

The distortionary effects of incentives in government: Evidence from China's "death ceiling" program

Raymond Fisman, Yongxiang Wang*

Boston University, University of Southern California

This Draft: June, 2016

Abstract

We study a 2004 program designed to motivate Chinese bureaucrats to reduce accidental deaths. Each province received a set of 'death ceilings' that, if exceeded, would impede government officials' promotions. For each category of accidental deaths, we observe a sharp discontinuity in reported deaths at the ceiling, suggestive of manipulation. Provinces with safety incentives for municipal officials experienced larger declines in accidental deaths, suggesting complementarities between incentives at different levels of government. While realized accidental deaths predict the following year's ceiling, we observe no evidence that provinces manipulate deaths upward to avoid ratchet effects in the setting of death ceilings.

JEL classification: D73; H75

Keywords: Bureaucracy; Incentive design; Workplace Safety; Regression discontinuity

*Raymond Fisman: Boston University, Room 304A, Boston University, Boston, MA 02215; email: rfishman@bu.edu; Yongxiang Wang: Marshall School of Business, University of Southern California, Los Angeles, California 90089; email: Yongxiang.Wang@marshall.usc.edu. We would like to thank Eric Hardy, Xingchen Ji and Weicen Zhang for excellent RA work. Wang also thanks the China National Science Foundation (Grant Number: 71472180) and Australian Research Council (DP150102339) for financial support.

1 Introduction

The tradeoff between motivating agents and distorting their efforts is central to incentive design. As emphasized in Holmstrom's (1979) classic model, this tradeoff is starker for difficult to observe objectives, or those subject to manipulation. In these cases weak or even flat rewards may be optimal. Holmstrom's insight in turn has implications for the design and assignment of tasks within an organization (Holmstrom and Milgrom, 1991), and even the assignment of tasks across organizational types: when viewed through the lens of incentive design, non-profits or government may be seen as a contractual solution to the problem of imperfect observability (see, for example, Easley and O'Hara (1984) and Dixit (1997)).

Against this backdrop, there has been an active debate - both in academic and policy circles - on the appropriateness of using performance incentives in bureaucracies. Such practices are common: teachers are assessed based on standardized test-based value-added scores (Springer et al, 2014), FBI agents on the production of intelligence leads (Gulati et al, 2009), and police officers on arrest and crime data (Eterno and Silverman, 2012). In each of these settings, much of the debate has centered around the extent to which outcome metrics are manipulable.

We study the effect of incentives on the manipulation of reported outcome measures in the context of a high-stakes safety targeting program implemented by the Chinese central government. Our primary contribution is to show that, even under conditions where ex ante one would expect that gaming of the outcome measure would be difficult or impossible, there is still widespread manipulation. This naturally raises concerns about the ever wider application of incentives in bureaucracies.

In February 2004, the State Administration of Work Safety (SAWS), the national safety regulatory authority, announced a set of province-level safety targets, or 'death ceilings,' across a range of causes of accidental death, mostly workplace related.¹ SAWS' head at the time, Wang Xianzheng, described the measure as a response to a worrying increase in severe (and well-publicized) accidents in the preceding months. The new rules targeted a reduction in accidental deaths of 2.5 percent for 2004, with further targeted reductions expected in future years. Provinces were required to develop their own plans for getting accidental deaths below the mandated ceilings, and to provide quarterly reports on workplace deaths for a number of industries (e.g., coal mining; agriculture; manufacturing), and also total deaths from traffic accidents or fire.

SAWS publicized provinces' quarterly death numbers by category, along with their year-end ceilings, via news releases that were in turn published in the national newspaper, the *People's Daily*. At least superficially, these reforms were a resounding success: deaths across

¹See http://www.chinadaily.com.cn/english/doc/2004-02/20/content_307974.htm accessed on June 2, 2016

all categories fell by about 45 percent between 2005 and 2012, an annual decline of nearly 8 percent.

When we delve more deeply into the public data, however, we show that it was very likely that manipulation of death statistics played a significant role in the reduction. Using McCrary’s (2008) test for manipulation of a running variable applied to the distributions of year-end deaths-to-ceiling ratios in each category, we find a clear discontinuous drop in the density function around one, an indication of ‘sorting’ of reported deaths around the ceiling.² By contrast, we observe no discontinuity for cumulative deaths through the first three quarters of the year.

We next examine cross-province heterogeneity in hitting safety targets, in particular examining the effects of “No safety, no promotion” (NSNP) laws for local regulators. As noted above, the 2004 SAWS directive gave considerable discretion to individual provinces in improving safety performance. Starting in 2005, provinces established rules for disaggregating province-wide death ceilings into municipal-level ceilings, and some went as far as tying local officials’ promotion to hitting their (local) targets. We show that the provision of local incentives via NSNP was effective in reducing province-level deaths (though this relationship is significant only at the 10 percent level). This is consistent with a complementarity between local and province-level governments in achieving safety targets.

