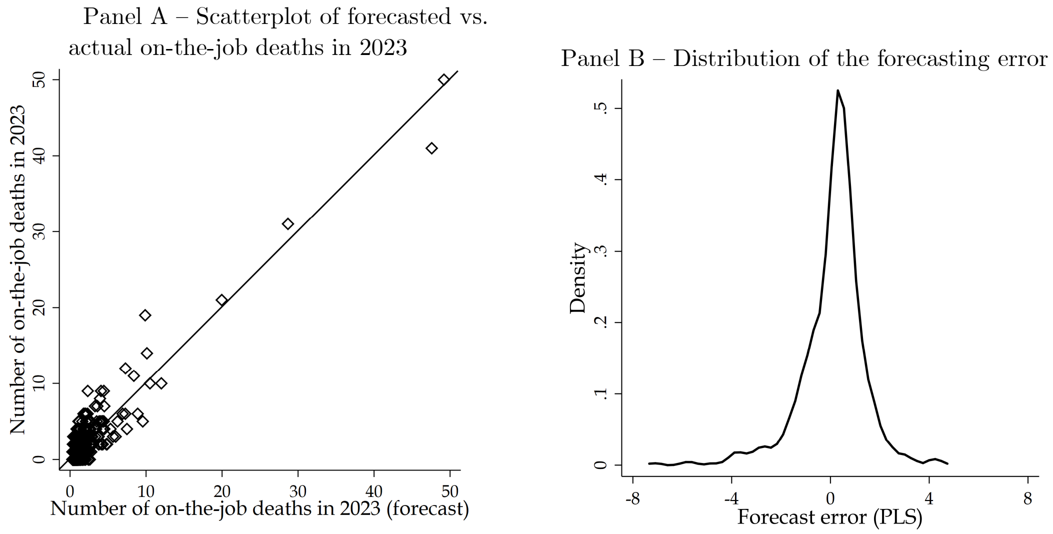

Notes: Data are at the LLM-level and refer to the held-out year 2023.

Since transparency is particularly important when ML tools are intended for use in public policy (Athey, 2017), we exploit Interpretable ML and Explainable AI (Lundberg & Lee, 2017; Lundberg et al., 2020; Molnar, 2020) to create SHapley Additive exPlanations (SHAP) values for the best-performing PLS algorithm. By decomposing the role played by key covariates, we make machine predictions more transparent, communicable, and interpretable. The SHAP values displayed in Figure A.4 in Supplemental Appendix A indicate that areas characterized by a lower number of plants and workers in general and in the construction sector in particular, as well as having smaller populations and a lower incidence of self-employment, tend to experience fewer fatal accidents. These results align with economic intuition, which strengthens the reliability of the models.

To further enhance interpretability and corroborate insights from the

SHAP values, we also present an intuitive 'local surrogate model' (Molnar, 2020) represented by a simple decision tree. In this model, we regress the predictions of the best-performing model (PLS) for 2023 LLM-level analysis on the set of predictors employed.[19] The purpose here is to replace the output of more complex and less interpretable models with an easy-to-understand decision rule for the prediction task. The resulting tree is reported in Figure A.5 and confirms that the number of workers in the construction sector and the total number of plants in a local economy in a given year are key predictors of the future number of on-the-job deaths.

In addition to the interpretability of the models, we further validate the high accuracy of the ML forecasts by mapping, as shown in Figure 2, the forecasts of on-the-job deaths for 2023 (central panel) alongside the observed numbers for the same year (left panel), demonstrating strong adherence between the two. This is confirmed by the forecasting error displayed in the right panel of Figure 2, which also indicates a lack of spatial autocorrelation (Moran's I index of +0.024).

Figure 2: Observed vs. forecasted outcome data for the testing set year (2023)

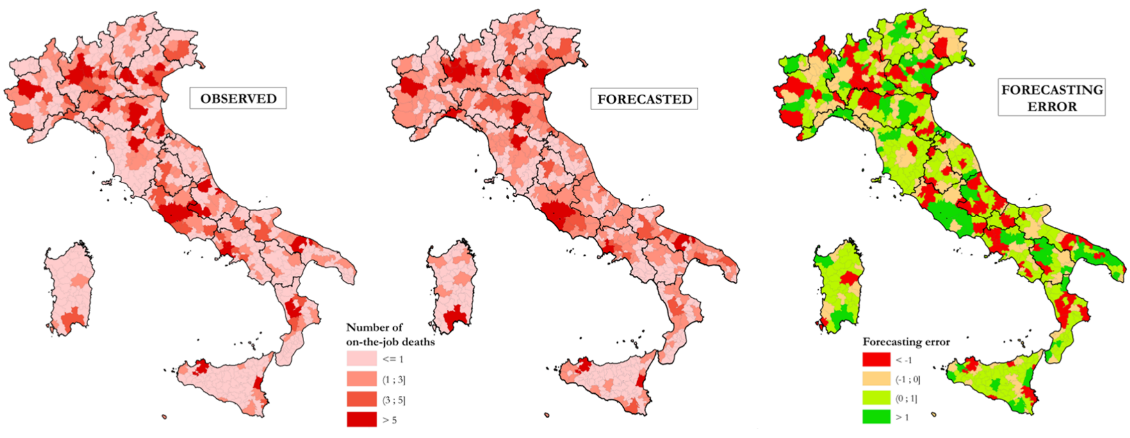

Notes: We have used Moran's I index, which is based on a queen contiguity spatial matrix, to test for the possible presence of spatial autocorrelation in the forecasting error. We obtained a Moran's I index of 0.024, which provides reassurance about the potential presence of spatial heterogeneity we did not account for.

We have also tested whether the forecasting approach based on ML algorithms is effective for alternative outcome measures. In particular, we used

[19] Following Molnar (2020), we set the maximum depth of the three to 3 to ensure interpretability, and the minimum number of observations that must exist in a node to attempt a split equal to 30.

the number of on-the-job deaths in each of the three sectors for which we disaggregated the data and considered only women as alternative dependent variables. The performance of the ML algorithms for these analyses is reported in Table 4. Importantly, for each outcome, every algorithm substantially outperforms the naïve estimator. Overall, this demonstrates the efficacy of a data-driven approach in targeting workplace deaths, even when a more detailed disaggregation of the data is considered. Although these outcomes are generally less frequent and therefore more challenging to forecast, ML algorithms, combined with a rich and informative dataset and a rigorous forecasting pipeline, are able to reliably anticipate the occurrence of workplace deaths. Finally, in Supplemental Appendix B, we provide detailed sensitivity checks on the consequences of not accounting for COVID-19-related workplace deaths.

Table 4: Performance with alternative outcomes

	Agriculture	Industry and services	Public sector	Women
LASSO	0.227	1.358	0.044	0.104
PLS	0.228	1.385	0.043	0.096
Random Forest	0.247	1.409	0.046	0.101
Stochastic gradient boosting	0.226	1.780	0.048	0.125
Neural network	0.237	1.681	0.057	0.113
Naïve	0.426	2.516	0.077	0.190

Notes: The alternative outcome measures include the number of on-the-job deaths in each of the three sectors for which we have disaggregated data and the number of on-the-job deaths among women. In every analysis, the held-out year is 2023. Each cell reports the MSFE for the best performing version of each ML algorithm.

The forecasting analysis outlined above yields a distribution of risk across LLMs. Naturally, the regulator is often interested not only in gaining access to accurate forecasts of the areas at highest risk but also, and perhaps more importantly, in detailed insights regarding the main territorial drivers and predictors underpinning these forecasts. Moreover, understanding the key characteristics of the areas most at risk—specifically how they differ structurally across several critical dimensions compared to less risky areas—is paramount. Therefore, the forecasting analysis must necessarily be supplemented by additional investigations addressing the following question: What distinguishes high-risk areas in terms of socio-economic, demographic, and geographic characteristics?

To answer this question, we exploit the estimates for workplace deaths to flag fatality hotspots and derive a risk classification for each LLM, defining 'high risk' local economies as those with a predicted risk (i.e., with a forecasted number of workplace deaths) above the 90th percentile of the distribution, and 'low risk' as with all other LLMs. Based on this classification, Table 5 presents a descriptive analysis of the systematic differences between high-risk and low-risk LLMs.

High-risk LLMs have different economic structures, with a significantly greater share of agricultural and construction workers and a lower share of manufacturing employment. The latter result is somewhat unexpected, on the basis of our previous discussion that the highest share of on-the-job deaths occurs in the industry and services sectors (cf. Table 1). High-risk areas also exhibit greater entrepreneurial density, with approximately 44 more plants per 1,000 inhabitants. While employment rates appear comparable (differing by only 1.66 percentage points and not statistically significant at conventional levels), high-risk LLMs show significantly lower unemployment rates by approximately 1.89 percentage points. Local economies identified as hotspots also present a significantly younger population and a lower share of foreigners. Notably, we find no significant differences in terms of income per capita, geographical distribution (South and Islands), or income inequality as measured by the Gini index, suggesting that our risk classification is not merely capturing well-known North–South disparities or income-related differences.

This evidence complements insights from the SHAP values and corroborates the notion that, while typically unexplored, territorial data retain significant predictive potential for the phenomena under scrutiny. Nonetheless, this exploratory analysis of high-risk areas is notably constrained by the lack of data at a higher resolution, which hinders the ability to identify drivers and characterize areas with granular detail. This limitation stems from the data itself rather than the methodology employed, and it could be addressed by replicating the forecasting and descriptive analyses using more detailed and comprehensive datasets.

Table 5: Differences in LLM characteristics by risk classification

Variables	High-risk LLMs	Other LLMs	Difference	p-value
% workers in agriculture	0.210	0.104	0.107 (0.014)	0.000
% workers in construction	0.119	0.108	0.010 (0.005)	0.042
% workers in manufacture	0.122	0.209	-0.087 (0.017)	0.000
Plants per 1,000 inhabitants	159.059	114.769	44.290 (3.070)	0.000
Income per capita (€)	14,315.490	14,569.692	-254.202 (516.972)	0.623
Unemployment rate	6.966	8.855	-1.889 (0.620)	0.002
Employment rate	45.391	43.731	1.660 (1.078)	0.124
% population over 65	0.274	0.252	0.022 (0.004)	0.000
% foreigners	0.053	0.073	-0.019 (0.005)	0.000
Share of LLMs in Southern regions	0.475	0.457	0.018 (0.067)	0.787
Gini index	0.394	0.394	0.000 (0.004)	0.955

Notes: This table reports the mean values and the differences in LLM characteristics between high-risk LLMs and other LLMs. The last column reports p-values from two-sided t-tests of equality of means.

4.2 ML risk map vs. current policy rules

After demonstrating that our ML forecasting approach yields accurate estimates of workplace deaths and insights into the key characteristics of areas at highest risk, we now turn to an analysis of public policy targeting and resource allocation criteria. We exploit our estimates to create a granular risk map for 2023 that reports the forecasted number of deaths per 100,000 inhabitants[20], and compare it with the current allocation and geographic distribution of work inspections and public subsidies for workplace safety. In particular, the questions we address with

[20] We have chosen to use the number of inhabitants as the denominator instead of the number of workers due to the geographically heterogeneous presence of informal work in Italy. Anyway, we also replicate the same analysis per 100,000 workers (see below).

this comparison are as follows: i) is public effort concentrated in the areas where it is most needed? ii) can the ML risk map be employed to refine the targeting of these policies? As the data concerning inspections and public subsidies are available only at the provincial level, we have re-run the main forecasting analysis using provincial-level data, as aggregating data from LLMs to provinces without imputation is not feasible, given that many LLMs span territories across two or more provinces. The performance of the ML forecasting pipeline for provincial data aligns with that reported for LLM data and is reported in Table A.8 in Supplemental Appendix A.

