Evaluating the effect of local monitoring on nuclear safety: evidence from France*

Romain Bizet†[1], Petyo Bonev[1], and François Lévêque[1]

[1]Mines ParisTech, PSL - Research University, CERNA - Centre for Industrial Economics, i3, CNRS UMR 9217

May 28, 2017

Abstract

This paper empirically studies the deterrence effect of monitoring intensity on safety care and compliance with a self-reporting mechanism. To do so, we analyse the consequences of a policy requiring French Departments to subsidize local commissions whose role is to monitor nuclear stations and to communicate with local populations. Our analysis uses a dataset of safety incidents reported by French nuclear stations between 2008 and 2015, and data describing the annual budgets of their monitoring agency. We address the endogeneity of the indirect measurement of local monitoring intensity by using forecasting errors in Department-level operating budgets as an instrumental variable. Contrarily to a large part of the literature on enforcement and compliance, we also account for the possibility of non-detection of non-compliance by plant managers, and disentangle the effects of the policy on safety care and on non-compliance with self-declaration guidelines. We find that plant managers react to informational incentives: although we observe no significant increase in safety care, we find significant increases in compliance with declaration criteria: a 2.000 € increase in the annual budget of a local commission leads, on average, to a 1% increase in the total number of events reported.

Keywords: nuclear power, safety, threat of regulation, local monitoring, transparency, incident data.

*We would like to thank the French nuclear safety authority and the personnel of all commissions for local information for providing us with the data that motivated this study and for providing us with multiple insights regarding the nature of their activities. We also thank all participants to the 6th ELAEE conference in Rio de Janeiro, to the 22nd YEEE seminar in Nürnberg Energy Campus, and to the 2017 ZEW Energy Conference in Mannheim for their feedbacks on previous versions of this paper. All errors are entirely our owns.

†Corresponding author: romain.bizet@mines-paristech.fr, 60 boulevard Saint-Michel, 75006, Paris, France.

1 Introduction

In this paper, we empirically assess the effect of a French policy passed in 2006, which enacts the monitoring of French nuclear power stations by local commissions whose purpose is to monitor nuclear plant managers and to disseminate information regarding nuclear power to local citizens. Through their monitoring activities and by disclosing information to the public, we hypothesize that local commissions may induce shocks on the costs faced by nuclear plant managers when reporting, or failing to report, safety events. Yet, as has been shown in the literature dedicated to environmental regulation through self-reporting mechanisms, the change in incentives induced by increased sanctions (e.g. the cost of reporting a non-compliant situation) or increased penalties (e.g. the expected cost of hiding a non-compliant situation) can be quite different. In particular, if increases in sanctions and penalties should theoretically lead to unambiguous increases in safety care, they may have countervailing effects on the decision to comply with self-reporting guidelines. Therefore, the aim of this empirical paper is twofold. First, we want to assess whether the activities of these commissions, measured by their annual budgets, influences the behaviours of French nuclear plant managers. If so, we also want to investigate how they influence these behaviours. In particular, we aim to disentangle the effort exerted by managers in order to improve the safety of their power station from their decision to comply with declaration guidelines when they detect significant safety events.

In the context of environmental regulation, informational policies such as public disclosure of information have been shown to increase the compliance of polluting firms with existing environmental regulations, and to decrease their levels of emissions (see e.g. Shimshack (2014) for a review). In this paper, we investigate whether these results carry forward to the nuclear industry, in which standard-based safety regulations are similarly implemented by centralized safety authorities, which deter non-compliance by conducting audits and levying penalties on non-compliant facilities. In addition to this classical command-and-control safety regulation, France implements since 2006 an informational policy by which Departments[1] hosting nuclear stations have to subsidize local commis-

[1]France's administration is organized in several levels below the national government. The French territory is first divided into thirteen administrative regions. Regions are then divided in a total of a hundred Departments, which are divided in over thirty-six thousands counties.

sions, who monitor the activity of their local nuclear station, and can communicate with the local population on the results of this monitoring.[2] We empirically study the effect of this monitoring policy on the decision of local plant managers to exert safety care and to comply with self-declaration criteria in the French nuclear industry.

Despite the importance of this question, it has remained largely unanswered by the literature. Its empirical evaluation is hampered by three major problems. The first one, inherited by the general specifics of empirical evaluation of nuclear safety, is the one of data scarcity. In particular, severe nuclear accidents are very rare, so that proper statistical analysis in this context is hardly available. The literature has dealt with this issue by either using Bayesian methods in the context of technical probabilistic risk assessment (Rangel and Lévêque, 2014), or by using extended sets of accidents from both nuclear power plants and fuel-cycle facilities (Sovacool, 2008; Hofert and Wüthrich, 2011; Burgherr and Hirshberg, 2014; Wheatley et al., 2016). The second problem stems from the unobserved nature of non-compliance. Non-compliance is economically meaningful only if perfect observability of the behaviour of the agent by the principal is not feasible. A first solution to this problem is to study detected non-compliance, as is done in Feinstein (1989). In our case, detected non-compliance is not observable to the econometrician. This implies that the information contained in observed reporting behaviours is ambiguous, as it is affected by both safety care and non-compliance. Finally, the intensity of the monitoring of an agent may be partially determined by how safe the agent is perceived by the principal. As a result, safety care and non-compliance are determined jointly with the intensity of the monitoring, which might induce a simultaneity bias in the estimation.

We contribute to the literature on nuclear safety and non-compliance in several ways, with a particular focus on the three aforementioned problems. First, we use a rich novel dataset on declared safety incidents in French nuclear power stations. These incidents, although of small magnitude, consist in deviations from the safety standards and operation guidelines set by the French safety authority. They are therefore considered as

[2]Most of these Departments had already created such agencies before 2006, on a voluntary basis. The law made the existence of these agencies mandatory in all Departments hosting either nuclear reactors or fuel-cycle facilities.

significant for nuclear safety.[3] The declaration of these events by plant managers to the safety authority is mandatory. Yet, these events may remain undetected, and managers face countervailing incentives when choosing whether to report detected events. Reporting an event may be costly to the manager, for instance if his salary is based on the safety performance of the plant, or if he has to incur some cost after the declaration in order to redeem compliance. Not reporting an event can be costly as well, in case of a public backlash or in case of more stringent regulations. To perform our analysis, we also use data on the monitoring exerted by local commissions between 2008 and 2015. This local monitoring can be more or less thorough, as Departments in charge of subsidizing local commissions provide them with very heterogeneous levels of resources. As a result, the intensity of the monitoring performed by these commissions may vary from organizing meetings with local managers where safety issues are discussed, to hiring independent experts to assess, for instance, the environmental impact of the operation of the power station. We use the annual budget of each commission as a measure of the intensity of their monitoring activities. The rich heterogeneity in the size of these budgets provides a powerful source of identification.

Second, to disentangle the effects of local monitoring on safety care and non-compliance with declaration guidelines, we present a formal model for the behaviour of plant managers adapted from the environmental literature (see e.g. Evans et al. (2009); Gilpatric et al. (2011)), in which the decisions to exert safety care and to comply with declaration guidelines are jointly determined by the perceived sanction incurred by the manager when reporting an event, and the expected penalty faced when non-reporting. We derive testable hypotheses from this model and test these by using several observables such as specific or general counts of reported safety events, and nuclear reactor reliability indices. In particular, in a strategy related to Hausman (2014), we look at automatic shut-downs of reactors, a type of events that is perfectly detected by managers and reported to the authority due to the resulting changes in the production of electricity. As non-compliance with declaration guidelines is not possible for this subset of events, the effect of local

[3] Through probabilistic risk and reliability analyses - a process of case-by-case scenario analysis performed by the operator and the regulator - an incremental probability of nuclear core meltdown is associated to each of these events.

monitoring on the sanction perceived by managers is identified. In a second step, the estimated effect of monitoring on safety care can be measured by observing evolutions of the reliability of nuclear power plants. Finally, conditionally on the results of the former regressions, we can identify the effect of monitoring on compliance by measuring the effect of monitoring intensity on the reports of significant safety events.

Third, we address the endogeneity of the measurement of monitoring intensity by using an instrumental variable (IV). Our instrument is based on the difference between the forecast and the realized annual operating revenues of the French Departments. This forecasting error has several attractive features. First, a forecasting error, once realized, might lead to a reassessment of the forecast for the current or coming fiscal year, and thus induce a change in the budget of the monitoring commissions. Second, such a forecasting error is almost per definition unanticipated, which prevents endogenous forward-looking behaviour of the local authorities. Finally, the source of the error is simply a financial miscalculation due to overall uncertainty or human failure related to tax returns, and thus it can be argued that it is not related to the unobserved factors affecting compliance and safety at the level of nuclear power stations. The second and the third property of these errors qualify them as a quasi-natural experiment. This instrumental variable is similar in spirit to the natural experiment used in Bressoux et al. (2009), who utilise random administrative mistakes to instrument for the endogenous assignment of teachers to schools in France.