If reducing deaths below government-mandated ceilings comes in part through manipulation, it raises the question of why we observe any discontinuity at all: why not manipulate reported deaths *well* below the ceiling? There are a number of potential explanations. First, death misclassification may be costly, and plausibly convex in the extent of manipulation, an issue we discuss in Section 2.2. Further, if ceilings are set as a function of past reported deaths, reducing deaths by “too much” could lead to a ratchet effect with ever-lower ceilings as a result of performing “too well” in the present year. As evidence that ratchet effects would present a concern to officials, we show that year-end deaths are an extremely strong predictor of the following year’s ceiling, even after accounting for lagged ceilings. Despite this, we do not find any evidence that officials went so far as to manipulate reported deaths *upward* to secure more favorable ceilings in future years, based on a comparison of pre- versus post-2004 growth rates in industrial workplace deaths. The bottom quartiles of the pre-2004 and post-2004 distributions are virtually identical, so that provinces were just as likely to have very low reported deaths before and after the advent of death ceilings. The distributions diverge primarily at positive growth rates, consistent with strong pressures to report declines in accidental deaths in the post-2004 period.

We conclude our overview of the results by emphasizing that we cannot rule out a real

²For severe accidents and fire deaths, we observe relatively modest discontinuities. We argue these categories may be less subject to manipulation and/or control.

response from the imposition of death ceilings. While the sharp discontinuities around the ceiling that we document in our main analysis are most easily reconciled with responses to safety incentives via manipulation rather than real safety improvements, we cannot distinguish whether safety improvements occurred as well. The clustering in reported deaths we observe around 90 to 100 percent of the ceiling could be the result of a real drop in high-mortality outcomes (accompanied by manipulation to ensure the final figure remains below the ceiling). It could alternatively be reconciled with a pure manipulation story with an audit probability that is greater just below the ceiling. In the absence of reliable data on industry output or safety expenditures, we cannot fully rule out these types of alternative explanations.

Our paper contributes to the large body of work within economics on the consequences of high-powered incentives, and the challenges of incentive design more broadly. While this work has primarily involved for-profits (see, for example, Oyer (1988) for a classic study on the topic), we examine the increasingly common phenomenon of high-powered incentives in governmental and/or non-profit settings. Recent work in economics has considered specifically how these incentives lead to manipulation of reported outcomes, in China and elsewhere. Most notably, Young (2003) argued that China's extraordinary GDP growth during the last two decades of the twentieth century was the result of distortions in reporting.³ Two recent papers study manipulation of outcomes around a threshold to study agents' responses to incentives, using approaches similar to ours. Duflo et al (2013) analyze the manipulation of pollution data by third-party auditors of industrial plants in Gujarat, while Dee et al (2011) study whether New York City teachers alter students' Regents' exam scores (the authors attribute manipulation to teachers' desires to help their students rather than concern for their own evaluations).

The results we present have particular relevance for the literature on the tradeoffs associated with decentralization. As Mookherjee (2006) emphasizes in his survey on decentralization in hierarchies, central administrators may have limited knowledge of local conditions, and their interests may diverge from those of local officials. This may lead to moral hazard that undercuts the benefits of local knowledge. Our study brings this tradeoff into sharp relief in an empirically important setting, given its distorting effects on reporting behavior. A further benefit of decentralization emphasized by more recent research is that it allows higher-level policymakers to assess bureaucrats' quality (Mookherjee, 2015). Our results suggest, however, that the tournament model of government promotion may be undermined if performance metrics are subject to manipulation. In our context reported safety outcomes may still contain some information - as in the earnings management literature, a manager only misses the mandated target if his organization is in truly terrible shape (Burgstahler and Dichev, 1997) - but it does not provide the fine-grained signal that would be obtained in

³Others have strongly disputed these claims. See Chow (2006) and Holz (2014).

the absence of manipulation.

Finally, our findings also may serve as a counterweight to the trend toward increased focus on evaluation based on observable performance metrics, providing a cautionary tale on the complications associated with the use of metrics-based evaluation and incentives for hard-to-measure outcomes that are subject to manipulation. The outcome we focus on here - accidental deaths - is one that is plausibly less subject to manipulation than many other outcomes of social interest.

In the next section, we provide background information on the safety reforms of 2004 and an overview of the data. In section 3 we present our main results. We provide concluding comments in section 4.

2 Background and Data

2.1 The death ceiling system in China

On January 9, 2004, the State Council of China, the country's highest administrative authority, issued a document whose title translates roughly as *Decision of the State Council on Further Enhancing Work Safety*.⁴ In addition to re-emphasizing the importance of workplace safety, the decision stated as an objective that by 2020, China's safety record should be comparable to those of other middle-income countries. The document also specified that safety targets would need to be set at the national and subnational levels, starting in 2004. Finally, the decision specified that beginning in 2004, provincial governments would have to provide quarterly reports on accidental fatalities to the State Administration of Worker Safety (SAWS), which in turn would publicize these reports. (SAWS is the country's workplace safety regulator, reporting directly to the State Council.)

As a response to the State Council decision, the Work Safety Commission of the State Council issued a further directive, *Opinions on Setting up a Death Ceiling System*, which provided a specific target of 2.5 percent for accidental death reductions nationally.⁵ Working with SAWS, the Work Safety Commission further set safety targets for each province in China, providing ceilings for a variety of industries, including: Agriculture; Chemical production; Coal mining; Construction; Fireworks; Railway; and Industrial, Non-coal mining, commercial, and trade (IMCT). Additionally, ceilings were assigned by severity of workplace accident,

⁴The *Decision*'s text may be found at <http://www.china.com.cn/chinese/PI-c/483862.htm> (accessed on June 2, 2016).