Figure 3 below provides the answers to these questions. Specifically, the map compares the risk map computed as the forecasted on-the-job death rates per 100,000 inhabitants with the number of inspections per 100,000 inhabitants and the amount of public subsidies per 100,000 inhabitants for the year 2023.[21] Visually, the risk map indicates greater risk in the North, while inspections appear more concentrated in the South, and public funds are predominantly allocated to the central regions. This misalignment between public efforts and the risk map is confirmed by the weak correlation of the risk map with the number of inspections per 100,000 inhabitants (+0.2271) and the amount of public subsidies per 100,000 inhabitants (+0.1708). The weak correlation with the amount of public subsidies is not surprising: indeed, their assignment mechanism does not take into account the past number of workplace injuries and deaths; rather, it is based on other requirements, such as company size and the type of proposed project. If we normalize with respect to the number of workers instead of inhabitants (see Figure A.6 in Supplemental Appendix A), the correlation strengthens (the correlation between the risk map and the number of inspections per 100,000 workers is +0.4919, whereas the correlation between the risk map and the amount of public

[21] We emphasize that the risk map is intended to be available to policymakers *in advance*; therefore, data on observed deaths contemporaneous to the policy data cannot be used to build it. While other methods, such as the naïve forecaster, could be employed, our ML methods have been shown to produce much more accurate and reliable risk estimates.

subsidies per 100,000 workers is +0.2778). We note that normalizing by the number of workers is not a sensible choice given the heterogeneous prevalence of the shadow economy and informal work in the country, hinted at by the mismatch between the number of deaths and injuries in the comparison reported in Figure A.1, which is also the main reason for our focus on fatalities, as discussed. Even setting aside this caveat, a significant misalignment persists between the risk map and public resources for workplace safety. In addition, these findings are independent of the ML algorithm selected, as demonstrated by Tables A.9 and A.10 in Supplemental Appendix A. These tables show that the risk maps produced by other ML algorithms strongly align with the one presented here, and the correlations are almost identical, irrespective of the algorithm used. In addition, Table A.10 complements Table A.7, which reports forecasting performances for the alternative year 2022 by showing the correlations between risk maps and public resources for 2022 as well. While the values are slightly higher than those for 2023, they still suggest that there is substantial room for refining the targeting of these measures and that the core findings are not driven by the choice of a particularly 'lucky' year in terms of forecasting performance or an 'unlucky' year for resource allocation.

By comparing the maps, we can conclude that the targeting of current policies is misaligned with the actual risk in the territory. Can this misaligned targeting also affect the effectiveness of policies aimed at increasing occupational health and safety? In the next subsection, we answer this question by examining the results of the ex-post analysis of the increase in the number of inspections and the amount of public subsidies per 100,000 inhabitants between 2019 and 2023 on the change in the number of workplace deaths per 100,000 inhabitants.

Figure 3: Comparison of the risk map and public efforts at the provincial level in 2023, normalized per 100,000 inhabitants

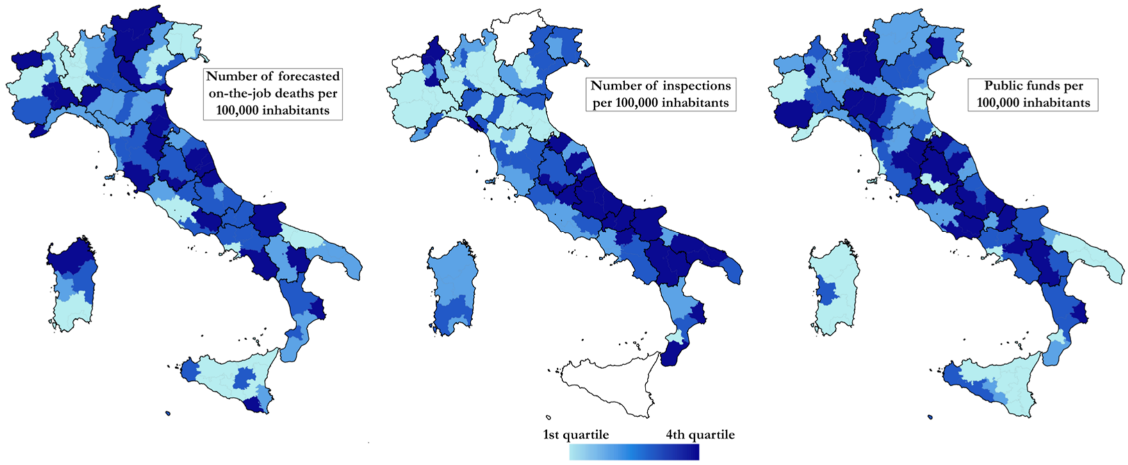

Notes: We lack data on 12 provinces for the inspections, as the INL covers the whole Italian territory with the exception of two three special-status regions: Aosta Valley (1 province), Sicily (9 provinces) and Trentino-South Tyrol (2 provinces). Each variable is segmented into four classes (quartiles). The forecasted on-the-job death rates per 100,000 inhabitants for 2023 exhibit a positive but weak correlation with the number of inspections per 100,000 inhabitants (+0.2271) and the amount of public subsidies per 100,000 inhabitants (+0.1708).

4.3 Effects of public policies on workplace safety

Finally, we move from the ex-ante targeting analysis to the ex-post policy evaluation. We start by providing counterfactual estimates for the effect of on-site inspections, and then summarize the results on the impact of public subsidies, which are presented in Supplemental Appendix C. For both analyses, we include, among the set of confounders, all the variables identified as structurally different between high-risk areas and other areas in the descriptive statistics provided at the end of the forecasting analysis (cf. Table 5). Tables 6 and 7 below report the results of the ex-post analysis described in subsection 3.2.1 for the effect of on-site inspections. Table 6 presents the main analysis, where we consider all provinces that experienced an increase in on-site inspections between 2018–2019 and 2022–2023 as treated. Table 7 connects the counterfactual analysis with the previous risk forecasting analysis by focusing on a heterogeneity analysis, which examines the impact on provinces that experienced an increase in inspections and belong to the highest risk decile according to the best-performing algorithm (LASSO) from the independent ML forecasting analysis run at the province level.

Table 6: Impact of an increase in inspections on workplace fatalities

	ATE estimate	Standard error	Lower CI	Upper CI
OLS	-0.238	0.279	-0.785	0.309
Lasso	-0.037	0.254	-0.535	0.461
Ridge	-0.057	0.267	-0.578	0.464
Elastic net	-0.039	0.254	-0.536	0.459
Random forest	-0.046	0.280	-0.596	0.503
Boosting	-0.027	0.266	-0.548	0.494
Neural networks	-0.120	0.275	-0.658	0.419

Notes: The table shows the province-level impact of an increase in the number of inspections between 2018-2019 and 2022-2023 per 100,000 inhabitants on the change in the number of workplace fatalities per 100,000 inhabitants between 2019 and 2023. All provinces experiencing such an increase are considered treated. Treatment and outcome regressions for debiasing include the change in the lagged values (2018-2022) of the set of confounders described in Table A.4 in Supplemental Appendix A. The number of observations is 95. Except for LASSO, Ridge, and Elastic Net, for which cross-validation is performed, all other models use default values for the tuning hyperparameters. Cross-fitting with 5 folds was employed for all models except OLS.

Table 7: Impact of an increase in inspections on workplace fatalities in high-risk areas

	ATE estimate	Standard error	Lower CI	Upper CI
OLS	-1.874***	0.725	-3.295	-0.453
Lasso	-1.733**	0.719	-3.142	-0.325
Ridge	-1.541**	0.709	-2.931	-0.151
Elastic net	-1.733**	0.719	-3.142	-0.325
Random forest	-2.267***	0.741	-3.719	-0.814
Boosting	-1.64**	0.728	-3.064	-0.211
Neural networks	-2.010***	0.684	-3.350	-0.669

Notes: The table shows the province-level impact of an increase in the number of inspections between 2018-2019 and 2022-2023 per 100,000 inhabitants on the change in the number of workplace fatalities per 100,000 inhabitants between 2019 and 2023 in areas classified as 'high-risk', i.e., *in the highest decile of the probabilistic risk distribution for 2023, according to the best-performing algorithm (LASSO) of the independent ML forecasting analysis run at the province level. Only these provinces are considered treated. Treatment and outcome regressions for debiasing include the change in the lagged values (2018-2022) of the set of confounders described in Table A.4 in Supplemental Appendix A. The number of observations is 95. Except for LASSO, Ridge, and Elastic Net, for which cross-validation is performed, all other models use default values for the tuning hyperparameters. Cross-fitting with 5 folds was employed for all models except OLS. Stars denote statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.*

By comparing the treatment effect estimates of the two tables, all the different algorithms employed to orthogonalize the treatment and the outcome support the same qualitative story: the overall impact of the increase in inspections registered in the five years of our sample on workplace deaths is negative but negligible and statistically insignificant. In contrast, this impact is negative, sizable, and statistically significant for high-risk areas, ranging from a reduction in the number of deaths per 100,000 inhabitants from -1.464 for boosting to -2.267 for random forest.

Tables C.1 and C.2 in Supplemental Appendix C report the corresponding estimates for the effect of the increase in public subsidies for occupational safety and health. Qualitatively, the results are consistent and suggest that, in terms of impact magnitude, an increase in public subsidies generally exerts a greater negative effect in areas flagged as high-risk by machine forecasts. From a statistical significance perspective, the results are more nuanced. Although the reduction is small, it is significant for two models (OLS and neural networks) in the baseline estimates. In contrast, although the impact is much greater in high-risk areas experiencing an increase in subsidies, it is significant only for OLS, Ridge, and random forest. This is mainly because the coefficients are more noisily estimated than they are in the inspection analysis, and the standard errors are larger, with confidence intervals covering zero for most estimates.