To the best of our knowledge, Davis and Wolfram (2012), Hausman (2014) and Feinstein (1989) are the only papers that analyse the impact of economic incentives on nuclear safety and non-compliance. Davis and Wolfram (2012) and Hausman (2014) identify the effect of market deregulation in the U.S. on the reliability and safety levels of some US nuclear reactors. The proxy used in Hausman (2014) for safety consists of automatic reactor shut-downs and is thus closely related to our identification strategy. Likewise, our use of reliability indices is similar to the empirical assessment performed by Davis and Wolfram (2012). Our main focus, however, is on non-compliance instead of safety and we use the reliability indices and counts of automatic shut-downs only to disentangle the effect of monitoring on safety care from its effect on non-compliant behaviours in a back-door-identification-type strategy. Feinstein (1989) uses data on inspections of US power plants to study the factors of non-compliance and the effect of non-compliance on safety. His identification depends crucially - much in the spirit of the time in which his paper was written - on strong parametric assumptions on the distribution of the unobservables. These assumptions, however, are not guided by economic theory and are in this sense rather arbitrary.

More generally, within the literature dedicated to monitoring on compliance, our paper is closely related to the work of Duflo et al. (2013) and Telle (2013), who use randomized controlled trials to assess the effectiveness of monitoring programs on self-reporting behaviours. Telle (2013) studies the effect of deterrence on self-reporting, and shows that whereas specific deterrence (e.g. fines) does increase self-reporting, an increased frequency of audit does not. Duflo et al. (2013) show that preventing conflicts of interests in audit mechanisms leads to less non-compliance and more mitigation efforts. In the environmental literature, another related paper is Lin (2013), who assesses the effect of increases in the probability of being monitored on mitigation efforts and truthful self-reporting, using rainfalls as an instrument for the monitoring probability. Compared to these three papers, we investigate similar questions in a different industry, and use an instrumental variable different from the one used in Lin (2013), and instead of a randomized experiments. We also contribute to this literature by introducing possible non-detection of events by the firm and imperfect audit results for the regulator.

Our main results are that nuclear plant managers react to informational incentives provided by this French local monitoring program. These incentives do not induce increases in safety care, but significantly reduce non-compliant behaviours. In particular, a 2.000 €increase in the annual budget of a commission leads local managers, in expectation, to increase by 1% their level of compliance with declaration guidelines. Our results are robust to a variety of specifications. While the non-significant impact of monitoring intensity on safety care contradicts the findings of Hausman (2014) and Duflo et al. (2013), the positive effect on compliance is in line with the findings of Feinstein (1989); Telle (2013) and Duflo et al. (2013).

2 Institutional setup

2.1 Nuclear-power safety in France

The French nuclear fleet is constituted of 58 reactors, located in 19 sites (or plants in the following), owned by a single utility, EDF. Nuclear safety is regulated by the Nuclear Safety Authority (ASN in the following) who sets technical standards regarding the construction, operation and maintenance of all nuclear reactors. In addition, the safety authority establishes declaration criteria which characterize a set of events considered as significant deviation from safety. Upon the detection of any of these events, a nuclear plant manager has to declare the event to the safety authority. This self-reporting mechanism aims to foster knowledge spillovers across reactors, and to detect generic design weaknesses or organizational failures, in order to improve nuclear safety.

Non-compliance with this self-reporting mechanism is deterred by the use of periodic and random inspections by ASN inspectors. During these inspections, inspectors access the paperwork describing the events detected by plant managers but not considered significant enough for reporting. The firm can be prosecuted for failing to declare significant safety events. In addition, anecdotal evidence suggests that failing to declare safety events can have other costly consequences for plant managers such as production losses due to public backlashes.[4]

Although all French nuclear plants are owned by a single firm, many decisions are delegated to the management of each plant. For instance, the reporting of safety events has to be done rapidly after detection, and is thus left to the discretion of power plant managers. Thus, the reporting of significant safety events captures the incentives faced by plant managers when deciding on how much care to dedicate to the limitation of their occurrences, and whether to report observed events to the safety authority. Though, these incentives may be countervailing.

First, if the occurrences of significant safety events lead to extended maintenance periods, or to the shut-down of a reactor, exerting care to limit these occurrences may

[4]For instance, the French station Fessenheim was invaded by Greenpeace activists in 2014, after it became public that its managers had understated the magnitude of an incident that happened earlier that year.

yield private benefits, for instance through an increased level of reliability of the power stations (e.g. larger production levels). Yet, these private benefits will be offset by the costs of exerting care, such as hiring more staff, fostering safety culture and skills through dedicated trainings, or investing in better equipments.

Second, reporting can be considered costly to plant managers, as they face sanctions when reporting safety events. Regulatory sanctions can consist in necessary investments required by the authority to mitigate the causes of the reported safety events.[5] Internal sanctions could also deter reporting, for instance if managers are incentivized to enhance their plant's safety performance, based on the counts of these events.[6] On the other hand, failing to report a safety event may also be costly to the manager, due to the possibility of undergoing prosecution if the authority discovers the event, or in case of a public backlash against the power station due to the lack of transparency of its managers.

Finally, it is to be noticed that a significant event can remain unreported for two distinct reasons. First, after observing the event, an operator can deliberately decide not to report it, because the cost associated with the report is larger than the expected cost associated with hiding the event. On the other hand, an event may also not be reported because the manager failed to detect it in the first place. Nevertheless, the consequences of non-reporting seem to be unaffected by the cause of the non-reporting.

In the next paragraphs, we describe the organization of the French local monitoring commissions, and the way they may interplay with the reporting of safety events.

2.2 Local monitoring

Since 1981 and the partial meltdown that occurred at the Saint-Laurent-des-Eaux nuclear power station - two years after the Three Mile Island accident - some French departments hosting nuclear stations have been involved in the organization a form of local monitoring of nuclear power stations through dedicated commissions whose purpose was to foster transparency regarding nuclear power.[7] In 2006, a law made the existence of these moni-

[5]In the environmental literature, these sanctions are usually thought of as a tax on reported emissions.

[6]Although the existence of such incentive schemes is not clear, anecdotal evidence on the distribution of event counts suggest it.

[7]In a note circulated in 1981 to local prefects - local state representatives - the French Prime Minister suggested the creation of local commissions dedicated to the monitoring of industrial activities prone to

toring commissions compulsory in all French Departments hosting nuclear power reactors or fuel cycle facilities.

These monitoring commissions are now composed of four groups of members: locally elected officials (mayors from cities neighbouring the power station or regional counsellors), members of local environmental associations, members of the nuclear plant workers unions, and competent local citizens.[8] These members are not remunerated for their participation, and some restrictions regarding the composition of the commissions are set by law. Elected officials must represent at least 50% of the commission, while each of the other three groups has to constitute at least 10% of the members.

Commissions are funded by the French Departments, as well as by the Nuclear Safety Authority. As the law does not set any rule regarding their budgets, the ASN typically matches, for each commission, the endowment granted by its Department. Local commissions obtain very heterogeneous budgets, which span between 5,000 €/year to more than 190,000 €/year. Due to these variations in endowments, commissions undertake heterogeneous activities, which we now describe in more details.

First, every commission has to organize at least two periodic meetings a year, during which plant managers and the safety authority present the main actions undertaken in the nuclear station to the commission. Commission members are provided with a set of documents regarding the operation of the nuclear facility to prepare the meeting, and may ask for specific topics to be addressed during the meetings. In particular, they receive an account of the occurrences of significant safety events within each reactor of their local station. Some commissions voluntarily set this number of periodic meetings to three per year.

Based on these meetings, commissions disseminate the information obtained to the public. To do so, most commissions invite the press to the periodic meetings, and often make public statements regarding the major decisions made by the plant managers or by

large risks and to the information of the public. At the time, this suggestion aimed to promote a sharing of responsibilities among local collectivities and the State regarding the information of populations about the nuclear risk. The original note authored by Prime Minister Pierre Mauroy (in French) can be downloaded on this website: http://www.cli-gravelines.fr/Services-en-ligne/Espace-documentaire/Documents-a-telecharger/Les-textes-reglementaires/Circulaire-MAUROY-du-15-decembre-1981.

[8]Additional examples can be found on the website of the commission of the Paluel and Penly nuclear power stations.

the safety authority. Depending on their budgets, commissions also publish contents on dedicated websites, distribute journals in city halls, or mail periodic information letters to neighbouring populations. A minority of commissions even organize additional open meetings for interested local inhabitants, and invite local populations from neighbouring countries[9].

