The Economics of Death Ceilings

Raymond Fisman, Yongxiang Wang*
Boston University, University of Southern California

This Draft: January, 2016

Abstract

We study the effects of a 2004 program designed to motivate Chinese bureaucrats to reduce accidental deaths. Each province was given a set of 'death ceilings' that, if exceeded, would impede government officials' promotions. We analyze the extent to which death figures were manipulated to achieve incentive targets. For each category of accidental deaths, we observe a sharp, significant discontinuity in reported deaths at the ceiling. Provinces with safety incentives for municipal officials experienced more rapid declines in accidental deaths, suggesting a complementarity between incentives at different levels of government in achieving targets. There is no effect on industrial production when provinces are close to their death ceilings, as would be the case if real safety effort (through, say, additional spending on workplace safety or reduced production rates) accounted for the discontinuities we observe. Finally, we examine whether the dynamic incentives created by death ceilings affect death reporting. While realized accidental deaths predict next year's ceiling, we observe no evidence that provinces manipulate deaths upward to avoid a ratchet effect 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: rfisman@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, when output is difficult to observe or subject to manipulation, weak or even flat rewards may be optimal. This 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. In particular, 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 nonetheless been a push to utilize performance metrics in evaluating employees at organizations of all types (World at Work, 2014). 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); police officers on arrest and crime data (Eterno and Silverman, 2012). In all of these settings, incentives based on hard-to-quantify or manipulable outcomes naturally raises concerns around Goodhart's Law, that "When a measure becomes a target, it ceases to be a good measure" (Goodhart, 1975). Yet there have been few empirical analyses of the impact of incentive provision in bureaucracies where such concerns are likely to loom large.[1]

We aim to begin filling this gap with an analysis of a high-stakes safety targeting program implemented by the Chinese central government. 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.[2] 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 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 paper, the *People's Daily*. On the face of it, these reforms were a resounding success: Deaths across all categories fell by about 45 percent between 2005 and 2012, an annual decline of nearly 8 percent.

[1] See, for example, Dee et al (2014) for manipulation of Regents' exam scores by New York City teachers, though in this case the authors attribute manipulation to teachers' desire to help students rather than improve their own evaluations.

[2] See http://www.chinadaily.com.cn/english/doc/2004-02/20/content_307974.htm accessed on June 4, 2015

When we delve more deeply into the public data, however, we find evidence which indicates that manipulation of death statistics likely played a significant role in the reduction. Our analysis begins with McCrary's (2008) test for manipulation of a running variable, applied to the distributions of year-end deaths-to-ceiling ratios in each category. A discontinuous drop in the density function around one is an indication of 'sorting' of reported deaths around the ceiling. For all accident categories, McCrary's test indicates the presence of a sharp and statistically significant discontinuity in this ratio at one, providing clear evidence of responsiveness to the ceilings. [3] (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. A few provinces went as far as to publicize the municipal-level ceilings and death figures in province-owned newspapers or provincial SAWS websites, allowing us to collect year-end reported deaths and death ceiling data on 158 municipalities in nine provinces.[4] We use these data to show that there is a sharp discontinuity in the deaths-to-ceiling ratio at the municipal level for these cities. Having established that incentives affected reported outcomes at the local level, we provide tentative evidence that the provision of local incentives via NSNP was effective in reducing province-level deaths (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.

We may also exploit heterogeneity in local incentives induced by NSNP to assess whether bureaucrats' career advancement is impacted by meeting safety targets, using hand-collected information on local officials' careers during the period 2005 - 2012. Overall, we find no relationship between meeting safety targets and subsequent promotions. The average effect may be attenuated for (at least) two reasons: first, in non-NSNP province-year observations that comprise the bulk of our data, safety incentives at the province-level may have had only a modest impact on local officials' promotions. Second, officials face a classic multitask problem and, while safety may be included in performance evaluation, its role may be dominated by the importance of increasing output. We therefore compare the impact of safety on promotion in NSNP provinces, where we posit safety is, on the margin, given greater weight for local officials, versus non-NSNP provinces. For NSNP province-year observations, there is a large

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

[4] Guangdong, Guizhou, Heilongjiang, Liaoning, and Inner Mogolia provided relatively good coverage of such data, while Hubei, Hunan, Shaanxi, Shanxi ans SIchuan only provided such information for a few quarters.

and statistically significant impact of meeting safety targets on promotion. Interestingly, the actual number of deaths in a province has no explanatory power in predicting promotion, only the binary outcome of meeting the safety target.

While the very sharp discontinuities we observe in accidental deaths around annual ceilings is suggestive of manipulation, it could also be the result of extremely focused efforts by enforcement authorities at accident reduction *precisely around the ceiling*. To examine the presence of real safety reductions, we merge industrial cost and output data with workplace fatality figures. If bureaucrats meet targets through aggressive safety efforts rather than manipulation, we predict there will be reduced output or increased per unit production costs in province-quarter observations when worker deaths approach the ceiling. While the production data are noisy - and hence these results need be interpreted with considerable caution - we find no evidence of production slowdowns or increased expenditures when provinces are close to their ceilings, arguing against real effort as playing a central role in generating the discontinuities we document.