In conclusion, while changes in these policy measures have, on average, been mostly ineffective in reducing fatalities, they can produce significant reductions in workplace deaths when implemented mainly in areas that are more vulnerable to the phenomenon. This treatment effect heterogeneity, depending on the ML-predicted risk levels, underscores the place-sensitive nature of workplace safety policies. The implications of this analysis are twofold: i) focusing on average effects when there is imperfect targeting can be misleading and lead to the dismissal of effective policies; ii) the apparent lack of broader impacts is most likely due to inadequate targeting and the misallocation of policy interventions (cf. Figure 3) in the territory. Our results suggest that shifting focus and concentrating policy efforts in red-flag areas identified by machine forecasts could boost the real-world impacts of these interventions. In more detail, a back-of-the-envelope calculation based on the most conservative estimate for the inspection analysis (boosting) suggests that counterfactually increasing the number of on-site inspections by 100 in *all* areas flagged as high-risk by our ex-ante machine forecasts, but not in other areas not classified as high-risk, might prevent a total of 86 workplace fatalities per year, or about 10% of the average annual on-the-job deaths registered during the years included in the analysis.[22] This reduction, in turn, would substantially help Italian policymakers reach the EU *Vision Zero* target as well as their stated goal of reducing the total number of workplace deaths (INAIL, 2022). In economic terms, directing this increase in policy only to provinces in the highest risk decile distribution (or another cutoff chosen by the regulator) would mean that the cost of these additional inspections in high-risk areas would be more than counterbalanced by reducing ineffective inspections in less risky areas. By focusing resources on areas flagged as high-risk, the overall cost-effectiveness of the policy would be enhanced.[23]

Importantly, the above considerations pertain solely to one of the potential objectives of the two policies—among other interventions in occupational safety and health—that we examine: the reduction of on-the-job fatalities. On-site inspections and public subsidies might also be implemented to achieve other goals that we do not consider. This relates to the issue of the so-called 'omitted payoffs' (Kleinberg et al., 2018) that the policymaker might be trying to achieve, which could explain the limited overlap of the allocation of inspections and public subsidies with our risk map, as well as the overall ineffectiveness in reducing

[22] This calculation is inevitably the result of an approximate estimate and should therefore be treated with caution. First, we convert to absolute terms the average change in inspections per 100,000 inhabitants in high-risk areas, resulting in a total of 36.48 more inspections in these areas over the period considered. Next, we also convert the most conservative estimate (boosting), which is expressed as the reduction in on-the-job deaths per 100,000 inhabitants yielding a reduction of 3.94 deaths associated with the increase in inspections. We then multiply this number by the absolute increase in inspections to obtain the total number of deaths (31.52) associated with 36.48 more inspections, assuming such an increase occurs in all high-risk areas. Finally, we derive from this estimate the total number of deaths avoided by conducting 100 additional inspections, instead of 36.48 more, in each high-risk area.

[23] However, since this conclusion is drawn from aggregate data, the ecological fallacy cautions that it might not hold true for individual plants within high-risk areas.

workplace deaths. Therefore, we stress that our judgment of the policies pertains solely to workplace deaths. Nevertheless, among the possible goals of these policies, the one we study is arguably one of the most important goals pursued by the regulator.

Overall, although this ex-post analysis comes with caveats associated with the estimation of causal effects due to the strong unconfoundedness assumption, we believe that, when coupled with the previous ML risk comparison, it makes a strong case for revising the targeting mechanisms of these policy measures.

5. Discussion and conclusion

Despite government agencies spending billions each year on workplace safety inspections (Johnson et al., 2023), fatal workplace accidents continue to pose a significant policy challenge across both developing and developed countries (ILO, 2023). Thus, devising new solutions, revising current allocation criteria, and framing alternative strategies are necessary to guide and complement existing measures. Using Italy as a case study, we demonstrated that coupling AI-powered risk assessment with a place-based approach explicitly centered on granular workplace death data can accurately forecast areas most at risk. This can also inform and target policy interventions on occupational health and safety in the territory, such as the local deployment of work inspections, which are currently misaligned with the actual distribution of risk and are not as effective as they could be if they specifically targeted areas at highest risk. The implication of this mismatch is that our AI-powered territorial targeting approach can reduce the incidence of this chronic issue while lowering the costs of policy implementation.Importantly, our conclusions rest on key caveats concerning the partial scope of our policy data (cf. Section 3.1.3) and the acknowledgment that the Italian policy framework for occupational safety and health encompasses more than the aspects analyzed in this study. Nonetheless, we believe our findings represent a significant initial step towards fostering a more data-driven, quantitative, and ultimately effective approach to policy discussion and implementation on these critical issues.In Italy, the topic is currently under intense political discussion in the country due to the Italian Legislative Decree 103/2024, which was effective as of July 12, 2024 and introduced significant reforms aimed at simplifying administrative controls on businesses. The decree redefines the framework for inspections concerning labor, safety, and health and streamlines how public administrations regulate economic activities. It reduces bureaucratic burdens on businesses while preserving essential oversight of safety and compliance. Within such a regulatory framework, it is possible to envision that the ML tools we developed in this paper could be employed to guide the implementation of the new rules at the local level and avoid misalignment in targeting, which, as we have shown, has characterized previous policy efforts, damaging their effectiveness. We are aware that translating predictions into policies comes with non-trivial issues concerning the transparency and interpretability of the measures implemented by the regulator (Athey, 2017). However, we believe that the latest advancements in the fields of Interpretable Machine Learning and Explainable AI (Lundberg & Lee, 2017; Lundberg et al., 2020; Molnar, 2020) can effectively address most of these concerns by opening the black box of machine predictions and making their fundamentals accessible and understandable to non-experts and policymakers.

Replacing current rules with those based on machine learning forecasts in hotspot areas might elicit behavioral changes in less inspected or less subsidized areas not at high risk, triggered by the lack of deterrence effects and unobservable in our historical data covering periods with old criteria. We deem this possibility, related to the so-called Lucas critique (Lucas, 1976), unrealistic. This has already been shown to be very unlikely for non-fatal injuries at the plant level (Johnson et al., 2023). Within our place-based approach, this hypothesis seems even more remote, as the distribution of risk is more likely driven by economic factors, such as the heterogeneous structural characteristics of local economies and the prevalence of the shadow economy, rather than by variations in behavioral patterns or inspection frequencies across different areas. In any case, the Lucas critique applies only to a budget-constrained scenario where the regulator can allocate only a fixed number of inspections or amounts of subsidies, and cost trade-offs exist. It does not pertain to scenarios where ML rules are employed to allocate additional resources to high-risk areas only without reducing pre-existing public efforts in less risky areas. Furthermore, our approach is intended to be based on dynamic learning and subject to continuous monitoring and updating so that any changing pattern in the risk distribution across areas can be promptly incorporated into the ML forecasting models and the targeting rules can be updated accordingly, thus preventing any significant risk spillover to less inspected areas.

The rationale for the place-based approach we advocate and the reason for the accuracy of the ML forecasts are rooted in the systematic territorial patterns associated with the phenomenon, which have thus far been largely ignored by the literature despite the growing sectorial specialization and clustering of local economies. The focus on deaths rather than injuries is also an important element that enhances the reliability of the forecasts we produced, as it minimizes the complex issues involving predictive tasks with imperfect labels (Cannings et al., 2020). These issues arise when injury data affected by underreporting are used, a particularly thorny problem in contexts such as Italy, with a high and clustered prevalence of the shadow economy.

Regarding potential future developments, several additional avenues warrant exploration. For instance, access to microdata on workplace fatality occurrences at the firm and plant level would enable extending the forecasting and targeting approach proposed in this paper to achieve substantially higher resolution.[24] Additionally, overcoming the limitations of the inspection data at our disposal—such as the inability to disaggregate the number of inspections by NACE-2 digit sectors or across other dimensions and categories (e.g., detailed data on the number of inspectors at the local level, as well as the capacity and characteristics of territorial inspection units)—would enable the conduct of heterogeneity analyses by both sector and region. Such analyses could assess both the adequacy of targeting rules and the effectiveness of these policies, which we were unable to explore in this work. Finally, having access to data on other policies that indirectly address occupational safety and health in Italy would allow evaluating targeting mechanisms and causal effects of these additional policies.. In turn, this would offer a more comprehensive understanding of the state of the Italian policy framework on this critical issue and enhance the external validity of our findings.

The general estimation model underlying the forecasting pipeline of Table 2 is the one reported in Equation 1:

$$Y_{i,t} = f(Y_{i,t-1}, Y_{i,t-2}, \boldsymbol{X}_{i,t-1}, \boldsymbol{X}_{i,t-2}) + \varepsilon_{i,t} \tag{1}$$

This forecasting model aims to predict the number of workplace fatalities Y in LLM i and year t, where the outcome is assumed to be a function of highly predictive Y outcome lags and lagged covariates $\boldsymbol{X}$ from the previous two years.[25] We remain agnostic regarding the data-generating process and flexibly allow for possibly arbitrary complexity of the function f to be addressed via a range of ML algorithms characterized by varying degrees of complexity.

In conclusion, our findings can be leveraged to optimize public efforts and more effectively address an issue that is a primary target of occupational safety

[24] A promising extension at the micro level would involve examining an alternative outcome—specifically, the number of fatalities normalized by the number of hours worked—which would provide a more precise definition of the outcome variable.

[25] Where the set of predictors includes only those selected at the end of step 2 in Table 2.

and health regulation and yet remains challenging to address. More generally, our place-based approach could be applied to other important issues, such as gender violence, suicides, and gambling addiction, for which the policy payoffs of combining algorithmic tools with increasingly rich and granular data to fine-tune targeting and pinpoint drivers of these complex phenomena might also be considerable.

References

Andini M, Ciani E, de Blasio D, D'Ignazio A, Salvestrini V (2018) Targeting with machine learning: An application to a tax rebate program in Italy. *Journal of Economic Behavior & Organization*, 156: 86–102.

Antonelli MA, Castaldo A, Forti M, Marrocco A, Salustri A (2024) Workplace accidents, economic determinants and underreporting: An empirical analysis in Italy. *International Journal of Manpower*, 45(8): 1555−1572.

Antulov-Fantulin N, Lagravinese R, Resce G (2021) Predicting bankruptcy of local government: A machine learning approach. *Journal of Economic Behavior & Organization*, 183: 681–699.

Arkhangelsky D, Imbens G (2024) Causal models for longitudinal and panel data: A survey. *Econometrics Journal*, *27*(3), C1-C61.

Ash E, Galletta S, Giommoni T (2024) A machine learning approach to analyze and support anti-corruption policy. *American Economic Journal: Economic Policy*, forthcoming.

Athey S (2017) Beyond prediction: Using big data for policy problems. *Science*, *355*(6324): 483-485.

Athey S, Keleher N, Spiess J (2025) Machine learning who to nudge: causal vs predictive targeting in a field experiment on student financial aid renewal. *Journal of Econometrics*, 105945.

Athey S, Wager S (2019) Estimating treatment effects with causal forests: An application. *Observational studies*, *5*(2), 37-51.

Autor DH (2015) Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, *29*(3), 3-30.

Barile B, Forti M, Marrocco A, Castaldo A (2024) Causal impact evaluation of occupational safety policies on firms' default using machine learning uplift modelling. *Scientific Reports*, 14: 10380.

Barone G, Schivardi F, Sette E (2024) Interlocking directorates and competition in banking. *Journal of Finance*, forthcoming.

Battaglini M, Guiso L, Lacava C, Miller DL, Patacchini E (2024) Refining public policies with machine learning: The case of tax auditing. *Journal of Econometrics*, available online 23 September 2024, 105847.

Blesse S, Diegmann A (2022) The place-based effects of police stations on crime: Evidence from station closures. *Journal of Public Economics*, *207*, 104605.

Boone J, Van Ours JC (2006) Are recessions good for workplace safety? *Journal of Health Economics*, 25(6): 1069–1093.

Britto DG, Pinotti P, Sampaio B (2022) The effect of job loss and unemployment insurance on crime in Brazil. *Econometrica*, *90*(4): 1393-1423.