Finally, if their budgets allow it, local commissions can hire independent experts in order to carry out assessments of some aspects of the operation of the plant. For instance, past investigations have assessed the environmental impacts of the operation of nuclear stations through radioactivity measurements in local water streams. Results of these investigations can then be discussed during the periodic meetings of the commission with the plant managers and the safety authority.

3 A model of monitoring and compliances

In summary, the existence of countervailing incentives faced by plant managers is important to our analysis as the variety of actions undertaken by local monitoring commissions may alter the different costs and benefits incurred by managers when exerting safety care and when deciding whether to report significant safety events. To shed light on the effect of local monitoring on the behaviour of plant managers, we present in the next section a theoretical framework derived from a classical principal-agent model, which captures the interplay of local monitoring with the incentives for reporting. We derive from this framework the hypotheses that enable the identification of our empirical results. Contrarily to the existing theoretical literature on audit mechanisms[10], we do not model explicitly the optimization problem of the regulator, and only model the best-response of an agent to the exogenous audit mechanism set by the principal. The determination of the optimal audit mechanism is irrelevant in our context, as our empirical estimation will consist in using an instrumental variable method to assess the effect of exogenous changes in

[9]Since 2015 and France's new energy transition law, the organization of a third meeting, open to the public, is mandatory for each commission. But as this law has not been implemented yet, the existence of these public meetings is out of the scope of our study.

[10]See for instance e.g. Macho-Stadler and Pérez-Castrillo (2006); Evans et al. (2009); Gilpatric et al. (2011), and Zahran et al. (2014).

monitoring intensity on the behaviour of plant managers.

3.1 The model

In the following, we suppose that an agent (the manager) operates a nuclear power reactor subject to a self-reporting mechanism enforced by a principal (the safety authority). This model is adapted from Evans et al. (2009) and Gilpatric et al. (2011) who introduced the possibility of imprecise monitoring technology in environmental emission auditing mechanisms.

Formally, let E_{tot} be a continuous variable capturing the total number of events that occur during a year in a nuclear power reactor. For tractability of the model, we forego the count nature of these events, and assume that this quantity E_{tot} decreases when the agent increases his level of safety care. In other words, we assume that the agent can choose the number of events E_{tot} that occur each year in his nuclear reactor.[11] We further assume that safety care is costly, but provides some private benefits, such as increased reliability of the power station. Hence, we assume there is a function $B(E_{tot})$ concave in E_{tot} with $B'(0) > 0$ and $B'' < 0$, that captures the costs of safety care and its private benefits. In the absence of a principal, the agent would privately choose a level of care associated with a number of events $\bar{E}$, satisfying $B'(\bar{E}) = 0$. We assume that B is a concave function that reaches its maximum, e.g. that $\bar{E}$ is finite.

Second, in order to model the detection abilities of the agent, and his level of compliance with declaration guidelines, we assume that a fraction ρ of events are privately observed by the agent, and that a fraction z of privately observed events are declared to the principal. In other words, the agent observes $E_{obs} = \rho E_{tot}$ and declares to the principal the quantity $zE_{obs} = z\rho E_{tot}$. In the following, we assume that the detection ability of the agent is exogenous. In this case the agent is only left with the discretion of choosing safety care (E_{tot}) and the level of compliance (z).[12] In the appendices, we let

[11] If the agent can exert a costly level of safety care s, and if E_{tot} is a decreasing function of s, and if s only affects E_{tot}, thn choosing s or choosing E_{tot} is equivalent in terms of the agent's preferences.

[12] It is to be noticed that holding ρ constant does not mean that detection ability is constant, but rather that the rate of privately observed events to the total number of events remains constant. This can be interpreted by saying that increasing safety care reduces the total number of events occurring in a reaction and reduces proportionally the quantity of events which are not privately observed by the plant manager.

ρ be endogenously chosen by the agent, and comment on how the existence of this third behavioural channel affects our results.

The fact that some events are not observed by the agent may be due to an imperfect knowledge of his equipment, or to a limited time spent trying to detect these events. Yet, as the agent knows that he may fail to detect a certain number of events, ρ is assumed to be known by the agent. In other words, the agent observes E_{obs}, but knows the total number of observable events E_{tot}. Yet, given the nature of the reporting mechanism, which requires to provide numerous details about the causes and consequences of each event, we assume that the agent cannot report an event which he did not really observe.

In addition, as inspectors have limited time to perform their inspections, we will assume that audits do not perfectly reveal all unreported events. More specifically, we assume that audits may reveal any safety event, regardless of whether they were observed by the agent. To do this, we adapt the model developed by Evans et al. (2009) and Gilpatric et al. (2011) to capture the imprecision that characterizes CO_2 emissions measurement technologies. Note that a crucial difference from their frameworks of Evans et al. (2009) and Gilpatric et al. (2011), is that the imprecision of the audit describes both the fact that the agent may fail to detect safety events, and the fact that the principal has only limited audit resources. The novelty in our case is the fact that the agent may endogenously affect the outcome of the audit.[13]

Hence, let u be a random variable distributed according to a cumulative distribution F and density f over $[0;1]$. The value taken by u represents the fraction of E_{tot} detected by the principal during audits. When u takes values between z and 1, the audit reveals a number of events larger than what the manager publicly reported. When u takes values between ρ and 1, the audit reveals a number of events larger than what the manager privately observed. This captures both the imperfection of inspections performed by the principal, and the possibility of non-detection of safety events by the agent. The expected

[13]More specifically, when choosing ρ, the agent can affect the extent to which the principal may find more events than the quantity declared by the agent. The distribution of u, however, cannot be affected by the agent.

quantity Q of unreported events revealed by the audit is:

$$Q = E_{tot} \int_{z\rho}^{1} (u - z\rho) f(u) du. \tag{1}$$

Upon the observation of an event, the agent can report the event and face a sanction α, which embodies the direct consequences of reporting, such as mandatory investments. Using previous notations, the sanction associated with reporting $z\rho E_{tot}$ events is $\alpha z \rho E_{tot}$. The agent may also decide not to report the detected event to avoid the sanction. To deter this behaviour, the principal audits the agent with a given probability q, and levies a penalty β when unreported events are detected. The probability of inspection can be thought of as the frequency of planned or unplanned inspections. The penalty embodies the consequences of non-declaration, such as legal prosecution, public backlashes or increases in regulatory oversight. The expected penalty faced by the agent is thus $q\beta Q = q\beta E_{tot} \int_{z\rho}^{1} (u - z) f(u) du$.

Under this self-reporting mechanism with imperfect audits and imperfect observation of events, a risk-neutral agent maximizes the following quantity:

$$\max_{E_{tot}, z} B(E_{tot}) - \alpha z \rho E_{tot} - q\beta E_{tot} \int_{z\rho}^{1} (u - z\rho) f(u) du \tag{2}$$

3.2 Comparative statics and testable hypotheses

Suppose that detection ability ρ is exogenous. Let $z^\star$ and $E_{tot}^\star$ be the best response played by the agent given an exogenous auditing mechanism characterized by α and $q\beta$. Provided $\alpha < q\beta$, the existence of an interior solution for $z^\star$ is ensured. This condition captures the fact that if the perceived sanction for reporting is higher than the expected penalty for non-reporting, then the agent never reports and an interior $z^\star$ cannot exist. Likewise, provided $\mu(z^\star) < B'(0)$, there exist an interior $E_{tot}^\star$ that maximizes equation (2).[14]

We can then derive the following comparative statics describing the effect of a change

[14] If $\mu(z^\star) > B'(0)$, the agent exerts a level of safety care associated with no occurrences of safety events. In other words, we can interpret this assumption as the fact that reducing the number of safety events to 0 would be infinitely costly, so that there exist an interior solution for the optimal level of safety care.

in the value of parameters α and $q\beta$ on safety care $E^{\star}_{tot}$, compliance $z^{\star}$ and the total observed quantity of reports $z^{\star}\rho E^{\star}_{tot}$, which we note $z^{\star}E^{\star}$ for simplicity.