⁵The Work Safety Commission is a unit of the State Council. The Commission's director is typically a vice-premier of China, and the chief officer of SAWS serves a vice-director. Other non-director members are vice-ministers of other central government departments. The text of the Commission's opinion may be accessed at <http://www.safetyhome.org/news/displaynews.asp?id=1130> (accessed on June 2, 2016).

including Severe (3 or more deaths) and Very Severe (10 or more deaths). Two final categories - which included both workplace and non-work fatalities in their totals - were also assigned ceilings: Road and Fire. We conjecture that these were included because they are more visible and salient to the general population than, say, fatalities due to poisoning or falls, which are nonetheless the cause of a high proportion of accidental fatalities. Finally, a ceiling was assigned for the overall number of deaths across all industries and categories (Overall).

In practice, figures for a number of less prominent categories went unreported, and in what follows we focus on the year-end ceilings in the eight categories where at least half of province-year death totals and ceilings were available, including Overall; Agriculture; Coal mining; Fire; Road; Railway; IMCT; and Severe. Given the non-independence of these categories (e.g., many severe accidents take place in coal mining, and all subcategories sum to the overall figure) and potential serial correlation, we will allow for province-level clustering of standard errors.

Starting in the second quarter of 2004, SAWS collected reported deaths and death ceilings on a quarterly basis, which were then published in the *People's Daily*. In the program's first year, most categories went unreported, so we begin our sample in 2005. SAWS stopped publishing reported deaths and death ceilings at the end of 2011, and the figures were not reported in the first and third quarters of 2007 and 2008, as well as the first half of 2005. We requested information on reported deaths and ceilings from SAWS for the dates that were not publicly available, and received supplemental data for 2007 and 2008, as well as the fourth quarter of 2012.⁶ While the target system remained in place at the time of writing, data were no longer reported publicly after 2012. Thus, our data include quarterly observations from 2005 through to the end of 2011, as well as the year-end total for 2012.

As part of their efforts at meeting province-level ceilings, some provinces adopted "No safety, no promotion" (NSNP) policies that made promotion of local government officials contingent on meeting the death ceilings set for their region by the provincial government. Guangdong was the first to adopt this policy in 2005, while others followed suit only several years later. As of June 2015, 20 provincial authorities, out of a total of 31, had adopted NSNP policies. Note that NSNP applies to local officials - that is, it is set by *provincial* officers to evaluate *local* officials. We list the dates of the passage and implementation of NSNP policies for each province in online Appendix Table A1, which shows that a number of NSNP laws were passed during our sample period of 2005 - 2012. This will allow us to exploit within-province variation in identifying the effects of NSNP on accidental deaths.

⁶SAWS stated that, since figures were not published in the first two quarters of 2005 and 2012, they could not provide the information to us, and we cannot infer from their response whether they simply do not have the information, or have it but are unwilling to share it. For the third quarter of 2012, SAWS did not collect accidental death information from provincial governments, because of the 18th National Congress of the Communist Party of China.

2.2 Anecdotal evidence of gaming of deaths

From the outset, the death ceiling system was controversial, with manipulation of reported figures thought to be widespread. An in-depth report published in the *Southern Weekly* (Hao, 2007), for example, described one common (albeit costly) means of manipulating road deaths: If an individual was severely wounded in a traffic accident, it would be classified as a road accident (and count toward the road ceiling) only if the individual died within seven days. If the death occurred eight or more days after the accident took place, it would be classified as a natural death. This created an incentive to keep accident victims alive for at least a week, especially when close to the year-end ceiling.

For workplace deaths, problems around ambiguities in death classifications led SAWS to issue a *Note on problems in identifying workplace accidents* on October 9, 2007.⁷ The note's issuance provides evidence in itself of the inevitable discretion involved in classifying deaths. It clarifies, for example, that accidental deaths at unlicensed firms should be classified as workplace deaths. However, as pointed out by Su (2011), the note had no legal authority, and the various liberties in classifying deaths that it describes likely continued after its publication.

Deaths were also thought to be under-reported by the local governments responsible for providing these figures to provincial officials (Shanghai Law Journal, 2007), following the Chinese adage that translates as "Village tells a lie to town, town then tells another lie to county, until the lie reaches the state council."⁸ As we argue below, the imposition of "No safety, no promotion" laws may have helped to better align the interests of officials at different levels of government, thus improving their ability to collude on reported deaths.

2.3 Data

We hand collected province-level quarterly reported deaths and corresponding death ceiling data for each category from the *People's Daily* for 2005 - 2012. For most categories in most years, we have data for all quarters except, as noted in Section 2.1 above, the first two quarters of 2005 and the first three quarters of 2012.

Reported deaths in category c in province p and year y is denoted by $Deaths_{cpy}$. For quarterly data we add an appropriate suffix to the variable. For example, $Deaths_{Q3cpy}$ denotes *cumulative* deaths through to the end of the third quarter. We similarly append the relevant suffix when a set of analyses focus on a specific category (e.g., $Deaths_{IMCT_{py}}$).