Brynjolfsson E, Mitchell T (2017) What can machine learning do? Workforce implications. *Science*, *358*(6370), 1530-1534.

Brynjolfsson E, Mitchell T, Rock D (2018) What can machines learn and what does it mean for occupations and the economy? *AEA papers and proceedings*, 108, 43-47.

Campo G, Cegolon L, De Merich D, Fedeli U, Pellicci M, Heymann WC, Pavanello S, Guglielmi A, Mastrangelo G (2020) The Italian national surveillance system for occupational injuries: conceptual framework and fatal outcomes, 2002-2016. *International Journal of Environmental Research and Public Health* 17(20): 7631.

Cannings TI, Fan Y, Samworth RJ (2020) Classification with imperfect training labels. *Biometrika*, *107*(2), 311-330.

Carrieri V, Lagravinese R, Resce G (2021) Predicting vaccine hesitancy from area-level indicators: A machine learning approach. *Health Economics*, 30(12): 3248–3256.

Castaldo A, Ragazzi E, Sella L (2023) È possibile incentivare la sicurezza sui luoghi di lavoro? Concezione, contesto, e implementazione dei Bandi ISI Inail. Giappichelli.

Cerqua A, Di Stefano R, Letta M, Miccoli S (2021) Local mortality estimates during the COVID-19 pandemic in Italy. *Journal of Population Economics*, 34(4): 1189–1217.

Cerqua A, Letta M, Menchetti F (2024) Causal inference and policy evaluation without a control group. *arXiv preprint arXiv:2312.05858*.

Cerqua A, Letta M, Pinto G (2024) On the mis(use) of machine learning with panel data. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5014594

Cerqua, A, Letta, M (2022) Local inequalities of the COVID-19 crisis. *Regional Science and Urban Economics*, *92*, 103752.

Cerqua A, Pellegrini G (2022) Decomposing the employment effects of investment subsidies. *Journal of Urban Economics*, 128:103408.

Charles KK, Johnson MS, Stephens Jr M, Lee DQ (2022) Demand conditions and worker safety: Evidence from price shocks in mining. *Journal of Labor Economics*, 40(1): 47–94.

Chernozhukov V, Chetverikov D, Demirer M, Duflo E, Hansen C, Newey W, Robins J (2018) Double/debiased machine learning for treatment and structural parameters. *Econometrics Journal*, 21(1): C1-C68.

Chernozhukov V, Newey WK, Singh R (2022) Automatic debiased machine learning of causal and structural effects. *Econometrica*, *90*(3), 967-1027.

Christensen P, Francisco P, Myers E, Shao H, Souza M (2024) Energy efficiency can deliver for climate policy: Evidence from machine learning-based targeting. *Journal of Public Economics*, 234: 105098.

Clarke P, Polselli A (2025). Double machine learning for static panel models with fixed effects. *Econometrics Journal*, forthcoming.

Colombo S, Golzio LE, Bianchi G (2019) The evolution of health-, safety- and environment-related competencies in Italy: From HSE technicians, to HSE professionals and, eventually, to HSE managers. *Safety Science*, 118: 724–739.

Corinth K, Coyne D, Feldman N, Johnson C (2025). The Targeting of Place-Based Policies: The New Markets Tax Credit Versus Opportunity Zones. *National Bureau of Economic Research (NBER) Working Paper,* No.33414.

Criscuolo C, Martin R, Overman HG, Van Reenen J (2019) Some causal effects of an industrial policy. *American Economic Review*, 109(1): 48–85.

DeLeire T, Levy H (2024) Worker sorting and the risk of death on the job. *Journal of Labor Economics*, 22(4): 925–953.

Dillender M, McInerney M (2020) The role of Mexican immigration to the United States in improved workplace safety for natives from 1980 to 2015. *Journal of Health Economics*, 70: 102280.

Drescher K, Janzen B (2025) When weather wounds workers: The impact of temperature on workplace accidents. *Journal of Public Economics*, *241*, 105258.

EASHW (European Agency for Safety and Health at Work) (2023) The Role of Vision Zero and Related Occupational Safety & Health Strategy, Interventions, and Tools in Reducing EU Work-Related Fatalities, Accidents and Ill Health. Available at https://oshwiki.osha.europa.eu/en/themes/role-vision-zero-and-related-occupational-safety-health-strategies-interventions-and-tools

Faggio G (2019) Relocation of public sector workers: Evaluating a place-based policy. *Journal of Urban Economics*, 111: 53–75.

Felten EW, Raj M., Seamans R (2018) A method to link advances in artificial intelligence to occupational abilities. *AEA Papers and Proceedings*, Vol. 108, 54-57.

Filomena M, Picchio M (2024) Unsafe temperatures, unsafe jobs: The impact of weather conditions on work-related injuries. *Journal of Economic Behavior & Organization*, 224: 851–875.

Freedman M (2015) Place-based programs and the geographic dispersion of employment. *Regional Science and Urban Economics*, 53: 1–19.

Gallé J, Overbeck D, Riedel N, Seidel T (2024) Place-based policies, structural change and female labor: Evidence from India's Special Economic Zones. *Journal of Public Economics*, *240*, 105259.

Girma S, Paton,D (2024) Using double-debiased machine learning to estimate the impact of Covid-19 vaccination on mortality and staff absences in elderly care homes. *European Economic Review*, *170*, 104882.

Glaeser EL, Hillis A, Kominers SD, Luca M (2016) Crowdsourcing city government: Using tournaments to improve inspection accuracy. *American Economic Review*, *106*(5), 114-118.

Hastie T, Tibshirani R, Friedman JH (2009) *The elements of statistical learning: data mining, inference, and prediction.* New York, Springer.

Hyndman RJ, Athanasopoulos G (2021) Forecasting: principles and practice. Third Edition. Melbourne, Australia: OTexts.

Holden K, Peel D (1990) On testing for unbiasedness and efficiency of forecasts. Manchester School, 58: 120–127.

International Labor Organization (ILO) (2023) A call for safer and healthier working environments. Available at: https://www.ilo.org/sites/default/files/wcmsp5/groups/public/%40ed_protect/%40protrav/%40safework/documents/publication/wcms_903140.pdf

INAIL (Istituto Nazionale per l'Assicurazione contro gli Infortuni sul Lavoro) (2022) Relazione Annuale. Available here: https://www.inail.it/portale/it/inail-comunica/pubblicazioni/rapporti-e-relazioni-inail/rapporti-e-relazioni-inail-dettaglio.2023.09.relazione-annuale-2022.html

Jäger S, Heining J (2022) How substitutable are workers? evidence from worker deaths. NBER Working Paper, No. 30629.

Johnson MS (2020) Regulation by shaming: Deterrence effects of publicizing violations of workplace safety and health laws. *American Economic Review*, 110(6): 1866–1904.

Johnson MS, Levine DI, Toffel MW (2023) Improving regulatory effectiveness through better targeting: Evidence from OSHA. *American Economic Journal: Applied Economics*, 15(4): 30–67.

Johnson MS, Schwab D, Koval P (2024) Legal protection against retaliatory firing improves workplace safety. *The Review of Economics and Statistics*, 106(5): 1236–1253.

Kleinberg J, Ludwig J, Mullainathan S, Obermeyer Z (2015) Prediction policy problems. *American Economic Review*, 105(5): 491–495.

Kleinberg J, Lakkaraju H, Leskovec J, Ludwig J, Mullainathan S (2018) Human decisions and machine predictions. *Quarterly Journal of Economics*, *133*(1): 237-293.

Knutson V, Aleshin-Guendel S. Karlinsky A, Msemburi W, Wakefield J (2023) Estimating global and country-specific excess mortality during the COVID-19 pandemic. *The Annals of Applied Statistics*, *17*(2), 1353-1374..

Le Moglie M, Sorrenti G (2022) Revealing "mafia inc."? Financial crisis, organized crime, and the birth of new enterprises. *Review of Economics and Statistics*, *104*(1):142-156.

Lee, JM, Taylor, LO (2019) Randomized safety inspections and risk exposure on the job: Quasi-experimental estimates of the value of a statistical life. *American Economic Journal: Economic Policy*, *11*(4), 350-374.

Liu Q, Lu R, Sun ST, Zhang M (2024) Unintended workplace safety consequences of minimum wages. *Journal of Public Economics*, *239*, 105247.

Lu Y, Wang J, Zhu L (2019) Place-based policies, creation, and agglomeration economies: Evidence from China's economic zone program. *American Economic Journal: Economic Policy*, *11*(3), 325-360.

Lucas Jr RE (1976) Econometric policy evaluation: A critique. In *Carnegie-Rochester conference series on public policy* (Vol. 1, pp. 19-46). North-Holland.

Lundberg S, Lee SI (2017) A unified approach to interpreting model predictions. A*dvances in Neural Information Processing Systems, 30.*

Lundberg SM, Erion G, Chen H, DeGrave A, Prutkin JM, Nair B, ... Lee SI (2020) From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, *2*(1): 56-67.

Mayer T, Mayneris F, Py L (2017) The impact of urban enterprise zones on establishment location decisions and labor market outcomes: evidence from France. *Journal of Economic Geography*, 17(4): 709–752.

McManus TC, Schaur G (2016) The effects of import competition on worker health. *Journal of International Economics*, 102: 160–172.

Melchior C, Zanini RR, Guerra RR, Rockenbach DA (2021) Forecasting Brazilian mortality rates due to occupational accidents using autoregressive moving average approaches. *International Journal of Forecasting*, 37(2): 825–837.

Molnar C (2020) *Interpretable Machine Learning*. Lulu. com.

Msemburi W, Karlinsky A, Knutson V, Aleshin-Guendel S, Chatterji S, Wakefield J (2023) The WHO estimates of excess mortality associated with the COVID-19 pandemic. *Nature*, *613*(7942), 130-137.

Mullainathan S, Obermeyer Z (2017) Does machine learning automate moral hazard and error?. *American Economic Review*, *107*(5), 476-480.

Organization for Economic Cooperation and Development, International Labor Organization (OECD/ILO) (2019) Tackling vulnerability in the informal economy. Development Centre Studies, OECD Publishing, Paris, https://doi.org/10.1787/939b7bcd-en.

Petropoulos F, Apiletti D, Assimakopoulos V, Babai MZ, Barrow DK., Taieb, SB, ... & Ziel F (2022). Forecasting: theory and practice. *International Journal of Forecasting*, *38*(3), 705-871.

Picchio M, Van Ours JC (2017) Temporary jobs and the severity of workplace accidents. *Journal of Safety Research*, *61*, 41-51.

Schweiger H, Stepanov A, Zacchia P (2022) The long-run effects of R&D place-based policies: evidence from Russian science cities. *American Economic Journal: Economic Policy*, *14*(3), 322-351.

Supplemental Appendix A – Descriptive statistics and robustness analyses

Table A.1: Yearly number of on-the-job deaths in Italy for 2020 and 2021 – imputed and actual values

	2020 actual	2020 imputed	2021 actual	2021 imputed
Industry and Services	1,298	802	979	766
Agriculture	120	131	127	141
Public Sector	74	15	49	12
Total	1,492	948	1,155	919

Notes: We have not imputed the number of "in itinere" deaths for 2020 and 2021, as the reduction is quite homogeneous across the country and therefore does not pose a significant threat to the territorial analysis.