Proposition 3.1 *Comparative statics*

At an interior solution, the following results hold:

- $\frac{\partial E^{\star}}{\partial \alpha} < 0$ *a marginal increase in* α *leads to a decrease in* $E^{\star}$*,*
- $\frac{\partial z^{\star}}{\partial \alpha} < 0$ *a marginal increase in* α *leads to a decrease in* $z^{\star}$*,*
- $\frac{\partial E^{\star}}{\partial q\beta} < 0$ *a marginal increase in* $q\beta$ *leads to a decrease in* $E^{\star}$*,*
- $\frac{\partial z^{\star}}{\partial q\beta} > 0$*: a marginal increase in* $q\beta$ *leads to an increase in* $z^{\star}$*,*

This result is identical to the result derived by Evans et al. (2009). A proof is proposed in the appendices. Two interesting direct corollaries from this proposition are interesting to notice. First, the first two comparative statics in proposition 3.1 show that a marginal change in the level of perceived sanctions α has an unambiguous effect on the total quantity of reports $z^{\star}E^{\star}$. Second, the last two comparative statics show that a marginal change in $q\beta$ has an ambiguous effect on $z^{\star}E^{\star}$. To see this, we can write:

$$\frac{\partial z^{\star}E^{\star}}{\partial q\beta} = E^{\star}\frac{\partial z^{\star}}{\partial q\beta} + z^{\star}\frac{\partial E^{\star}}{\partial q\beta} \quad (3)$$

We know that the first term in the right-hand side of (3) is positive, while the second term is negative. Therefore, the variation in observed reports induced by a marginal change in $q\beta$ is determined by the relative size of these two terms, and in particular by the relative amplitude of the variations in compliance and in safety care. For instance, if $\frac{\partial E^{\star}}{\partial q\beta}$ is small enough, then an increase in $q\beta$ should be followed by an increase in the quantity of observed reports.

Proposition 3.1 and its two corollaries depict the fundamental problem of identification in our setup. First, any observed change in the observable outcome $z^{\star}E^{\star}$ can be due to changes in either perceived sanctions α or expected penalties $q\beta$. Moreover, conditionally on a constant perceived sanction α, the effect of a marginal change in $q\beta$ on the observable outcome $z * E*$ consists of two effects: the effect on compliance: $E^{\star}\frac{\partial z^{\star}}{\partial q\beta}$, and the effect on safety care: $z^\star \frac{\partial E^\star}{\partial q\beta}$. As these two effects have opposite signs, any change in $z^\star E^\star$ can be explained by either an increase or a decrease in $q\beta$.

Thus, the fundamental problem of identification consists in assessing whether the activities undertaken by local monitoring commissions induce changes in perceived sanctions or changes in expected penalties, and then to disentangle the different channels which determine the compound effect of monitoring on the observable outcome $z^\star E^\star$. We refer to this identification issue as the *channel identification problem*. It adds additional complexity to the problem of potentially endogenous level of monitoring and motivates our identification strategy, which we outline in section 5.

4 Data and descriptive statistics

The data we use to conduct this study emanates from three sources: the French Nuclear Safety Authority, the French utility EDF, and fourteen local monitoring commissions. Our unit of observation is set at the level of the reactor.year. In other words, each observation in our dataset consist of a pair (reactor, year). Our dataset consists in an unbalanced panel of 234 observations of reactor-years, spread across 50 different nuclear reactors observed between 2008 to 2015. As the French fleet contains 58 nuclear reactors, the largest possible dataset that we could have gathered over the same time period would have contained 464 observations. The 8 missing reactors are located in XX sites, whose commissions could not provide us with any data regarding their activities. The rest of the missing data is due to the fact that many commissions could not provide us with financial data prior to 2010. The following paragraphs describe our different variables as well as their sources.

4.1 Treatment variables

We first gathered data describing the activities of local monitoring commissions. This dataset was constituted based on the annual activity reports of these commissions, which contain information regarding their annual budgets, e.g. the endowment received from the Department councils and the subsidies granted by the ASN. We also retrieved in-

Table 1: Descriptive statistics: treatment variables.

Variable		Mean	Std. Dev.	Min.	Max.
Commission controls	*budget*	52.415	48.146	4	198
	meet	2.271	0.446	2	3
	multiple	0.169	0.376	0	1
	SaintLaurent	0.051	0.22	0	1

236 observations in 50 reactors from 2007 to 2015 (522 possible)

Note: This table describes the activities of local monitoring commissions, measured by several observables characterizing the frequency of their meetings and their communication strategies.

formation regarding the administrative statuses and composition of the commissions. We finally gathered data regarding the frequency of their meetings, and whether these meetings are open to the press or the public. We also counted the independent studies mandated by these commission. Finally, we identified the commissions which have multiple facilities to monitor, as some nuclear sites in France host more than one nuclear facility. Descriptive statistics regarding these commissions are gathered in table 1. The budgets of the commissions varies in the data from 4 000 €/year to 198 000 €/year. Notice also that most commissions only organize two meetings a year, and that only two of them have multiple sites to monitor.

We use the annual budgets of local commissions as a proxy for the intensity of their monitoring activities. We deem this variable to be a good proxy for the intensity of the commissions as these budgets are used to finance environmental impact assessments by independent experts, in order to train commission members, or to pay for the diffusion of the information gathered by these commissions. Commissions endowed with larger budgets are thus likely to be able to induce greater shocks on the perceived sanctions and expected penalties faced by plant managers when deciding how much safety care to exert and whether to report observed events.

On the other hand, these budgets are an endogenous measure of the intensity of the monitoring performed by the commissions. Commissions budgets and the behaviour of plant managers may be simultaneously determined, as budgets can increase because commission members expect the safety of their local plant to be diminished, while managers may increase safety care and compliance because when they expect local commissions to exert more monitoring pressure on them thanks to larger budgets. This endogeneity issue is addressed in section 5.

The other treatment variables described above are also prone to endogeneity. For instance, the number of yearly meetings can also be affected by the reverse causality described in the previous paragraph. Yet, it can be noticed that only the budget variable is varying over time and across nuclear sites, whereas the other variables, such as the number of meetings or the variables describing the communication of each commission, are time invariant. This means that if these variables are endogenous, they have to be related to the expectations held by commission members when their values were set. In this sense, the endogeneity of these variables is less troublesome than the one characterizing the commissions' budgets.

4.2 Reporting and reliability

As a proxy for nuclear safety, we use a dataset obtained from the French Nuclear Safety Authority which contains the significant safety events reported by plant managers. Although these events only have minor consequences, their number is substantially larger than the number of nuclear accidents. This dataset contains over 19.000 safety events, declared between 1972 and 2015 in currently operated nuclear power stations. We use this dataset on the period 2008-2015, to match our data regarding the local commissions' activities.

Within this dataset, we focus on counts of events annually reported in the French reactors. In order to implement the identification strategy described in the following section, several counts of events will be considered: the count of all events declared during a reactor-year (ALL), and specific counts of events declared during a reactor-year, such as automatic shut-downs (ASD) or unplanned uses of safeguard mechanisms (SFG). These two types of events were identified jointly with the safety authority as being subject to perfect detection and declaration, a property which we will use to disentangle the effect of monitoring intensity on safety care and compliance with declaration guidelines. Automatic shut-downs have an impact on the electrical output of the power station, and are thus impossible to hide. Events requiring the use of safeguard mechanisms are deemed particularly severe and easy to detect by the authority. These two categories are jointly referred to as perfectly detected and declared events, and measured by the variable PDD. It is to be noticed that we have $PDD = ASD + SFG$.

In order to control for the various differences across reactors that may also explain the occurrences of safety events, we rely on two datasets obtained from the Nuclear Safety Authority and the French utility EDF. These datasets contain detailed information regarding the annual production levels, as well as information regarding the reliability of nuclear reactors. In particular, we use data on the annual length of maintenance activities conducted in each reactor, and on the share of electricity lost due to unplanned maintenance extensions (K_{ipr}) or due to fortuitous stops (K_{if}).

In addition, we construct several variables that account for the history and technological design of the reactors. We first construct an age variable that describes the age of a reactor during the calendar year of observation. Age is defined here as the duration between observation and the year of connection of a plant to the electricity grid.[15] We also construct three design-fixed-effects dummy variables that match the three power plant designs that coexist in the French fleet. In order to capture possible learning-by-doing effects, we finally construct dummy variables which identify the first reactors built within each nuclear site and the first reactors built within the groups of reactors sharing a common plant design. These variables are described further in table 2. It appears that the age of the reactors considered in this dataset ranges from 8 to 37, that sites include from 2 to 6 reactors, each of which produced an average of 7 TWh per year over the elapsed period of time. We can finally see that, on average, reactors undergo 68 days of maintenance every year.

4.3 Attrition bias, local monitoring and reporting behaviours

As our study relies on an unbalanced panel of pairs of reactor-year, table 3 proposes comparisons of means for several observables between the sample studied and the sample of excluded observations, for which we could not obtain data regarding the budget of the local commission. Table 3 shows that some variables take significantly different values

[15] Other possible definitions of the age of a reactor is the time since the beginning of its construction, the time since the first divergence of its nuclear core, or the time since its first sale of electricity.

Table 2: Descriptive statistics: reactor-level data.