Our main interest will be in analyzing the relationship between reported deaths and

⁷The text is available at <http://wenku.baidu.com/view/aff5de48e518964bcf847c18.html> (accessed on June 2, 2016).

⁸See "Should death ceiling system be established?" (Siwang Zhibiao Gaibugai Sheli?), *Shanghai Law Journal*, May 30, 2007.

mandated death ceilings, denoted by $Ceiling_{cpy}$, at the province level. Our main outcome of interest is the deaths-to-ceiling ratio, $Deaths_{cpy}/Ceiling_{cpy}$, where officials faced an incentive to keep this ratio at less than or equal to one. In Online Appendix Table A2, we provide a matrix showing, by year, the categories where year-end data on $Deaths_{cpy}/Ceiling_{cpy}$ were available.

Table 1 provides summary statistics on annual deaths, $Deaths_{cpy}$, as well as the deaths-to-ceiling ratio, $Deaths_{cpy}/Ceiling_{cpy}$. Road deaths constitute nearly 80 percent of deaths in the Overall category. Given that Road and Fire are the only non-workplace categories included in Overall, a ratio of 80 percent is in line with comparable figures from other countries. (For example, the ratio of traffic deaths to workplace deaths in the U.S. is about eight to one.⁹)

The figures in Table 1, when put in per capita terms, also suggest that under-reporting was likely a serious problem even in the earliest years of the death ceiling program. For example, there were 98,738 reported Road deaths in 2005, or about 7.6 per 100,000 population, about half the traffic fatality rate in the United States in that year.¹⁰ Such differences are one reason that public health researchers question the veracity of China’s safety figures. A 2011 article in *The Lancet* argued that the problem of misrepresenting traffic fatality figures has increased since 2004, precisely the period we consider in our analysis (Alcorn, 2011). Aggregating across all categories of workplace fatalities, the official figures in Table 1 imply a Chinese workplace death rate in 2005 of 1.8 per 100,000 population, virtually identical to the U.S. rate.¹¹ Given the two countries’ very different levels of development, this is again suggestive of large-scale under-reporting.

We next define an indicator variable $NSNP_{py}$ to denote the province-year observations when “No safety, no promotion” regulations were in place; it takes on a value of one starting in the year following the passage of NSNP regulations (though our results are not sensitive to this treatment of the year of implementation).

For most accidental death categories, data are unavailable outside the public reporting window of 2005 - 2012. However, deaths in the IMCT category are available through a variety of sources for 1993 - 2003, excluding 1998. The 1993 - 1997 IMCT death data were collected from the China Labor Yearbooks; 1999 from the China Economics and Trade Yearbook; and 2000 - 2003 from China’s Worker Safety Yearbooks.¹² These data will allow us to compare

⁹See <http://www.cdc.gov/nchs/fastats/accidental-injury.htm> accessed on June 2, 2016.

¹⁰See <http://www.iihs.org/iihs/topics/t/general-statistics/fatalityfacts/state-by-state-overview/2005> accessed June 2, 2016.

¹¹See http://www.bls.gov/iif/osh_nwrl.htm#choi accessed June 2, 2016

¹²1998 is unavailable because it is a transitional year, during which the China Labor Yearbook stopped providing these data, while the China Economics and Trade Yearbook had not yet started to report them. Also note that the province of Chongqing was classified as a municipality of Sichuan Province prior to 1997, so we have one observation less for these early years.

the distributions of changes in fatalities pre- and post-2004 for the IMCT category.

Finally, province-year level economic indicators such as GDP per capita are drawn from the *Regional Economy Database* at GTA, a Shenzhen-based data vendor which provides finance and economics databases on Chinese listed firms and the Chinese economy. We also obtained the real GDP growth targets that each province set for itself via government annual reports. For each province these are available online. We combine the GDP growth targets with realized real GDP data ($GDP_{Growth_{py}}$) to define $GDP_{py}/GDPTarget_{py}$, the ratio of actual to targeted (real) GDP. This ratio provides an output-based analog to our measure of safety compliance (though for economic output a high value indicates greater compliance).

3 Results

3.1 Discontinuities around the targeted number of deaths

We begin with a series of graphs showing of the distribution of $Deaths_{cpy}/Ceiling_{cpy}$, disaggregated by death category c , and test for “manipulation” around the SAWS-mandated ceiling (recall that p indexes province and y indexes year). We use McCrary’s (2008) test for a discontinuity in the distribution of $Deaths_{cpy}/Ceiling_{cpy}$ around a value of one, the threshold for meeting the safety target. McCrary’s test involves estimating local linear regressions of the density on either side of the cutoff (in our case, around one), and using a Wald test to check whether the predicted number of observations at the cutoff is the same for the two local linear regressions. If the predicted frequencies are different, it suggests that sorting around the cutoff is non-random.¹³

In Figure 1, we show a histogram of $Deaths_{cpy}/Ceiling_{cpy}$ for overall province-level accidental deaths, with a bin width of 0.077, chosen to match the bin width that will be used by McCrary’s density test. A discontinuity is clearly discernable at one. Figure 2 presents the McCrary density test, providing the local linear regressions on either side of 1 (as noted previously with a bin width of 0.0077). The gap between the density estimates on either side is -2.72, and a z-statistic of the difference of -4.23 (p-value < 0.001). In Online Appendix Figure A1, we provide graphs of the local linear regressions and discontinuity estimates for $Deaths_{cpy}/Ceiling_{cpy}$ for all other death categories. In every case there is a statistically significant discontinuity at one. Notably, the discontinuity is smallest for Severe accidents (workplace accidents that lead to three or more fatalities) and Fire, which are plausibly less amenable to reclassification or concealment¹⁴

¹³This test has been employed primarily in research using regression discontinuity methods, where it is used to assess whether the distribution of the running variable is smooth around the cutoff, which is necessary for the RD design to be valid. The Stata code to implement the test is available at <http://eml.berkeley.edu/~jmccrary/> accessed on June 2, 2016.