Table A.2: Variable details for the targeting analysis

Dependent variable		Source
Number of on-the-job workplace deaths		INAIL

Predictors	Lags	Source
Economic classification dummies (e.g., highly specialized urban LLMs, multi-specialized urban LLMs, touristic LLMs, agri-food LLMs)	Time-invariant (17 dummies)	Istat
Geographical dummies (North–East, North–West, Centre and South)	Time-invariant (4 dummies)	Istat
Population dummies: ≤ 10,000 inhabitants; (10,000; 50,000]; (50,000; 100,000]; (100,000; 500,000]; >500,000 inhabitants	Time-invariant (5 dummies)	Istat
Share of population located in municipalities considered as peripheral or ultra-peripheral according to the SNAI classification	Time-invariant	Istat
LLM considered as rural	Time-invariant (1 dummy)	Istat
LLM with at least one industrial district	Time-invariant (1 dummy)	Istat
Average price per square meter (house and villa separately)	1 lag + 1 lag of the difference	Osservatorio del Mercato Immobiliare – Agenzia delle Entrate
Average age of the mayors	1 lag + 1 lag of the difference	Ministry of the Interior
% of mayors with a university degree	1 lag + 1 lag of the difference	Ministry of the Interior
% of female mayors	1 lag + 1 lag of the difference	Ministry of the Interior
Number of voters at the European elections	1 lag + 1 lag of the difference	Ministry of the Interior
Turnout at the European elections	1 lag + 1 lag of the difference	Ministry of the Interior
Gender gap in turnout at the European elections	1 lag + 1 lag of the difference	Ministry of the Interior
Income per capita	1 lag + 1 lag of the difference	Ministry of Economy and Finance (MEF)
Overall declared income	1 lag + 1 lag of the difference	MEF
Declared income - employed	1 lag + 1 lag of the difference	MEF
Declared income - entrepreneurs	1 lag + 1 lag of the difference	MEF
Declared income - self-employed	1 lag + 1 lag of the difference	MEF
Declared income - retirees	1 lag + 1 lag of the difference	MEF
Declared income - equities	1 lag + 1 lag of the difference	MEF
Declared income – lands and buildings	1 lag + 1 lag of the difference	MEF
Municipal personal income surcharge revenues	1 lag + 1 lag of the difference	MEF
Regional personal income surcharge revenues	1 lag + 1 lag of the difference	MEF
Gini index	1 lag + 1 lag of the difference	MEF
Unemployment rate	1 lag + 1 lag of the difference	Istat
Employment rate	1 lag + 1 lag of the difference	Istat
Population	1 lag + 1 lag of the difference	Istat
Population - foreigners	1 lag + 1 lag of the difference	Istat
Population - 65+ years old	1 lag + 1 lag of the difference	Istat
Number of newborns	1 lag + 1 lag of the difference	Istat
Number of newborns per 1,000 inhabitants	1 lag + 1 lag of the difference	Istat
Number of plants	1 lag + 1 lag of the difference	Infocamere

Number of plants in agriculture	1 lag + 1 lag of the difference	Infocamere
Number of plants in manufacturing	1 lag + 1 lag of the difference	Infocamere
Number of plants in the construction sector	1 lag + 1 lag of the difference	Infocamere
Number of plants per 1,000 inhabitants	1 lag + 1 lag of the difference	Infocamere
Number of employees	1 lag + 1 lag of the difference	Infocamere
Number of employees in agriculture	1 lag + 1 lag of the difference	Infocamere
Number of employees in manufacturing	1 lag + 1 lag of the difference	Infocamere
Number of employees in the construction sector	1 lag + 1 lag of the difference	Infocamere
Number of employees per 1,000 inhabitants	1 lag + 1 lag of the difference	Infocamere
Average number of employees per plant	1 lag + 1 lag of the difference	Infocamere
Share of workers in agriculture	1 lag + 1 lag of the difference	Infocamere
Share of workers in manufacturing	1 lag + 1 lag of the difference	Infocamere
Share of workers in the construction sector	1 lag + 1 lag of the difference	Infocamere
Number of on-the-job workplace deaths	1 lag + 1 lag of the difference	INAIL
Number of on-the-job workplace deaths - agriculture	1 lag + 1 lag of the difference	INAIL
Number of on-the-job workplace deaths – industry and services	1 lag + 1 lag of the difference	INAIL
Number of on-the-job workplace deaths – public sector	1 lag + 1 lag of the difference	INAIL
Number of on-the-job workplace deaths – women	1 lag + 1 lag of the difference	INAIL
Number of in-itinere workplace deaths	1 lag + 1 lag of the difference	INAIL
Number of in-itinere workplace deaths - agriculture	1 lag + 1 lag of the difference	INAIL
Number of in-itinere workplace deaths – industry and services	1 lag + 1 lag of the difference	INAIL
Number of in-itinere workplace deaths – public sector	1 lag + 1 lag of the difference	INAIL
Number of in-itinere workplace deaths – women	1 lag + 1 lag of the difference	INAIL

Notes: All variables have been collected at the LLM level. All time-varying variables are available for all years between 2017 and 2023.

Table A.3: Descriptive statistics for the targeting analysis

Variable	Mean	Std Dev
Rural LLM (%)	82.30	38.18
LLM with an industrial district (%)	23.11	42.16
Population	97,432.78	268,652.60
Number of in-itinere workplace deaths	0.52	1.50
Number of in-itinere workplace deaths - agriculture	0.01	0.10
Number of in-itinere workplace deaths – public sector	0.02	0.16
Number of in-itinere workplace deaths – industry and serv.	0.46	1.40
Number of in-itinere workplace deaths - women	0.08	0.33
Number of on-the-job workplace deaths	1.49	3.44
Number of on-the-job workplace deaths - agriculture	0.23	0.60
Number of on-the-job workplace deaths – public sector	0.02	0.16
Number of on-the-job workplace deaths – industry and serv.	1.25	3.25
Number of on-the-job workplace deaths - women	0.10	0.38
Average price per square meter – house (€)	1,064.98	628.64
Average price per square meter – villa (€)	1,265.59	723.59
Average age of the mayors	52.74	4.59
Female mayors (%)	12.67	13.67
Mayors with a university degree (%)	49.66	23.01
Turnout at the European elections (%)	54.79	13.27
Gender gap in turnout at the European elections (pps)	4.35	2.29
Number of employees	31,049.20	104,533.70
Number of employees per 1,000 inhabitants	286.55	89.65
Number of plants	10,407.96	28,696.79
Number of plants per 1,000 inhabitants	116.89	25.20
Workers in agriculture (%)	11.84	10.95
Workers in manufacturing (%)	20.25	12.66
Workers in the construction sector (%)	10.31	3.47
Population with 65 or more years (%)	24.65	3.30
Foreigners (%)	6.95	3.69
Number of newborns	687.32	1991.36
Number of newborns per 1,000 inhabitants	6.73	1.30
Income per capita (€)	13,191.67	3,598.02
Employment rate (%)	42.88	8.21
Unemployment rate (%)	10.69	5.76
N	610	
T	7	
N·T	4,270	

Table A.4: Variable details for the ex-post analysis

Dependent variable	**Source**
Number of on-the-job workplace deaths	INAIL
Covariates	**Source**
Number of on-site work inspections	INL
Public subsidies for occupational safety and health	INAIL
Average price per square meter (house and villa separately)	Osservatorio del Mercato Immobiliare – Agenzia delle Entrate
Average age of the mayors	Ministry of the Interior
% of mayors with a university degree	Ministry of the Interior
% of female mayors	Ministry of the Interior
Turnout at the European elections	Ministry of the Interior
Gender gap in turnout at the European elections	Ministry of the Interior
Income per capita	Ministry of Economy and Finance
Unemployment rate	Istat
Population	Istat
% of foreigners	Istat
% of 65+ years old	Istat
% of males	Istat
Number of newborns per 1,000 inhabitants	Istat
Number of deaths per 1,000 inhabitants	Istat
Number of employees per 1,000 inhabitants	Infocamere
Number of plants per 1,000 inhabitants	Infocamere
Average number of employees per plant	Infocamere
Share of workers in agriculture	Infocamere
Share of workers in manufacturing	Infocamere
Share of workers in the construction sector	Infocamere
The Quality of Life Index and its six thematic pillars: i) wealth and consumption; ii) business and employment; iii) environment and services; iv) demography, society, and health; v) justice and security; vi) culture and leisure time	Il Sole 24 Ore

Notes: All variables have been collected at the provincial level. Each year, the Quality of Life Index gauges well-being across Italian provinces on the basis of a comprehensive analysis of various socio-economic indicators. See this link for more details: https://lab24.ilsole24ore.com/qualita-della-vita/ (in Italian).

Table A.5: Descriptive statistics for the ex-post analysis

Variable	Mean	Std Dev
Population	562,296	631,991
Number of on-the-job workplace deaths	8.63	8.72
Number of on-site work inspections	203.97	166.39
Public subsidies for occupational safety and health (€)	2,345,228	2,285,465
Average price per square meter – house (€)	1,044	360
Average price per square meter – villa (€)	1,241	445
Average age of the mayors	52.71	2.32
Mayors with a university degree (%)	47.58	11.05
Female mayors (%)	15.49	8.31
Turnout at the European elections (%)	58.96	9.57
Gender gap in turnout at the European elections (pps)	3.81	1.70
Number of employees per 1,000 inhabitants	312.10	64.01
Number of plants per 1,000 inhabitants	109.66	13.34
Average number of employees per plant	2.87	0.62
Workers in agriculture (%)	6.81	5.16
Workers in manufacturing (%)	22.50	9.29
Workers in the construction sector (%)	9.30	1.76
Population with 65 or more years (%)	24.58	2.35
Male population (%)	48.84	0.47
Foreigners (%)	8.35	3.39
Number of newborns per 1,000 inhabitants	6.60	0.81
Number of deaths per 1,000 inhabitants	11.86	1.72
Income per capita (€)	15,097	3,197
Unemployment rate	9.16	3.88
Quality of Life Index - Overall	495.11	52.56
Quality of Life Index - Wealth and consumption	489.39	103.86
Quality of Life Index - Business and employment	479.19	69.59
Quality of Life Index - Environment and services	515.57	92.66
Quality of Life Index - Demography, society, and health	594.88	95.69
Quality of Life Index - Justice and security	511.99	211.33
Quality of Life Index - Culture and leisure time	379.62	92.34
N	95	

Notes: All variables have been analyzed at the provincial level. We consider the value of each variable in the years used in the empirical analysis, i.e., *2018, 2019, 2022 and 2023.*

Table A.6: Performance by population size

	Small LLMs (<= 50,000 inh.)	**Medium-sized LLMs (50,000; 200,000] inh.**	**Large LLMs (>200,000 inh.)**	**All LLMs**
LASSO	0.546	2.053	8.367	1.741
PLS	0.510	2.030	8.148	1.694
Random Forest	0.532	2.144	8.706	1.794
Stochastic gradient boosting	0.536	2.227	11.871	2.091
Neural network	0.496	2.099	10.320	1.893
Naïve	1.015	3.881	12.824	3.031
N	340	219	51	610

Notes: Each cell reports the MSFE for the best performing version of each ML algorithm.