Variable		Mean	Std. Dev.	Min.	Max.
Event counts	ALL	12.856	4.778	2	27
	SDD	1.017	1.13	0	5
	ASD	0.809	0.955	0	5
	SFG	0.208	0.492	0	3
	SDD	1.017	VAR	MIN	MAX
Reactor reliability	K_{ipr}	MEAN	VAR	MIN	MAX
	K_{if}	MEAN	VAR	MIN	MAX
Reactor controls	age	28.169	5.659	8	37
	$size$	3.966	1.38	2	6
	$FOAS$	0.559	0.498	0	1
	$FOAK$	0.008	0.092	0	1
	$production$	6.866	1.747	2.165	11.622
	$maintenance$	67.568	49.839	0	279

236 observations in 50 reactors from 2007 to 2015 (522 possible)

Note: the first four variable describe counts of safety events reported each year by each reactors. The second two variables capture the reliability of each nuclear reactor, measured by the share of electricity lost to unplanned maintenance works and to fortuitous stops. Finally, the last variables are reactor controls that describe their yearly production levels and some of their technical features.

in each sample. The difference in mean age is intuitive, as most missing data would characterize the period 2008-2010. The difference in reports (ALL) suggests that there may be some degree of attrition bias, but the difference is not highly significant.

In addition, and similarly to the descriptive lowess regression provided by Hausman (2014), we conducted a simple analysis of the quantity of events reported per reactor and per year before and after the creation of their local monitoring commission. This test aims to provide some tentative evidence of the effect of local commissions on the reporting behaviour of nuclear plant managers.

To do so, we first regress the yearly reports of each reactor on yearly dummies, to obtain a time-corrected estimation of the quantity of events reported each year in each reactor. This time correction controls for changing declaration criteria, which may evolve over time. We then plot this quantity against a time indicator, calculated as the difference between the year of declaration and the year of creation of the local monitoring

Table 3: Descriptive statistics: attrition bias.

Variable	Sample mean	Out-of-sample mean	t-statistic	p-value
ALL	12.85	13.67	1.81	0.07
SDD	1.09	1.19	0.89	0.38
ASD	0.87	0.89	0.23	0.82
SFG	0.23	0.30	1.49	0.14
K_{ipr}	0.04	0.05	1.25	0.21
K_{if}	0.03	0.04	2.24	0.03
Age	29	24	9.49	0.00
Production	6.86	7.28	2.50	0.01
Maintenance	66.98	68.59	0.34	0.73

Note: In this table, we compare the mean value of some observables for the pairs of reactor-year within our panel (sample means), with the pairs of reactor-year for which we did not obtain data regarding the budgets of local commissions (out-of-sample mean). For most variables, the difference is not significant. The difference in mean age is normal, as most missing data would characterize the period 2008-2010. The difference in reports (ALL) suggests that there may be attrition bias, but the difference is not highly significant.

commission.[16] Figure **??** shows this plot. Two lowess fits are calculated on the two subsamples defined by the date of creation of local commissions. These lowess fits suggest that the introduction of local monitoring commission was correlated with an increase in the number of events reported annually by nuclear plant managers. Yet, this statistic accounts for neither the intensity of the local monitoring, nor for the endogeneity that characterizes its measurement, an issue that we tackle in the following two sections.

5 Empirical strategy

5.1 Endogeneity of monitoring intensity

Another identification problem that formally and logically precedes the Channel problem is the identification of the total effect (that is, regardless through which mediation channel) of the activity of the commissions on $z^{\star}E^{\star}$. This problem stems from two potential endogeneity sources. First, the budget of these commissions, which we use as a proxy for the intensity of their monitoring activity, is potentially endogenous due to reverse

[16] As was mentionned in the institutionnal setup, some local commissions were created as early as 1981, although there existence only became compulsory in 2006.

Figure 1: Descriptive statistics: local monitoring and reporting behaviours.

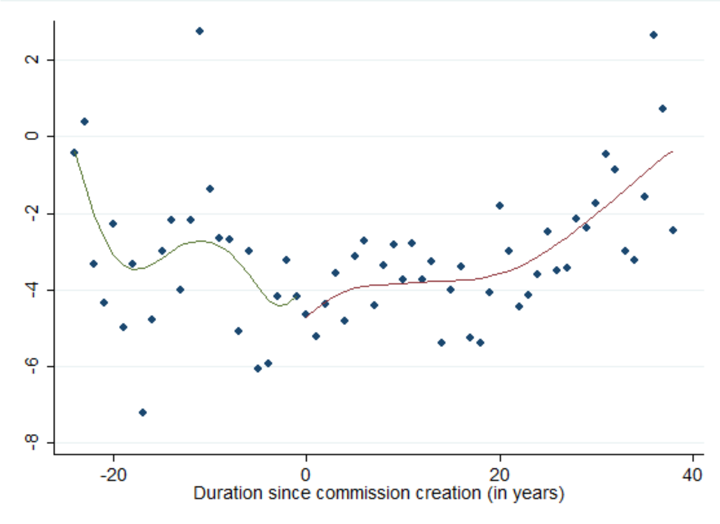

Note: Two-part lowess of the time-corrected annual quantity of events reported by French nuclear reactors. The x-axis shows years since creation of the local monitoring commission. Simple lowess fits suggest an increasing trend in declarations after the creation of local commissions.

causality: a local commission may have a high level of activity because the Department council - which allocates their budget - is aware that the power station has an abnormal level of declaration of events. Second, the incentives provided by the regulator to the plant managers might be related to the intensity of the local monitoring. In a period of intensified political or public debate regarding nuclear power, for instance, both the threat of new regulation and the expenditures for local monitoring might be increased. Since we do not measure the intensity of the regulatory oversight, the intensity of the commissions' activity could be endogenous due to an omitted variable bias (OVB).

In order to solve the two endogeneity issues mentionned above, we use an instrumental variable method. Our instrumental variable, *shock*, is based on a natural experiment triggered by forecasting mistakes. More precisely, we use primitive budget data published every year by Department councils, in which a forecast of the balance of their revenues and expenditures for the upcoming year is provided. We also use *ex post* data on the Departments realized financial revenues in order to compute the forecasting errors made by the Departments. These errors are attractive in many respects. They are first, by nature, unanticipated, which precludes endogenous forward looking behaviours of local authorities. Second, as these errors are the results of a failure from the local officials to predict accurately the revenues levied by local taxes, it seems fair to argue that this error will be independent from the unobserved factors influencing the commissions' budgets. Finally, failing to predict accurately their revenues may lead local officials to reassess the funding provided to local monitoring commissions. As was mentioned in the introduction, this instrument is similar to the natural experiment used in Bressoux et al. (2009).

The instrument *shock* is defined using public financial data from French Departments[17]. More precisely, we used the detailed reports describing both the anticipated and realized budgets for each Department hosting a nuclear power station. These budgets include two main sections: revenues and spendings, whose total are equal. Within each section of these budgets, one can find two main categories: investment revenues/expenditures, and operating revenues/spendings. Both categories are separated: investment revenues can only finance investment expenditures, while operating revenues are used to finance operating expenditures. Operating budgets account for approximately 85% of the budget of the French Departments. Yet, it appeared from a careful analysis of these datasets that the total forecast error is mostly driven by the investment revenue forecast error. Though, subsidies granted to the monitoring commissions are part of the operating budgets of the Departments. Therefore, the *shock* variable was defined for each reactor and every year between 2008 and 2015 as the two-year lagged value of the forecast error on Department-level operating revenues. The two-year lag is introduced because real budgets are usually published with a one-year delay. This suggests that forecasting errors made in year Y are only known at the end of year Y+1, and may affect the decision to subsidize local monitoring commissions at year Y+2.

5.2 The channel identification problem

5.2.1 Observable channels

Using the instrumental variable described above, we can now estimate the effect of an exogenous change in monitoring intensity on the observed number of events reported annually by each nuclear reactor. Yet, given the model developed in section 3, we cannot

[17]French local territories publish their budget forecasts and realized budgets every year on this governmental website.

at this stage disentangle the effect of a change in monitoring on safety care (e.g. changes in E_{tot}) from its effect on compliance (e.g. changes in z). This channel identification problem was presented in section 3 in equation (3). To solve this issue, we first need to identify whether an exogenous change in monitoring intensity affects the sanctions perceived by plant managers, or the expected penalty faced when choosing not to report safety events. We then need to disentangle the effect of monitoring intensity on safety care and compliance.