¹⁴There are certainly examples that have come to light of attempted suppression of news of high-fatality

In Online Appendix Figure A2 we show histograms for each category of the distribution of $Deaths_Q3_{cpy}/Ceiling_{cpy}$, i.e., reported accidental deaths at the end of the third quarter as a fraction of the year-end ceiling. In contrast to the year-end reported death ratios, the third-quarter distributions are “smooth.” We can, for example, reject the presence of a discontinuity at 0.75 at least at the 10 percent level for all categories except coal, where we obtain a p-value of the difference at the cutoff of 0.028 - an even sharper discontinuity occurs at a value of 1 for coal, however, as many observations are already close to the year-end threshold by the end of Q3. Thus, the discontinuity in the distribution of accidental deaths emerges only in the final quarter of each year.

3.2 Effect of local incentives

We next assess whether provinces are better able to improve their safety records (and meet their safety targets) after implementing “No safety, no promotion” (NSNP) incentives for local officials. We do so by comparing death rates and probabilities of meeting safety targets, pre- versus post-NSNP. Owing to the relative sparseness of the data, we perform this analysis for the pooled sample of all accidental death categories.¹⁵ We use a specification that controls for category-year and category-province effects, as well as the extent to which a province achieved its targeted level of economic growth:

$$I(Deaths_{cpy} > Ceiling_{cpy}) = \beta_1 * NSNP_{py} + \beta_2 * GDP_{py} / GDPTarget_{py} + \mu_{cy} + v_{cp} + \epsilon_{cyp} \quad (1)$$

The outcome variable, $I(Deaths_{cpy} > Ceiling_{cpy})$, is an indicator variable denoting whether deaths in category-province cp exceeded the ceiling in year y ; in all specifications we cluster standard errors at the province level. This addresses potential within province-category autocorrelation, as well as the covariance within each province-year grouping that results from the non-independence across categories.

In column 1 of Table 2, we show the results of specification 1, including only year, province, and category fixed effects. We find a weak negative relationship (p-value of 0.21) between $NSNP$ and the deaths-to-ceiling ratio. While we present results for the full sample, we are most interested in provinces that implement NSNP rules during our sample period, which allows for within-province variation in the province-year variable $NSNP_{py}$. In column 2

accidents in coal mining. See, for example, “Graft in China Covers Up Toll of Coal Mines,” (*New York Times*, April 10, 2009) for a discussion of a company’s attempt to coerce victims’ families into maintaining silence about a multi-fatality mine collapse. Yet the very fact that these were reported on in the media is some indication of the difficulty in concealing such information.

¹⁵For the high-death categories of overall accidental deaths and road deaths, there is not a single case that $Deaths_{cpy} > Target_{cyp}$ in province-year observations with $NSNP = 1$, though there are only 3 “failures” (of 66 data points) for $NSNP = 0$ observations, so it is hard to draw strong conclusions from this result.

we limit the analysis to provinces that implemented *NSNP* incentives during 2005 - 2011, which we refer to as the $Var(NSNP) > 0$ sample. The coefficient on *NSNP* doubles in magnitude to -0.089, though it is not significant at conventional levels (p-value of 0.11). In column 3, we include $GDP_{py}/GDPTarget_{py}$, which controls for the extent to which a province has met its GDP growth target. The coefficient on $NSNP_{py}$ increases to -0.098 (p-value 0.09); interestingly, the coefficient on $GDP_{py}/GDPTarget_{py}$ is negative, which suggests that provinces which meet (self-imposed) GDP growth targets also tend to meet the central government's safety targets. In column 4 we add year \times industry and province \times industry fixed effects, which has very little impact on the estimated coefficients. While the coefficient on $NSNP_{py}$ is noisily measured across all specifications, its economic magnitude is large, indicating that *NSNP* laws are associated with a 10 percentage point decrease in the likelihood of exceeding the year-end ceiling. Given the average probability that $Deaths_{cpy}/Ceiling_{cpy} > 0$ of 0.19 for $NSNP = 0$ province-year observations, this implies a 50 percent reduction in the likelihood of exceeding the ceiling as a result of *NSNP* incentives. (Our regression results are roughly in line with the simple pre-versus-post difference in means for provinces that implement *NSNP* incentives during our sample period, of 0.19 versus 0.12.)¹⁶

In Figure 3 we show the cumulative density functions for $Deaths_{cpy}/Ceiling_{cpy}$ across all categories, disaggregated by *NSNP* status. Interestingly, for the portion of the distribution below $Deaths_{cpy}/Ceiling_{cpy} = 1$, the cdf for $NSNP = 1$ province-year observations is virtually identical to that of the cdf for $NSNP = 0$ observations, before diverging just below one. This suggests that *NSNP* incentives reduced deaths primarily by lowering the likelihood of exceeding safety targets - conditional on meeting the target, the two distributions are quite similar.