Table A.7: Performance by number of predictors for 2022

	10 variables	20 variables	30 variables	ALL (160) variables	Smallest MSFE
LASSO	1.918	1.724	1.768	1.986	1.715
PLS	2.026	1.939	1.936	2.019	1.936
Random Forest	2.273	2.053	2.054	2.133	2.031
Stochastic gradient boosting	4.845	4.789	4.476	3.104	3.104
Neural network	2.268	1.806	1.810	3.343	1.806
Naïve			2.561		

Notes: Each cell reports the MSFE for the best performing version of each ML algorithm.

Table A.8: Performance by number of predictors for 2023 – provincial data

	10 variables	20 variables	30 variables	ALL (138) variables	Smallest MSFE
LASSO	11.409	12.577	13.071	18.941	11.409
PLS	11.830	13.326	13.797	15.529	11.830
Random Forest	12.637	12.523	12.670	12.644	12.523
Stochastic gradient boosting	13.549	13.818	12.519	11.657	11.657
Neural network	14.267	15.508	14.217	17.743	14.217
Naïve			18.794		

Notes: We report the parameters selected via the panel CV for the best performing case of each ML algorithm. For boosting, 1,500 trees were selected, 3 as the maximum depth of each tree, 10 as the minimum number of observations in terminal nodes and 0.01 for the learning rate. For LASSO, a penalty parameter equal to 0.9, 7 as the number of components (PLS), and 21 for the number of variables randomly sampled as candidates at each split (random forest) are used. In this analysis, we did not use the predictors only available at the LLM-level (e.g., the economic classification dummies). The neural network comprises two LSTM layers followed by a dense layer. The first LSTM layer has 500 units and feeds into a second LSTM layer with 50 units. A dense layer with a single unit produces the final output, using ReLU activation throughout.

Table A.9: Correlations between estimates produced by the five ML algorithms and provincial data

	LASSO	PLS	Random Forest	Stochastic gradient boosting	Neural network
LASSO	1				
PLS	0.9913	1			
Random Forest	0.9889	0.9753	1		
Stochastic gradient boosting	0.9682	0.9562	0.9683	1	
Neural network	0.9652	0.9441	0.9706	0.9476	1

Notes: Each cell reports the MSFE for the best performing version of each ML algorithm.

Table A.10: Correlations between ML estimates and public resources

Year 2023

	The number of inspections	**The amount of public subsidies**
	Normalized per 100,000 inhabitants	
Risk map (LASSO)	0.2271	0.1708
Risk map (PLS)	0.2860	0.1424
Risk map (Random Forest)	0.2981	0.0242
Risk map (Stochastic gradient boosting)	0.2182	0.0510
Risk map (Neural network)	0.3056	0.1760
	Normalized per 100,000 workers	
Risk map (LASSO)	0.4919	0.2778
Risk map (PLS)	0.5533	0.2568
Risk map (Random Forest)	0.5133	0.1359
Risk map (Stochastic gradient boosting)	0.4715	0.1627
Risk map (Neural network)	0.5276	0.2169

Year 2022

	The number of inspections	**The amount of public subsidies**
	Normalized per 100,000 inhabitants	
Risk map (LASSO)	0.4132	0.3188
Risk map (PLS)	0.3819	0.3470
Risk map (Random Forest)	0.3719	0.0626
Risk map (Stochastic gradient boosting)	0.3615	0.1299
Risk map (Neural network)	0.3750	0.1447
	Normalized per 100,000 workers	
Risk map (LASSO)	0.6323	0.3832
Risk map (PLS)	0.6177	0.3961
Risk map (Random Forest)	0.6016	0.1789
Risk map (Stochastic gradient boosting)	0.5874	0.2589
Risk map (Neural network)	0.5981	0.2780

Notes: Each cell reports the MSFE for the best performing version of each ML algorithm.

Figure A.1: Number of on-the-job deaths and injuries in 2023, normalized per 100,000 inhabitants and workers

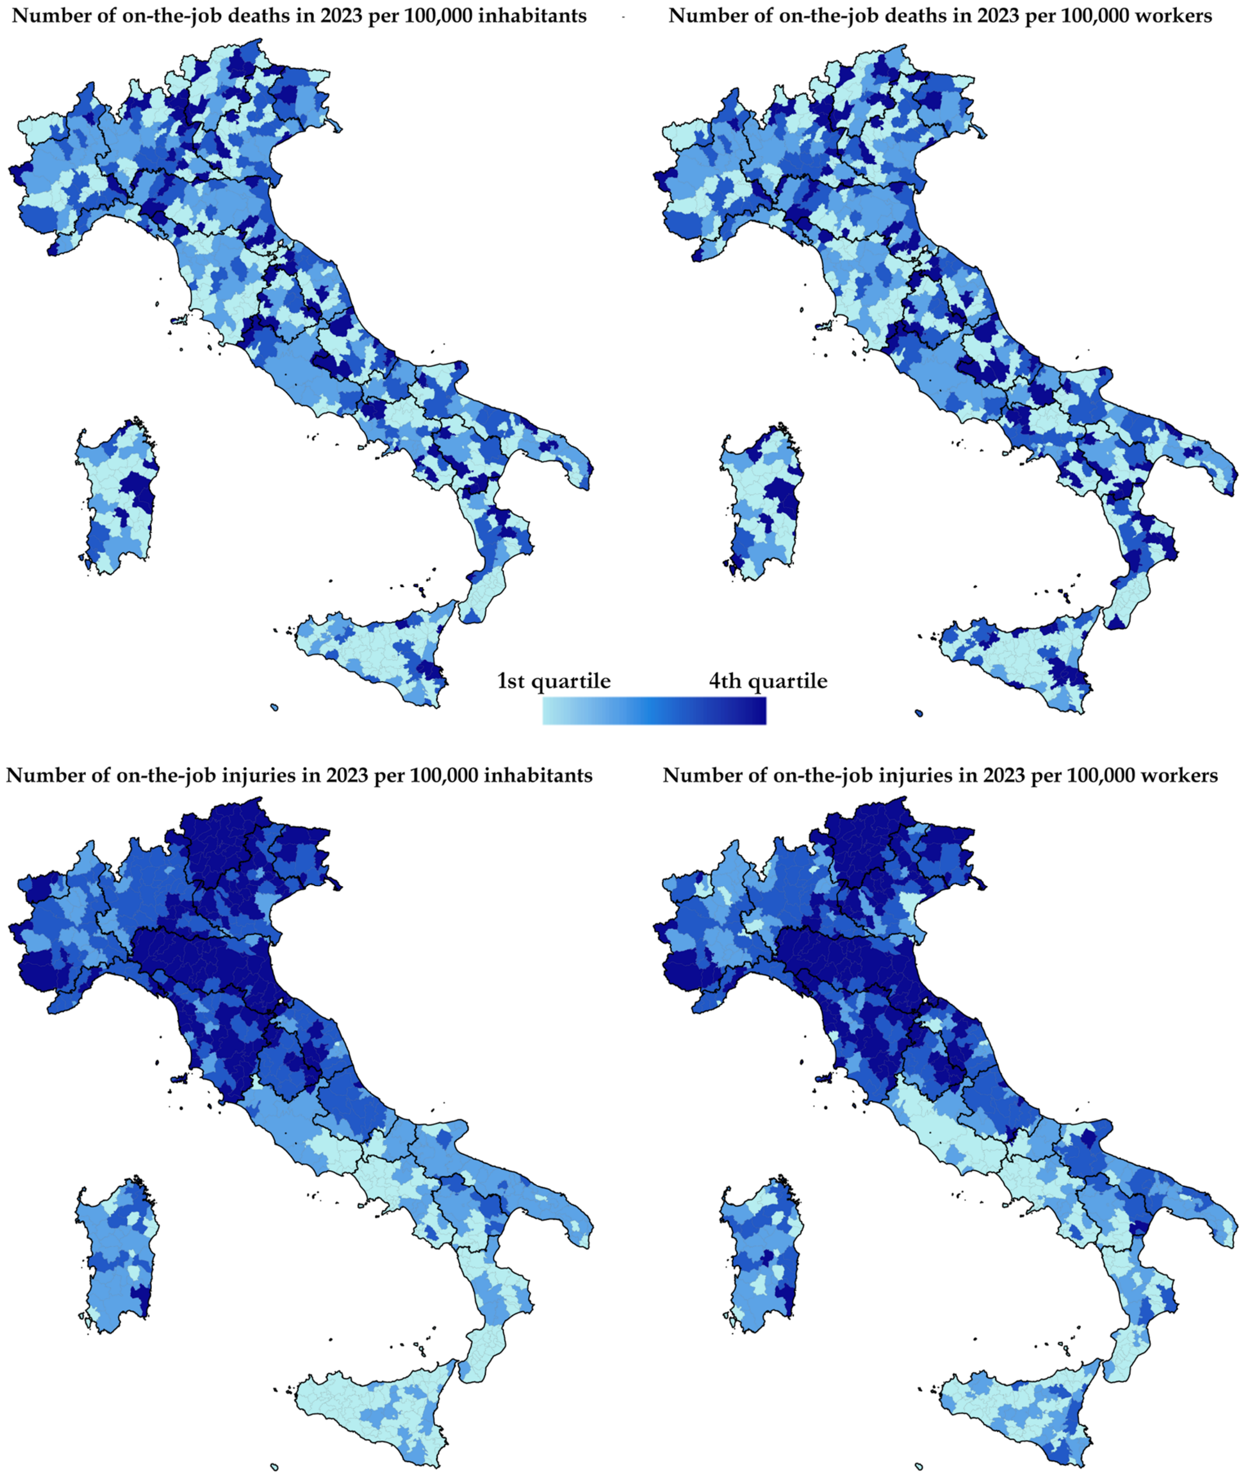

Notes: Data are at the LLM-level and refer to the year 2023.