To do this, we first identify within the reports of significant safety events a subset of events which are perfectly detected by managers and necessarily declared to the authority. These events first contain automatic shut-downs of reactors (ASD), which are perfectly detected and declared due to the fact that they instantaneously stop the electrical production of a reactor. They also contain events labelled "abnormal uses of safeguard systems" (SFG). These rare events are considered particularly severe by the safety authority, and are particularly simple to detect during audits.[18] We refer to these events as PDD events (for Perfectly Detected and Declared), and note their total yearly number per reactor E_{PDD}. For this subset of events, we have $z = 1$ and $\rho = 1$. Hence, assuming that perceived sanctions and expected penalties are constant across all types of events, the first order condition defining the optimal effort exerted by the manager to prevent these events is the following:

$$B'(E_{PDD}) = \alpha \tag{4}$$

Therefore, if variations in the intensity of the activity of local commissions lead to a change in the occurrences of these perfectly detected and declared events, then we can conclude that the activity of these commissions do induce a change in the sanction perceived by the plant managers, which led to a change in the level of safety care exerted by the agent. Conversely, the absence of an effect of the monitoring performed by local commissions on the occurrences of PDD events, combined with a significant effect of local monitoring on the unrestricted quantity of events reported to the authority will be consistent with the

[18]These two categories of events were selected jointly with the nuclear safety authority during informal discussions.

hypothesis that local monitoring affects the expected penalty faced by plant managers when failing to report an event.

In addition, we can also assess the effect of monitoring on safety care by estimating the effect of the monitoring performed by local commissions on the reliability of reactors, which is a proxy for the level of safety care exerted by plant managers. We expect reliability to be correlated with safety care since reducing the frequency of these events would limit the likelihood of fortuitous stops of the reactors or would limit the need to extend maintenance periods for safety reasons. Both of these effects would in turn increase reliability. Finally, as reliability is directly related to the profits made by the power station, we argue that a manager provided with incentives to increase safety care will do so, when possible, in a way that increases his profits.

Reliability is here measured by two variables: the annual share of electricity lost due to unplanned maintenance extensions, and the annual share of electricity lost due to fortuitous stops. These proxies capture the quantity of electricity that was not produced due to unplanned maintenance works or fortuitous stops.

5.2.2 Identification

Our identification strategy is based on the combination of the results of the estimation of the effect of local monitoring on three observed variables: the quantity of events reported, the quantity of PDD events reported, and the reliability of power stations. The first estimation, using counts of reported events, provides the general direction of the effect of monitoring on compliance and safety care. If monitoring affects z or E_{tot}, this first estimation provides us with the sign and amplitude of the variations of zE_{tot}. The second estimation, using counts of perfectly detected and declared events, indicates whether local monitoring affects perceived sanctions α. The third estimation, using reliability indicators, provides us with the sign of the effect of monitoring on safety care. We now discuss under what circumstances these estimations allow us to identify the various channels described previously.

First, if monitoring affects the total counts of events zE but does not significantly affect sanctions, or if the effect on reports and the effect on sanctions are not coherent with the first corollary of proposition 3.1, then we can unambiguously claim that monitoring affects expected sanctions. In all other cases, these two estimations are insufficient to identify whether local monitoring affects expected penalties, and only the joint effect of monitoring on compliance and safety care can be identified.

If the effect of monitoring on sanctions and safety care have the same sign, then the third estimation confirms the predictions of our model, and bring no new information. On the other hand, if the sanction channel is closed by the first two estimations, or if the effects of monitoring on sanctions and on safety care have different signs, then it follows that monitoring has to have an effect on expected penalties. The sign of this effect can be determined by the sign of the effect of monitoring on safety care, using proposition 3.1.

Quantitatively, if the third estimation fails to show that monitoring has a significant effect on safety care, then we can neglect the second term of the right-hand side of equation (3). It follows that we can interpret the results of the first estimation as the effect of monitoring on compliance z. To see this, note i the intensity of local monitoring, and β_{budget} the coefficient of the *budget* variable used as a proxy for monitoring intensity in the different estimations. When considering the estimation of the effect of local monitoring on observed reports of events, we have $\beta_{budget} = \frac{\partial zE}{\partial i} = z\frac{\partial E}{\partial i} + E\frac{\partial z}{\partial i}$. Then, if $\frac{\partial E}{\partial i} = 0$, we have that $\beta_{budget} = E\frac{\partial z}{\partial i}$. Unless the sanction channel is convincingly closed by the results of the second estimation, the effect of monitoring intensity on compliance can in general be caused by both a change in perceived sanctions and in expected penalties.

When local monitoring intensity significantly affects expected penalties, such that the resulting effects on safety care and compliance have different signs[19], then the results of the first estimation can still be interpreted quantitatively. Indeed, as $\beta_{budget} = \frac{\partial zE}{\partial i} = z\frac{\partial E}{\partial i} + E\frac{\partial z}{\partial i}$, then we have $\beta_{budget} > z\frac{\partial E}{\partial i}$ and $\beta_{budget} < E\frac{\partial z}{\partial i}$. Hence, observing a positive β_{budget} constitutes a lower bound for the (positive) effect of monitoring intensity on compliance. Likewise, observing a negative β_{budget} constitutes an upper bound for the (negative) effect of monitoring intensity on safety care.

[19]This can arise in multiple situations, for instance if observed reports increase but reliability decreases, or if sanctions are constant but observed reports vary.

5.2.3 Econometric framework

A linear specification allowing to carry out the estimations presented above is:

$$Y_{it} = \beta \cdot X + \beta_{budget} \cdot budgets + \eta_i + \delta_t + \epsilon_{it} \tag{5}$$

where indices i and t respectively refer to the reactor and the year of observation. As was described above, the dependant variable Y_{it} will be defined in turn as the total number of events reported per reactor and per year (ALL), as the number of reports of perfectly detected and declared events (ASD and SFG), and finally as the performance indicators describing the reliability of nuclear stations (K_{ipr} and K_{if}). δ_t represents year fixed-effects, and controls for potentially varying declaration guidelines, or particular time-varying factors, such as generic efforts exerted by EDF at a national scale. η_i represents reactor fixed-effects, which capture potentially varying local factors influencing the safety of nuclear reactors. In all regressions, control variables X include reactor age, electrical production[20] and the overall number of days of maintenance during the year.

Given the endogeneity of the *budget* variable, we estimate equation 5 using our instrument *shock* and a generalized method of moments(GMM-IV) estimator with robust standard-errors.[21] We choose to use a GMM-IV estimator as it is known to be approximatively unbiased and to achieve close to perfect nominal coverage in the just-identified case (see e.g. Angrist and Pischke (2009b,a)). Robustness checks regarding the linear specification of the model and on the nature of the estimator are carried out in Appendix 1.

Estimation results are presented in the following way. We first present in table **??** the four regressions in which the dependant variable Y is defined as, respectively, the annual number of reports (per reactor) of automatic shut-downs ASD, the annual number of reports of unplanned uses of safeguard systems SFG, the rate of lost production due to unplanned prolonged maintenance works K_{ipr} and the rate of lost production due to

[20] Production can be seen as a form of exposure, as all power stations do not produce the same amount of energy each year.

[21] As we only have one endogenous regressor and one instrument, the GMM-IV, two-stage least-square and limited information maximum likelihood estimators are equivalent, see for instance p.189 in Wooldridge (2002), or chapter 8.6 in Hayashi (2000).

fortuitous stops K_{if}. The first two regressions test whether increased monitoring intensity induce a change in the sanctions perceived by plant managers. The last two regressions test whether increases in monitoring intensity lead to changes in safety care.

Second, table 4 contains the results of four regressions. In the first two regressions, the dependant variable Y_{it} is defined as the number of reported significant safety events. In the last two regressions, carried out as a robustness check, Y_{it} is defined and as the logarithm of this quantity. In both cases, regressions differ by the inclusion or exclusion of reactor-fixed effects. These regressions aim to measure the effect of increased monitoring intensity on the overall number of events reported by local managers. The results of these regressions can then be interpreted conditionally on the results obtained in table **??**.

The results of the first-stage regressions are reported in the appendices. These first-stage regressions support our instrument as the coefficient of the *shock* variable is positive and significant. In addition, the test statistics reported in tables 5 and 4 support our instrumental variable.

6 Empirical results

6.1 Monitoring and reporting behaviours

The regressions presented in table 4 show that increased intensity of local monitoring leads nuclear managers to increase significantly the number of safety events declared. At this stage, this increase may be due to either an increase in the total number of occurrences of events, or to an increase in the level of compliance of plant managers with declaration guidelines.

The coefficient associated with the budget of local commissions reported in table 4 is highly significant for the linear specification without fixed-effects, and is only significant at the 10% level of confidence for the linear specification that includes reactor-fixed effects, and for the logarithmic specification without fixed-effects. Results are not robust to the logarithmic specification without fixed effects. In the Appendix, we detail other specifications to which our results are robust (namely GMM with exponential link and control function estimators).