3.3 Death targets and ratchet effects

If manipulation is responsible for much of the discontinuity around safety targets, it naturally raises the question of why bureaucrats do not manipulate figures downward to an even greater extent. Convex manipulation costs and, relatedly, higher audit probabilities for suspiciously low values, are possible reasons that reported deaths might not be manipulated far below the ceiling. There must be at least some cost to manipulation under some circumstances, else provinces would never exceed their ceilings at all (though we have no record of government investigations related to manipulation beyond the non-binding guidelines reported in Section 2.2). If officials wish to avoid attention or scrutiny, there would again be

¹⁶If we employ the continuous measure of safety compliance, $Deaths_{cpy}/Ceiling_{cpy}$, the estimated coefficients on *NSNP* are all negative, though none is significant at the 10 percent level. This is consistent with *NSNP* acting first and foremost to motivate provinces to get deaths figures below their ceilings, but the data are far too noisy to distinguish between a general decline in deaths and a drop only below the ceiling.

at least some check against extreme manipulation.

A further explanation may be due to concerns over a ratchet effect, where good performance - beating the death ceiling by a wide margin - induces a harsher (i.e., lower) ceiling the following year. Ratchet effects may appear in any setting where next period's target is a function of this period's performance. In Figure 4, we show the raw relationship between $\log(Deaths_{cpy-1})$ and $\log(Ceiling_{cpy})$. Observations line up on or near the 45 degree line. (The very strong one-to-one relationship persists even if we remove province, category, and year effects.) The same pattern could be generated by persistence in target-setting, for example, given the high concordance between targeted and reported year-end deaths. In Table 3, we attempt to distinguish between these two possibilities by including both $\log(Deaths_{cpy-1})$ and $\log(Ceiling_{cpy-1})$ as regressors in predicting the current target $\log(Ceiling_{cpy})$. As shown in columns 1 and 2, each variable is highly significant when included on its own (p-value < 0.001). When both are included together in column 3, the coefficient on $\log(Ceiling_{cpy-1})$ declines by about 85 percent, whereas the coefficient on $\log(Deaths_{cpy-1})$ is virtually unchanged (though its standard error increases somewhat, the p-value remains below 0.001). In column 4 we allow for the interaction of lagged reported deaths and $I(Deaths_{cpy-1} > Ceiling_{cpy-1})$, to account for the possibility that bureaucrats, recognizing concerns over ratchet effects, untether the link between reported deaths and next year's ceiling at lower deaths-to-ceiling ratios. The interaction term is small in magnitude (-0.006) and statistically insignificant; we thus find no evidence that the relationship between lagged deaths and current ceiling changes as lagged deaths move below the ceiling.

If the static impact of death ceilings is a downward manipulation of reported deaths (as documented in the first part of our paper), dynamic concerns may create incentives for upward manipulation in order to secure a more favorable ceiling the following year.¹⁷ We attempt to shed some light on the question of upward manipulation by comparing the pre- and post-2004 distributions of accidental deaths. As we note in the Section 2.3, the only category where pre-2004 death figures are available is IMCT. If upward manipulation takes place to secure favorable future targets, we expect that the accidental death growth rate, $\% \Delta Deaths_{cpy} = \frac{Deaths_{cyp}}{Deaths_{cpy}} - 1$, would have relatively few low values for post-2004 data, as compared to the pre-2004 data. In Figure 5, we show the cumulative density functions for the pre- and post-2004 distributions of $\% \Delta Deaths$. Interestingly, the two functions are near-identical over the bottom quartile of their distributions; if death figures had been inflated as a result of ratchet effects, we would have expected the post-2004 distribution to be shifted to the right at low levels. We thus find no obvious indication in the data of upward manipulation in deaths in response to ratchet effect concerns. At about the 25th percentile, just as

¹⁷In contrast to downward manipulation, we know of no clearcut anecdote of fabricated deaths to artificially raise the accidental death count closer to the ceiling.

the growth rate approaches -10 percent, the lines diverge sharply, a natural effect of death ceilings being set relatively close to lagged values of *Deaths* (and hence the ceiling constraint starting to bind at this level).

4 Conclusion

We examine the consequences of a high-stakes incentive program that aimed to improve safety practices in China. We argue that the sharp discontinuities in the deaths-to-ceiling distributions across all accident categories are strongly suggestive of a downward manipulation of reported deaths.

Our findings may be seen as a cautionary tale on the use of incentives based on performance metrics in bureaucracies. However, in the absence data on actual accidental deaths or credible data on industrial production, there are several caveats in considering the policy implications of our results. First, it is difficult to assess whether the incentive program we study had *any* real effect on safety (as distinct from the effect on manipulation) or economic activity. This limits our ability to make a general assessment of the welfare effects of the program. Nor can we say definitively whether the program was suboptimal in its design: that would require a full-blown cost-benefit assessment that included assumptions about what safety might have been like under alternative regimes.