Figure A.2: Panel cross-validation

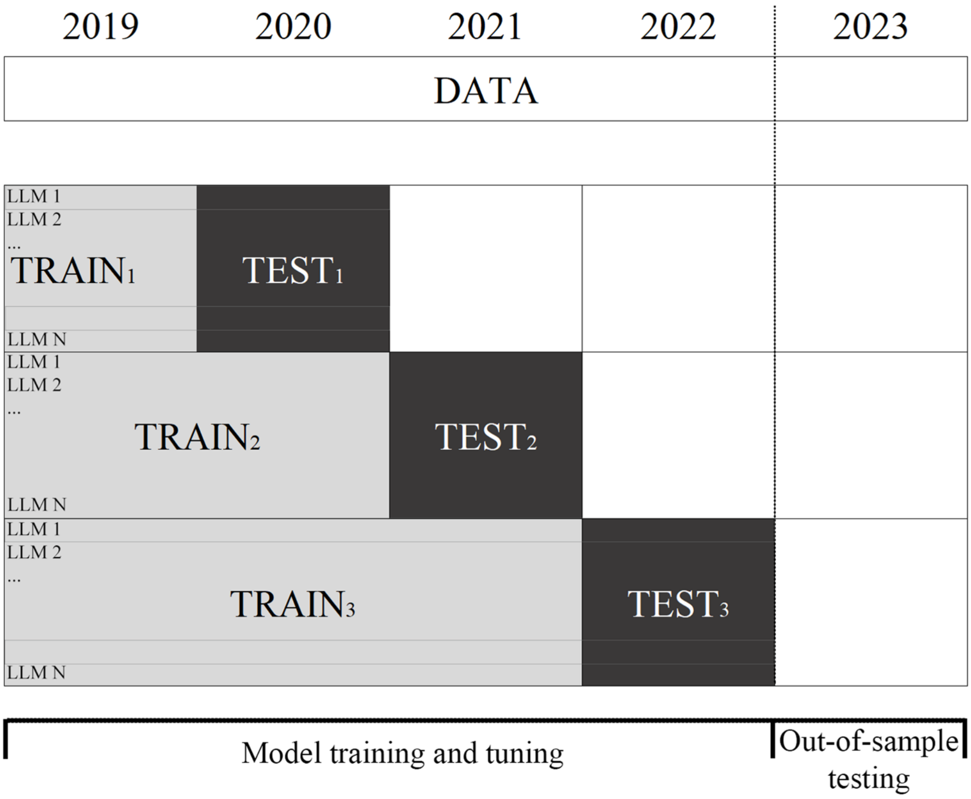

Figure A.3: Variable importance ranking – Preliminary random forest

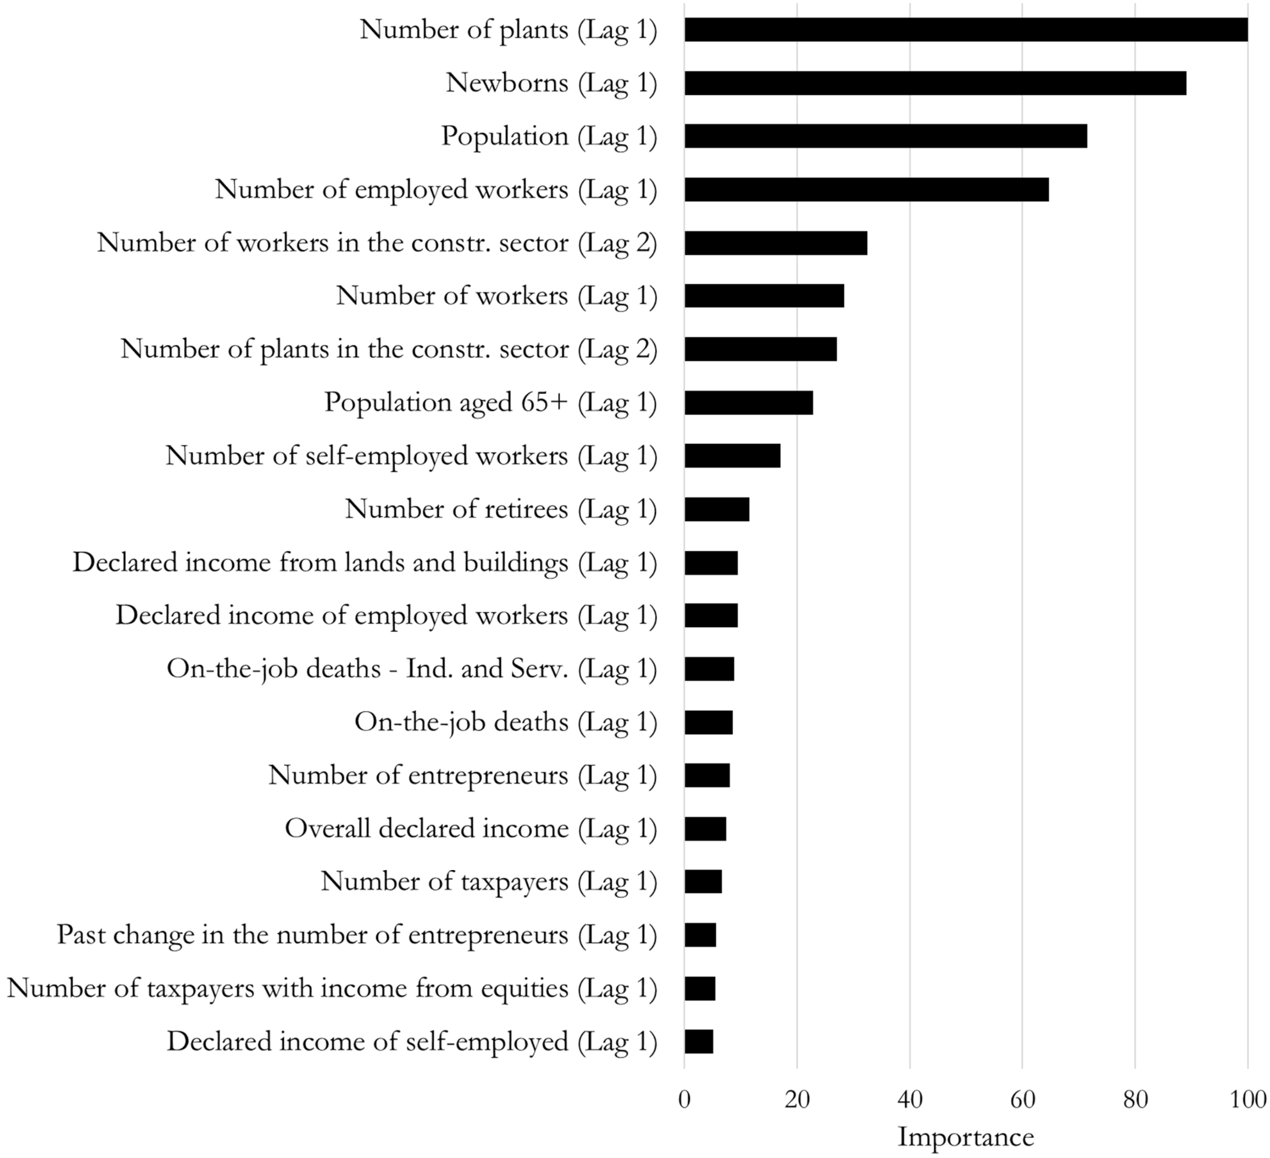

Notes: The figure shows the 20 most predictive covariates for the dependent variable "number of on-the-job deaths" on the basis of the preliminary random forest analysis conducted following the approach proposed by Athey and Wager (2019).

Figure A.4: SHAP values of the best performing model for forecasting on-the-job deaths in 2023

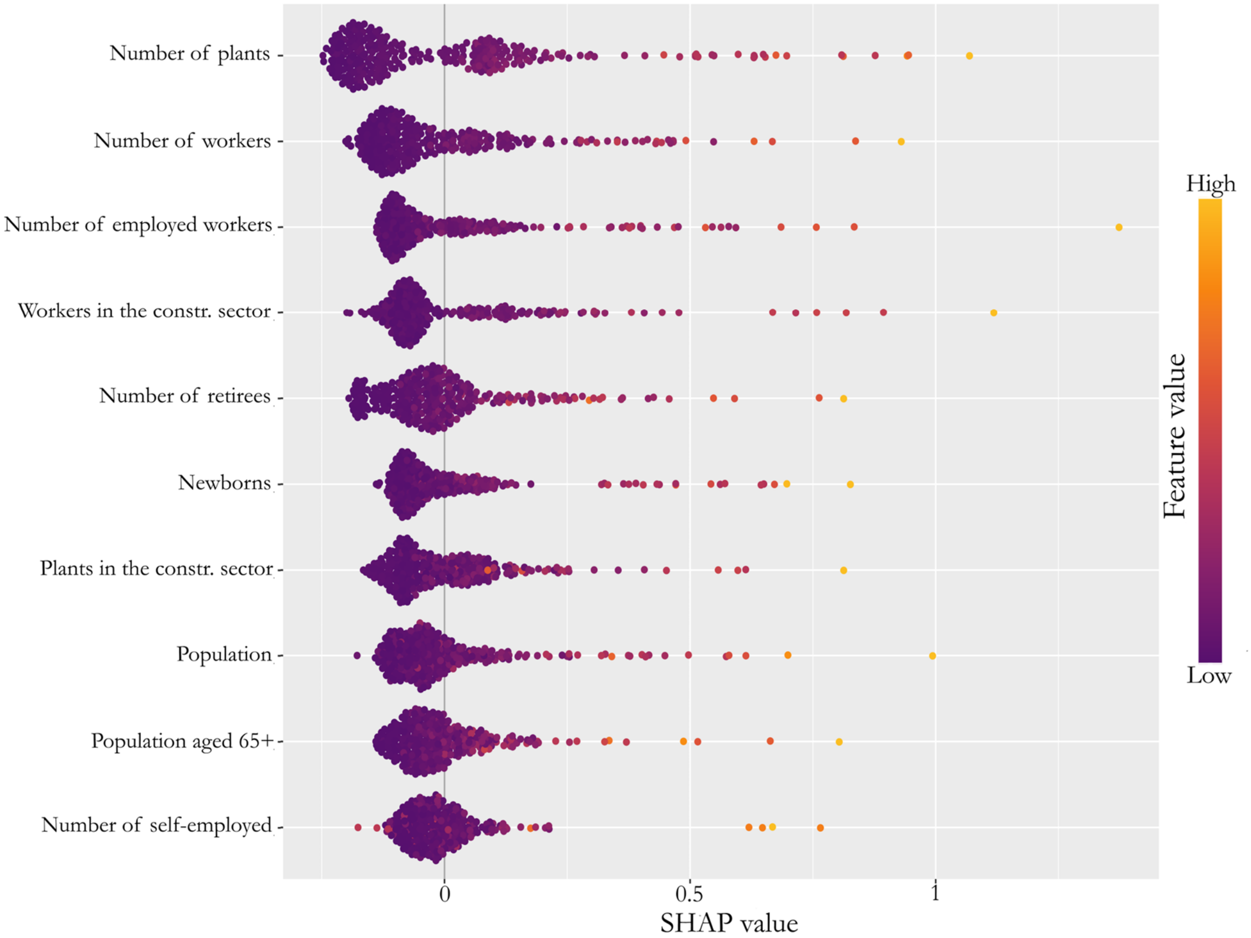

Notes: The best performing model is the PLS model with 20 variables and 9 components. This figure reports the SHAP values for the 10 most predictive variables based on the preliminary random forest analysis.