Table 4: Monitoring intensity and reports of significant safety events

	OLS		GMM-IV		
VARIABLES	ALL	ALL	ALL	log(ALL)	log(ALL)
budget	-0.0222	0.0571***	0.132*	0.00353**	0.00863
age	-0.183	0.178	-0.473**	0.0209	-0.0418**
production	-1.642***	-1.137*	-0.942	-0.0835*	-0.0723
maintenance	0.000883	0.0124	0.0135	0.00113	0.00117
Status		9.357***		0.641***	
multiple		-2.355**		-0.172*	
meet		7.513**		0.515**	
SaintLaurent		-1.958		-0.217	
size		1.554**		0.109	
FOAS		-0.481		-0.0552	
FOAK		-0.558		0.0263	
1300.Power_Group		9.191***		0.681***	
1450.Power_Group		20.20***		1.552***	
Constant	30.56***	-12.06		0.650	
Observations	234	234	234	234	234
Fixed effects	No	Y	R-Y	Y	R-Y
R-squared	0.388	0.170	0.157	0.209	0.197
KP rk Wald		26.83	10.08	19.37	10.03
Wu-Hausman		13.46	4.553	6.434	3.162

*** p<0.01, ** p<0.05, * p<0.1. Robust standard-errors

Table 4 also contains the result of a fixed effect OLS regression in which the endogeneity of the budget variable is unaccounted for. In this specification, the coefficient associated with the budget variable is negative, although not significant. This downward bias of the OLS regression is consistent with the reverse causality that may exist between the measurement of local monitoring intensity and the behaviour of local plant managers. Indeed, if local commissions lobby for higher budgets when they expect their local managers to hide information, then we would expect to see high budgets where declarations are rather low (e.g. a negative effect of monitoring on reporting).

The other results provided in table 4 are consistent with intuition. Production and maintenance durations are respectively negatively and positively correlated with reporting, which is consistent with the fact that most events occur during maintenance works. Age has a negative effect on declarations in the fixed-effects regressions, which may be a sign of learning-by-doing in safety. Reactors from the group of 1450 MW and 1300 MW reactors report more events than those in the 900 MW group, which is consistent with the increase in technical complexity of these reactors.

Table 5: First-stage results of GMM-IV regressions

VARIABLES	budget	budget
INSTR	0.298***	0.772***
age	3.332***	5.230***
production	-3.798*	-2.566
maintenance	-0.0647	-0.0659
Status		-13.79
multiple		43.59***
meet		-110.3***
SaintLaurent		-137.2***
size		-27.41***
FOAS		-7.981
FOAK		45.36***
1300.Power_Group		-41.85***
1450.Power_Group		-55.97*
Constant		275.9***
Observations	234	234
Fixed-effects	R-Y	Y

*** p<0.01, ** p<0.05, * p<0.1. Robust standard errors.

Finally, first-stage statistics are reported in table 5, and provide support for our in-

Table 6: Monitoring intensity and perceived sanctions

VARIABLES	ASD	SFG
budget	0.00464	-0.0277**
age	-0.0859*	0.0420*
production	-0.285**	-0.202**
maintenance	-0.00252	-0.00148
Observations	234	234
R-squared	0.076	-0.314
KP rk Wald	10.03	10.03
Wu-Hausman	0.549	8.108

*** p<0.01, ** p<0.05, * p<0.1. Robust standard errors.

strument, whose coefficient in the both first-stage regressions (with and without reactor fixed-effects) is positive and highly significant. The positive sign is coherent with the definition of the instrument: when department fail to forecast accurately their tax revenues, they adapt their spendings accordingly. In other words, departments that receive more budgets than expected also seem to allocate larger budgets to their local monitoring commissions. In addition the KP rank Wald F statistics reported in table 4 are all larger than 10 and thus provide additional support for our instrument. It is to be noticed that in fixed-effects regressions, the value of the Kleibergen-Paap statistics is lower than the Stock-Yogo-15% statistic.

6.2 Monitoring, perceived sanctions and safety care

The two regressions presented in table 6 use counts of perfectly detected and declared events as the dependent variable. The first regression shows that the intensity of local commissions' monitoring activities have no significant impact on the quantity of automatic shut-downs reported by plant managers. On the contrary, the second regression shows a small but statistically significant effect of local monitoring intensity on the occurrences of unplanned uses of safeguard systems.

Conditionally on the assumptions made in the identification section, the results obtained when using reports of ASD as a dependant variable suggest that local monitoring has no impact on perceived sanctions. This result is contrasted by the results obtained when using reports of unplanned safeguard events as a dependant variable, where the

Table 7: Monitoring intensity and reliability

VARIABLES	K_{ipr}	K_{if}
budget	-0.000106	-0.000322
age	-0.00134	0.00122
maintenance	0.000289***	-0.000282**
production	-0.0104*	-0.0158**
Observations	234	234
Fixed effects	R-Y	R-Y
R-squared	0.460	0.091
KP rk Wald	10.03	10.03
Wu-Hausman	0.0499	0.0271

*** p<0.01, ** p<0.05, * p<0.1. Robust standard errors

results indicate a small but significant increase in the sanctions perceived by local managers.

In the two regressions of table 7, we find no significant effect of the commissions' monitoring activity on the on reliability of nuclear reactors. As reliability is assumed to be correlated with safety efforts, these two regressions fail to reject the hypothesis that significant the level of safety care exerted by plant managers remains constant when the intensity of the local monitoring increases. Overall, we conclude from the results gathered in tables 6 and 7 that the intensity of the monitoring performed by local commissions has a small effect on perceived sanctions, but no significant effect on safety care.

These two estimations also show mixed results regarding the age of nuclear reactors, which has a significant positive impact on the declarations of SFG events, but a small negative impact on automatic shut downs. Conversely, the last two regressions show that age has no significant effect on reliability. In all four regressions, these results on age are small when compared to the results on production which significantly decreases the quantity of reports and the reliability of nuclear station. The results on reliability are coherent, given the fact that production and reliability ratios are closely related. The relation between production and declarations of PDD events is also coherent with the fact that significant safety events seldom occur when the power station is producing electricity at its nominal capacity.

6.3 Monitoring and compliance

Given the results presented in tables 4 and 6, observing an increase in the number of events reported can only be consistent with an increase in the expected penalty faced by nuclear plant managers when deciding not to report a safety event. Indeed, under constant expected penalties, the observed increase in perceived sanctions can only be consistent with a decrease in reports. Moreover, under the previous observation that monitoring intensity does not significantly affect safety care, the only way to observe an increase in reports is to have an increase in expected penalties, leading to an increase in compliance.[22]

Given the results presented in table 7, safety care is constant when monitoring intensity increases. As a consequence, the coefficient β_{budget} obtained in the second table can be interpreted as the total effect of monitoring intensity on compliance. From the third column of table 4, we conclude that a budget raise of 1.000€ leads to a 1% increase in compliance with declaration guidelines (assuming that the average number of declaration per reactor and per year is 13 events). As noted in the previous section, if $\frac{\partial E}{\partial i} \neq 0$, then the quantitative interpretation of β_{budget} is still a lower bound on the effect of monitoring intensity on compliance.

6.4 Endogenous detection abilities

These results rely on several assumptions. First, we neglect the fact that all French plant managers work for a single firm, which could provide them with collective incentives that are not captured in our formal model, or accounted for in our empirical estimation framework. Second, we assume that safety care cannot be specific to a particular type of event, and that the ratio of the number of events undetected by the operator to the total number of events occurring in a power station is independent from the level of care exerted by the operator[23].

[22]Indeed, a decrease in expected penalties would lead to a decrease in compliance, which would not be consistent with the observed increase in reports combined with the observed constant level of safety care.

[23]This is different from saying that safety care as no impact on the detection of safety events. In our model, exerting safety care diminishes both the total number of events and the number of events undetected. Yet, we assume that the ratio of these two quantities remains constant.

Assume now that the agent may also choose (at no cost) any detection ability ρ between 0 and 1. Then, all the results derived in section 3 can be applied to a new variable $Z' = z\rho$. Z' captures the *transparency* of nuclear plant managers, defined as the combination of their compliance rate and of their detection abilities.

The identification strategy described above is then valid, provided one replaces all mentions of z by the new variable Z', and substitutes transparency to compliance. Likewise, our empirical results can be reinterpreted by replacing the level of compliance z by the level of transparency Z': plant managers react to the incentives provided by local monitoring, and if they don't seem to increase their level of safety care, they show an increase in transparency, which can be explained by either an increase in detection abilities or by an increase in compliance. The unobserved nature of these two choice parameters makes us unable to distinguish them empirically with available data.