We note finally that it is not clear whether the death ceiling program failed to achieve its objective, from the perspective of the implementing authority, as that would require a further set of assumptions about the China State Council's objective function. China's central government is plausibly interested in dissent minimization. If this is accomplished by reporting year-on-year death reductions across all accident categories, regardless of realities on the ground, then performance improvements through manipulation may be a perfectly acceptable result.

These caveats nonetheless suggest a number of ways forward. First, more detailed data or a different setting may allow for a more fine-grained distinction between manipulation and effort as the source of reported improvements. More importantly, we may be able to use the choice of incentives to better understand the functioning and objectives of governments themselves.

References

- [1] Alcorn, Ted. "Uncertainty clouds China's road-traffic fatality data." *The Lancet* 378.9788 (2011): 305-306.
- [2] Baker, George P., Michael C. Jensen, and Kevin J. Murphy. "Compensation and incentives: Practice vs. theory." *The Journal of Finance* 43.3 (1988): 593-616.
- [3] Burgstahler, David, and Ilia Dichev. "Earnings management to avoid earnings decreases and losses." *Journal of accounting and economics* 24.1 (1997): 99-126.
- [4] Chow, Gregory. "Are Chinese official statistics reliable?." *CESifo Economic Studies* 52.2 (2006): 396-414.
- [5] Dee, Thomas , Brian Jacob, Jonah Rockoff, and Justin McCrary. "Rules and discretion in the evaluation of students and schools: The case of the New York regents examinations." *Columbia Business School Research Paper* (2011).
- [6] Dixit, Avinash. "Power of incentives in private versus public organizations." *The American Economic Review* (1997): 378-382.
- [7] Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan. "Truth-telling by Third-party Auditors and the Response of Polluting Firms: Experimental Evidence from India." *The Quarterly Journal of Economics* (2013)
- [8] Easley, David, and Maureen O'Hara. "The economic role of the nonprofit firm." *The Bell Journal of Economics* (1983): 531-538.
- [9] Eterno, John A., and Eli B. Silverman. *The crime numbers game: Management by manipulation*. CRC Press, 2012.
- [10] Goodhart, Charles. *Problems of monetary management: the UK experience*. 1975.
- [11] Holmstrom, Bengt. "Moral hazard and observability." *The Bell Journal of Economics* (1979): 74-91.
- [12] Holmstrom, Bengt, and Paul Milgrom. "Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design." *Journal of Law, Economics, & Organization* (1991): 24-52.
- [13] Carsten, A.H. "Deconstructing China's GDP statistics." *China Economic Review* 15.2 (2004): 164-202.
- [14] Hao, Jiawei. "Decoding Death Targets in China." *Southern Weekly* (May 4th, 2007).

- [15] Mookherjee, Dilip. "Decentralization, hierarchies, and incentives: A mechanism design perspective." *Journal of Economic Literature* (2006): 367-390.
- [16] Mookherjee, Dilip. "Political Decentralization." *Annual Review of Economics* (2015). Vol. 7: 231-249.
- [17] McCrary, Justin. "Manipulation of the running variable in the regression discontinuity design: A density test." *Journal of Econometrics* 142.2 (2008): 698-714.
- [18] Oyer, Paul. "Fiscal year ends and nonlinear incentive contracts: The effect on business seasonality." *Quarterly Journal of Economics* (1998): 149-185.
- [19] Rivkin, Jan W., Michael Roberto, and Ranjay Gulati. "Federal Bureau of Investigation, 2009." (2009).
- [20] Su Lin, "Death Ceilings and the dilemma of local safety regulations in China." Working Paper (2011).
- [21] Springer, M. G., Ballou, D., Hamilton, L., Le, V. N., Lockwood, J. R., McCaffrey, D. F., and Stecher, B. M. (2011). *Teacher Pay for Performance: Experimental Evidence from the Project on Incentives in Teaching (POINT)*. Society for Research on Educational Effectiveness.
- [22] *World at Work, Incentive Pay Practices Survey: Non-Profit/ Government Organizations, 2014*, mimeo
- [23] Young, Alwyn. "Gold into base metals: productivity growth in the People's Republic of China during the Reform Period." *Journal of Political Economy* 111.1 (2003): 1220-1261.

Table 1: Deaths at the province-year level by category, 2005-2012

	Mean	StdDev	Min	Max	Obs	Mean	StdDev	Min	Max	Obs
	<i>Deaths_{cpy}</i>					<i>Deaths_{cpy}/Ceiling_{cpy}</i>				
All	2960.6	1939.36	353	11298	248	.94	.06	.72	1.37	248
Agriculture	13.08	15.4	0	68	165	.72	.4	0	2	164
Coal	125.48	143.4	0	837	207	.82	.37	0	2	208
Fire	42.51	39.66	1	298	244	.94	.45	.1	2	244
IMCT	391.15	242.27	10	1192	248	.94	.14	.22	1.38	248
Railway	104.74	122.38	1	897	238	.8	.29	0	2	241
Severe Accidents	54.28	28.46	10	152	185	.97	.23	.54	2	185
Road	2280.38	1816.07	119	9959	217	.93	.07	.6	1.47	217

Notes: The first set of columns provide summary statistics on year-end accidental deaths in each listed category at the province-level. The second set of columns provide summary statistics on the year-end ratio of reported deaths to the government mandated death ceiling. The IMCT category includes workplace deaths in industrials, non-coal mining, commercial, and trade; severe accidents are deaths in accidents that involve three or more fatalities. Other death categories are self-explanatory (see text for details).