Figure A.5: Local surrogate model based on 2023 LLM-level forecasts with PLS

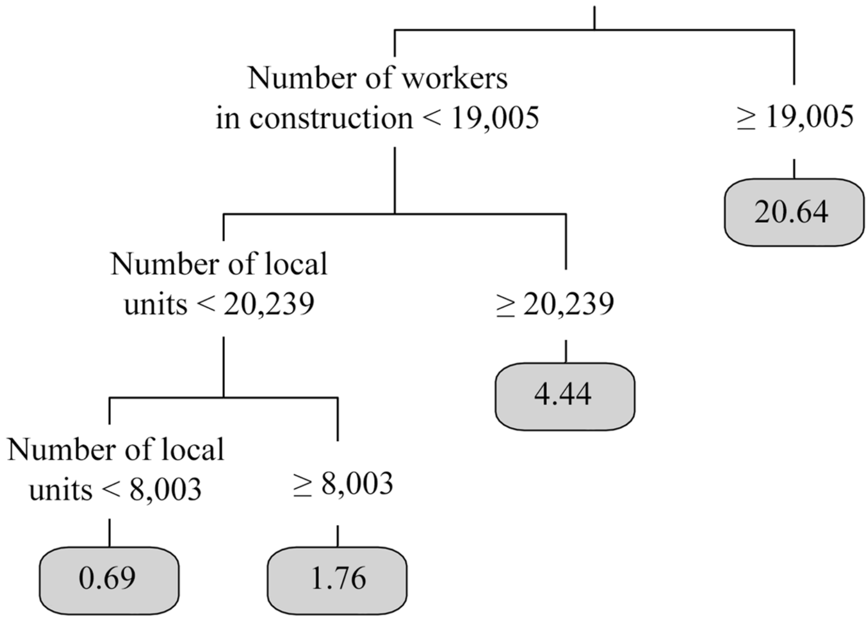

Notes: The numbers in the terminal nodes refer to the predicted annual number of on-the-job deaths in 2023 for all LLMs falling within that node. The predictions are made by the best-performing PLS model.

Figure A.6: Comparison of the risk map and public efforts at the provincial level in 2023, normalized per 100,000 workers

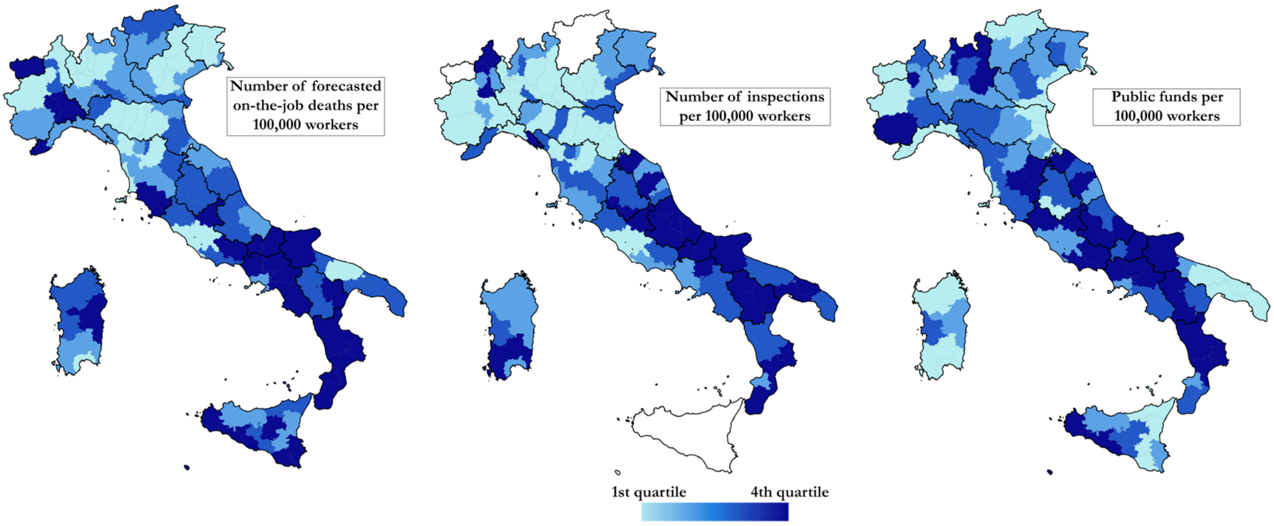

Notes: We lack data on 12 provinces for the inspections, as the INL covers the whole Italian territory with the exception of two three special-status regions: Aosta Valley (1 province), Sicily (9 provinces) and Trentino-South Tyrol (2 provinces). Each variable is segmented into four classes (quartiles). The forecasted on-the-job death rates per 100,000 workers for 2023 exhibit a positive correlation with the number of inspections per 100,000 workers (+0.4919) and a weaker correlation with the amount of public subsidies per 100,000 workers (+0.2778).

Supplemental Appendix B – Testing estimate sensitivity to imputed workplace deaths in 2020 and 2021

In this Appendix, we test the sensitivity of the estimates with respect to the imputation of the number of workplace deaths in 2020 and 2021 to account for the effects of COVID-19 during those years, following the WHO approach (Knutson et al., 2023; Msemburi et al., 2023). INAIL only provides the number of pandemic-related deaths at the national level (576 in 2020, 276 in 2021, 8 in 2022, and 0 in 2023), necessitating an alternative strategy for imputation at the sub-national level. To address this, we leveraged the strong correlation (more than +0.70) between the "excess deaths" data due to COVID-19 provided by Cerqua et al. (2021) and the "excess" workplace deaths calculated as the difference from the reported number of on-the-job deaths in the post-COVID-19 years and those reported in 2019 at the LLM level, which is our unit of analysis (see below). The imputation process also considered sector-specific characteristics and gender composition. As in the WHO approach, these corrections were made only for the years impacted by the pandemic, namely, 2020 and 2021. Following this process, the average number of deaths per sector and gender remained relatively stable between 2017 and 2023. Table A.1 in Supplemental Appendix A presents the number of deaths by year and sector before and after the imputation. Importantly, in the empirical analysis, the years in which we validated our methodological approach were 2022 and 2023 (see the Methodology section), years in which no imputation occurred. Moreover, the imputation exercise did not exploit data from 2022 or 2023 to avoid any possible data leakage, a serious challenge when applying ML to panel data (see Cerqua, Letta and Pinto, 2024).[26]

To test the sensitivity of the estimates with respect to the imputation, we replicated the analysis using only one lag of the covariates on workplace deaths. Consequently, the final forecasts for 2023 rely exclusively on non-imputed data, as workplace deaths in 2022 were not imputed at all. The MSFE, reported in Table B.1, are very close to those reported in Table 3. The lack of worsening performance is not unexpected, as most predictors are the

[26] In addition, in the robustness section presented in Supplemental Appendix B, we will replicate the empirical analysis without using imputed data as predictors and demonstrate that the ML model can still provide highly forecasts more accurate than the benchmark approach.

same as the most predictive covariates selected by the preliminary random forest (see Figure A.2 in Supplemental Appendix A).

We then repeated the analysis on the original INAIL data, i.e., without any imputation of workplace deaths. The MSFEs are reported in Table B.2 for each ML algorithm and the naïve estimator. We observe a sharp deterioration in the accuracy of the ML algorithms, which, under some circumstances, especially with respect to random forest, perform even worse than the naïve estimator. This is not unexpected, as all the ML algorithms were trained on "peculiar data" that do not accurately reflect the usual pattern of on-the-job workplace deaths. Notably, under these circumstances, the LASSO estimator generally performs better than the other ML algorithms do but only when 20 or 30 predictors are used. To understand why LASSO performs well with 20 predictors, we examined the coefficients of the final model and found that only 6 variables have a non-zero coefficient. None of these variables pertain to past data on workplace deaths.

Overall, it is commonly recognized that in circumstances as unique as the COVID-19 pandemic, data imputation is the preferred approach (Knutson et al., 2023; Msemburi et al., 2023). However, it is also important to acknowledge that imputed data can result in a loss of information and potentially diminish the ability of ML algorithms to leverage all available information fully. If one could employ actual data on workplace deaths directly attributable to COVID-19, the imputation would become even more realistic, likely resulting in even more accurate forecasts.

Table B.1: Performance with only 1 lag of on-the-job workplace death predictors

	10 variables	20 variables	30 variables	ALL (97) variables	Smallest MSFE
LASSO	1.744	1.746	1.712	1.780	1.712
PLS	1.723	1.704	1.855	2.022	1.704
Random Forest	1.800	1.802	1.835	1.835	1.800
Stochastic gradient boosting	2.299	2.216	1.821	2.015	1.821
LSTMs	1.913	1.938	2.065	2.110	1.913
Naïve			3.031		

Notes: We report the parameters selected via the panel CV for the best performing case of each ML algorithm.

Table B.2: Performance using non-imputed on-the-job death data

	10 variables	20 variables	30 variables	ALL (131) variables	Smallest MSFE
LASSO	4.022	1.725	2.006	3.765	1.725
PLS	50.831	3.766	13.707	2.350	2.350
Random Forest	8.576	5.020	4.867	7.992	4.867
Stochastic gradient boosting	4.239	2.998	3.357	2.900	2.900
LSTMs	2.072	4.299	4.098	3.760	2.072
Naïve			3.031		

Notes: We report the parameters selected via the panel CV for the best performing case of each ML algorithm.

Supplemental Appendix C – The effect of public subsidies

Table C.1: Impact of an increase in public subsidies for occupational safety and health on workplace fatalities

	ATE estimate	Standard error	Lower CI	Upper CI
OLS	-0.504*	0.287	-1.067	0.059
Lasso	-0.348	0.254	-0.846	0.149
Ridge	-0.217	0.251	-0.708	0.274
Elastic net	-0.349	0.254	-0.847	0.149
Random forest	-0.236	0.264	-0.754	0.281
Boosting	-0.366	0.259	-0.875	0.143
Neural networks	-0.534**	0.249	-1.023	-0.046

Notes: The table shows the province-level impact of an increase in the absolute amount of public subsidies for occupational safety and health received by a given province between 2018-2019 and 2022-2023 on the change in the number of workplace fatalities per 100,000 inhabitants between 2019 and 2023. All provinces experiencing such an increase are considered treated. Treatment and outcome regressions for debiasing include the change in the lagged values (2018-2022) of the set of confounders described in Table A.4 in Supplemental Appendix A. The number of observations is 95. Except for LASSO, Ridge, and Elastic Net, for which cross-validation is performed, all other models use default values for the tuning hyperparameters. Cross-fitting with 5 folds was employed for all models except OLS.

Table C.2: Impact of an increase in public subsidies for occupational safety and health on workplace fatalities in high-risk areas

	ATE estimate	Standard error	Lower CI	Upper CI
OLS	-2.216***	0.797	-3.778	-0.655
Lasso	-0.934	0.742	-2.388	0.520
Ridge	-1.657**	0.734	-3.095	-0.219
Elastic net	-0.934	0.742	-2.388	0.520
Random forest	-1.345*	0.775	-2.862	0.174
Boosting	-0.878	0.756	-2.360	0.604
Neural networks	-0.960	0.692	-2.315	0.396

Notes: The table shows an increase in the absolute amount of public subsidies for occupational safety and health received on the change in the number of workplace fatalities per 100,000 inhabitants between 2019 and 2023 in areas classified as 'high-risk', i.e., *in the highest decile of the probabilistic risk distribution for 2023, according to the best-performing algorithm (LASSO) of the independent ML forecasting analysis run at the province level. Only these provinces are considered treated. Treatment and outcome regressions for debiasing include the change in the lagged values (2018-2022) of the set of confounders described in Table A.4 in Supplemental Appendix A. The number of observations is 95. Except for LASSO, Ridge, and Elastic Net, for which cross-validation is performed, all other models use default values for the tuning hyperparameters. Cross-fitting with 5 folds was employed for all models except OLS. Stars denote statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.*