7 Conclusion and policy implications

This paper empirically studies the effect of a French informational policy that organizes the monitoring of nuclear power stations by dedicated local commissions. These commissions can provide no monetary incentives to local plant managers, but can communicate the results of their monitoring activities to local populations and to the safety authority. This paper investigates whether this policy induces changes in the behaviour of nuclear plant managers.

To do so, we design an original empirical strategy to identify the causal impact of this policy on the choices of local managers regarding both safety care and compliance with self-reporting guidelines. We first clarify the incentives faced by managers when exerting safety care and reporting significant safety events: the two main drivers of these decisions are the perceived sanctions incurred for reporting a safety event, and the expected penalty faced when an unreported event is discovered by the safety authority. Using this formal model, we estimate the effect of increased monitoring on safety care and compliance using an instrumental variable method. Our instrument, defined as budget forecasting errors at the Department level, aims to correct the possible biases induced when considering annual budgets of local commissions as a proxy for the intensity of their monitoring activities.

This instrument exhibits several interesting features that qualify it as a quasi-natural experiment.

Empirically, we study the effect of monitoring intensity on three observables: the annual level of reliability of nuclear reactors, the quantity of events declared per reactor and per year, and the quantity of events declared that belong to a certain category of events, characterized by perfect detection and declaration. Using this latter observable, we identify the effect of monitoring intensity on perceived sanctions. Results suggest that the sanctions perceived by plant managers only marginally increase under increased monitoring intensity. Second, we find that increased monitoring intensity leads to an increase in the total number of declarations of safety events, and to no significant changes in the reliability of nuclear reactors. We conclude from these observations that local monitoring significantly increases the expected penalty faced by managers who decide not to comply with reporting schemes, which in turn induces an increase in their compliance with self-reporting criteria. Quantitatively, a budget raise of 2.000 € leads to a one percent increase in the number of events reported.

These results rely on several strong assumptions. First, we neglect the fact that all French plant managers work for a single firm, which could provide them with collective incentives that are not captured in our formal model, or accounted for in our empirical estimation framework. Second, we assume that safety care cannot be specific to a particular type of event, and that the ratio of the number of events undetected by the operator to the total number of events occurring in a power station is independent from the level of care exerted by the operator[24]. When relaxing this assumption, our empirical evidence can still be interpreted as a positive and significant impact on the overall *transparency* of nuclear managers, where transparency is defined as the combination of a manager's ability to detect safety events and his propensity to declare them.

Regarding policy implications, we believe this paper's results call for optimism, as they suggest potentially cost-effective ways of improving the institutional design of nuclear safety regulation. At a national level, there has been a debate in France regarding the

[24]This is different from saying that safety care as no impact on the detection of safety events. In our model, exerting safety care diminishes both the total number of events and the number of events undetected. Yet, we assume that the ratio of these two quantities remains constant.

subsidies given to local monitoring commissions. Indeed, the French law passed in 2006 mentioned that Department ought to provide their local monitoring commissions with a fixed share of the special tax imposed on nuclear installations. As of this writing, this article of the law is yet to be implemented. Our findings suggest that implementing this measure could foster significant improvements in the level of compliance of nuclear operators. More generally, our results suggest that informational policies such as this French local monitoring policy could be used as efficient complements to traditional command-and-control nuclear safety regulation mechanisms, in France or abroad.

Appendix 1: Robustness checks

The results presented in this paper are robust to several other empirical specifications. They are first robust to several other specifications of the GMM estimator. In particular, given that our dependent variables are all necessarily positive, and that our variables ASD, SFG and ALL are count variables, we verify that our results are robust to the following specifications:

- the GMM exponential estimator with Year and Reactor dummies
- the GMM exponential estimator with time-invariant controls

In addition, they are robust to several control function approaches à la Wooldridge (2015). As a reminder, the control function approach is a two step procedures, whose first-step consists in predicting the residuals of a pooled OLS regression of the endogenous explanatory variable on the set of independant variables and excluded instruments. The second step of the control function approach consists in a estimation of a Poisson-QMLE regression where the residuals of the first estimation are included as an explanatory variable. Our results are robust to several such procedures, where the estimators are defined by:

- A pooled OLS first-stage and a Poisson-QMLE or negative binomial (NB2) second stage, including reactor fixed effects and robust standard-errors.
- A pooled OLS first stage including year fixed effects and poisson-QMLE second stage with time-invariant control variables.

However, our results are not robust to the following specifications of the control function approach:

- pooled OLS first stage with year FE (or year and reactor FE) and poisson-QMLE (or NB2) second-stage.

Appendix 2: Proof of proposition 1

Let $\mu(z) = \alpha\rho z + q\beta \int_{z\rho}^{1}(u - z\rho)f(u)du$. From equation (2), one can derive the following first-order condition characterizing the existence of an interior solution for the agent's choice:

$$z^\star = \frac{1}{\rho}F^{-1}(1-\frac{\alpha}{q\beta}) \quad (6)$$

$$B'(E^\star_{tot}) = \mu(z^\star) \quad (7)$$

The effect of a change in α or $q\beta$ on $z^\star$ derives directly from equation (6). Then, using the envelope theorem in equation (7), we can differentiate $\mu(z^\star)$ with respect to either α or $q\beta$, which yields the second part of the result. This result is identical to the results provided by Evans et al. (2009) and Gilpatric et al. (2011).

References

Angrist, J. and Pischke, J. (2009a). A Note on Bias in Just Identified IV with Weak Instruments (referering to Harmless Econometrics: An Empiricist's Companion).

Angrist, J. and Pischke, J. (2009b). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

Bressoux, P., Kramarz, F., and Prost, C. (2009). Teachers' training, class size and students' outcomes: Learning from administrative forecasting mistakes. *Economic Journal*, 119(536):540–561.

Burgherr, P. and Hirshberg, S. (2014). Comparative risk assessment of severe accidents in the energy sector. *Energy Policy*, pages S46–S57.

Davis, L. W. and Wolfram, C. (2012). Deregulation, consolidation, and efficiency: Evidence from US nuclear power. *American Economic Journal: Applied Economics*, 4(4):194–225.

Duflo, E., Greenstone, M., Pande, R., and Ryan, N. (2013). Truth-telling by third-party auditors and the response of polluting firms: Experimental evidence from India. *The Quarterly Journal of Economics*.

Evans, M. F., Gilpatric, S. M., and Liu, L. (2009). Regulation with direct benefits of information disclosure and imperfect monitoring. *Journal of Environmental Economics and Management*, 57(3):284–292.

Feinstein, J. S. (1989). The safety regulation of u.s. nuclear power plants: Violations, inspections and abnormal occurences. *Journal of Political Economy*, 97:115–154.

Gilpatric, S. M., Vossler, C. A., and McKee, M. (2011). Regulatory enforcement with competitive endogenous audit mechanisms. *The RAND Journal of Economics*, 42(2):292–312.

Hausman, C. (2014). Corporate incentives and nuclear safety. *American Economic Journal: Economic Policy*, 6(3):178–206.

Hayashi, F. (2000). *Econometrics*. Princeton University Press.

Hofert, M. and Wüthrich, M. V. (2011). Statistical review of nuclear power accidents. *Asia-Pacific Journal of Risk and Insurance*, 7:1–13.

Lin, L. (2013). Enforcement of pollution levies in china. *Journal of Public Economics*, 98:32–43.

Macho-Stadler, I. and Pérez-Castrillo, D. (2006). Optimal enforcement policy and firms' emissions and compliance with environmental taxes. *Journal of Environmental Economics and Management*.

Rangel, L. E. and Lévêque, F. (2014). How Fukushima Dai-ichi core meltdown changed the probability of nuclear accidents ? *Safety Science*, 64:90–98.

Shimshack, J. P. (2014). The economics of environmental monitoring and enforcement. *Annu. Rev. Resour. Econ.*, 6(1):339–360.

Sovacool, B. K. (2008). The costs of failure: A preliminary assessment of major energy accidents,1907-2007. *Energy Policy*, 36(5):1802–1820.

Telle, K. (2013). Monitoring and enforcement of environmental regulations: Lessons from a natural field experiment in norway. *Journal of Public Economics*, 99:24–34.

Wheatley, S., Sovacool, B. K., and Sornette, D. (2016). Reassessing the safety of nuclear power. *Energy Research & Social Science*, 15:96–100.

Wooldridge, J. M. (2002). *Econometric analysis of cross section and panel data*. MIT press.

Wooldridge, J. M. (2015). Control function methods in applied econometrics. *Journal of Human Resources*, 50(2):420–445.

Zahran, S., Iverson, T., Weiler, S., and Underwood, A. (2014). Evidence that the accuracy of self-reported lead emissions data improved: A puzzle and discussion. *Journal of Risk and Uncertainty*, 49(3):235–257.