Table 2: No safety, no promotion laws and accidental deaths

Dependent Variable	(1)	(2)	(3)	(4)
	$I(Deaths_{cpy} > Ceiling_{cpy})$			
$NSNP_{py}$	-0.051 (0.040)	-0.090 (0.053)	-0.098* (0.053)	-0.104* (0.057)
$GDP_{py}/GDPTarget_{py}$			-0.800* (0.388)	-0.903* (0.439)
Category and Year FEs	Yes	Yes	Yes	No
Province FEs	Yes	Yes	Yes	No
Cat. \times Year FEs	No	No	No	Yes
Cat. \times Province FEs	No	No	No	Yes
Full Sample	Yes	No	No	No
$Var(NSNP) > 0$ Sample	No	Yes	Yes	Yes
Observations	1755	855	855	855
Adjusted R-Squared	.128	.124	.125	.114


Notes: Standard Errors clustered at the province level. All analyses use province-year level data for the sample period 2005-2012. The dependent variable in all columns an indicator variable denoting whether $Deaths_{cpy}/Ceiling_{cpy}$ (strictly) exceeds one. NSNP is an indicator variable denoting whether the province has passed “No safety, no promotion” legislation by year y . See text for additional details on variable construction and definitions. The $Var(NSNP) > 0$ sample is the set of 15 provinces that passed NSNP legislation during the sample period 2005-2012 (and hence have within-province variation in NSNP). * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 3: Determinants of death ceilings

Dependent Variable	(1)	(2)	(3)	(4)
$\log(Deaths_{cpy-1})$	0.562*** (0.068)		0.549*** (0.071)	0.545*** (0.083)
$\log(Ceiling_{cpy-1})$		0.210*** (0.069)	0.059 (0.065)	0.065 (0.081)
$I(Deaths_{cpy-1} > Ceiling_{cpy-1})$				0.032 (0.101)
$I(Deaths_{cpy-1} > Ceiling_{cpy-1}) * \log(Deaths_{cpy-1})$				-0.006 (0.018)
Category \times Year FEs	Yes	Yes	Yes	Yes
Province \times Year FEs	Yes	Yes	Yes	Yes
Province \times Cat. FEs	Yes	Yes	Yes	Yes
Observations	1453	1449	1449	1449
Adjusted R-Squared	.991	.985	.991	.991


Notes: Standard Errors clustered at the province level. All analyses use province-year level data for the sample period 2005-2012. The dependent variable, $\log(Ceiling_{cpy})$ is the natural logarithm of province p 's assigned death ceiling in category c and year y , while $\log(Ceiling_{cpy-1})$ is the lagged value of this variable. $\log(Deaths_{cpy-1})$ is the lagged value of the natural logarithm of reported deaths. $I(Deaths_{cpy} > Ceiling_{cpi})$ is an indicator variable denoting whether reported deaths exceed the mandated ceiling in year $y - 1$. * significant at 10%; ** significant at 5%; *** significant at 1%.

Figure 1: Histogram for $Deaths_{cpy}/Ceiling_{cpy}$ for overall province-year level accidental deaths



Notes: The figure provides a histogram of the ratio of overall reported accidental deaths to the government-mandated ceiling, for province-year observations during 2005-2012.


Figure 2: McCrary density test for $Deaths_{cpy}/Ceiling_{cpy}$ for overall province-year level accidental deaths



McCrary Density Test (t-stat: -4.252)


Notes: The figure provides McCrary's density test for discontinuity in the distribution at 1, for the ratio of overall reported accidental deaths to the government-mandated ceiling. The sample includes province-year observations during 2005-2012.. The t-statistic for the significance of the discontinuity at 1 is listed just below the graph. The t-statistic for the significance of the discontinuity at 1 is listed just below the graph.

Figure 3: Pre vs Post-NSNP Cumulative density function of $Deaths_{cpy}/Ceiling_{cpy}$




Notes: Each line shows a cumulative density function for the ratio of reported deaths to government-mandated ceiling, $Deaths_{cpy}/Ceiling_{cpy}$. The solid line employs data from province-year observations where “No safety, no promotion” legislation had been passed (NSNP=0), while the dashed line employs observations where such legislation was in place (NSNP=1).

Figure 4: Relationship between lagged reported deaths and current death ceiling



Notes: The graph uses data from all category-year-province observations during 2005-2012. It shows the relationship between the natural logarithm of reported deaths in year $y - 1$ and the natural logarithm of the death ceiling in year y .

Figure 5: Pre vs. post-2004 cumulative distribution of % Δ Deaths for IMCT category



Notes: The graph depicts cumulative density functions for the annual change in accidental deaths, at the province-level, for industrial, non-coal mining, construction, and trade (IMCT) for the years 1993-1997, 1999-2003, and 2005-2012. We split the sample by 2004, the year of implementation of death ceilings. See Section 2 for details on the variable definition and sources.