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 go 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 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. Our data do not allow us to distinguish between these classes of 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 a governmental and/or non-profit setting. As Baker et al (1988) emphasize, incentives observed in practice often depart from what economic theory might have predicted. We take our work as part of the larger effort to document empirically the effects of incentives, to better understand how economic models may be reconciled with observed behavior.

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 undermines 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 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.*[5] 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. Additionally, the document 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 be required 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, which reports 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.[6]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, including Severe (3 or more deaths) and Very Severe (10 or more deaths). Two additional 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 fatal accidents. 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., a number of severe accidents take place in coal mining, and all subcategories

[5]The *Decision*'s name in Chinese is Guowuyuan Guanyu Jinyibu Jiaqiang Anquan Shengchan Gongzuo De Jueding, and its text may be found at http://www.china.com.cn/chinese/PI-c/483862.htm (accessed on June 4, 2015).

[6]The Work Safety Commission is a unit of the State Council. The Commision'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 Chinese text of the Commission's opinion (Guanyu Jianli Anquan Shengchan Kongzhi Zhibiao Tixi De Yijian) may be accessed at http://www.safetyhome.org/news/displaynews.asp?id=1130 (accessed on Feb 15, 2015).

sum to the overall figure) and potential serial correlation, we will allow for province-level clustering of standard errors below.

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 they were not publicly available, and received supplemental data for 2007 and 2008, as well as the fourth quarter of 2012. [7] While the target system remains in place, we have been unable to obtain data for 2013 or later.[8] 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 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.

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 were 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

[7]SAWS 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.

[8]This can be verified via the safety target directives issued by prefectural government for each county, which may be obtained online. See, for example, http://www.hengyangsafety.gov.cn/upfile/2014/05/28/20140528155846_529.pdf for a regulatory document issued by the SAWS branch in the city of Hengyang, Hunan Province, which specified the decomposition of the death ceiling received from the Hunan Provincial government.

a natural death. This created an incentive to keep accident victims alive for at least a week, especially when close to the year-end targeted maximum.

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.[9] 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 were thought to continue after its publication.

In addition to the exploitation of death classification loopholes, deaths were thought to be under-reported by the local governments responsible for providing these figures to provincial officials (Shanghai Law Journal, 2007), as in 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." [10] 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

For the main variables in our analysis, we hand collected province-level quarterly reported deaths and corresponding death ceiling data for each category from the *People's Daily*. Our sample covers the period 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.

The reported deaths in category c in province p in year y is denoted by $Deaths_{cpy}$. When we look at quarterly data we add an appropriate suffix to the variable (e.g., $Deaths\_Q3_{cpy}$), which denotes the *cumulative* deaths through to the end of the third quarter; the flow of deaths during this quarter is denoted by $\Delta Deaths\_Q3_{cpy}$. We similarly append the relevant suffix when a set of analyses focus on a specific category (e.g., $Deaths\_IMCT_{py}$). Figure 1 shows the accidental deaths for the country as a whole for the Overall category for the years 2005 - 2012. The graph shows an annual decline of between 4 and 11 percent, with steep declines experienced particularly in the earlier years.

[9] The Chinese text of this document ("Guanyu Shengchang Anquan Shigu Rending Ruogan Yijian Wenti De Han (Zhengfahan [2007] 39 Hao)") is available at http://wenku.baidu.com/view/aff5de48e518964bcf847c18.html (accessed on Feb 10, 2015).

[10] The expression in Chinese is "cun pian xiang, xiang pian xian, yizhi daole guowu yuan" (xiang, xian, and yuan rhyme in Mandarin). For reference to this problem in the implementation of the death ceiling program, see "Should death ceiling system be established?" (Siwang Zhibiao Gaibugai Sheli?), *Shanghai Law Journal*, May 30, 2007.

Our main interest will be in analyzing the relationship between reported deaths and mandated death ceilings, denoted by $Ceiling_{cpy}$, at the province level. Our main outcome of interest will be 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 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 Panel A provides summary statistics on annual deaths, $Deaths_{cpy}$, as well as the deaths-to-ceilings ratio, $Deaths_{cpy}/Ceiling_{cpy}$. Road deaths constitute nearly 80 percent of deaths in the Overall category. While this is puzzling at first blush, note that, apart from Road and Fire, the Overall category includes only workplace deaths. Given this, a ratio of 80 percent is in line with comparable figures from other countries. (The ratio of traffic deaths to workplace deaths in the U.S., for example, is about eight to one, even higher than what we observe in our sample.[11])

The figures in Table 1A, 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.>[12] 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 1A imply a Chinese workplace death rate in 2005 of 1.8 per 100,000 population, virtually identical to the U.S. rate.[13] 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 a small subset of municipalities (indexed by m) in NSNP provinces, we were able to obtain data on city-level deaths $Deaths_{cmy}$, as well as the municipal ceiling, $Ceiling_{cmy}$. We collect information from two sources: provincial SAWS branches' official websites, and official newspapers controlled by individual provincial governments. These municipal-level data are not as detailed as those published in the *People's Daily*, and typically only cover only a few death categories: Overall accidental deaths, Coal, IMCT, and Road. Appendix Table A3 provides the number of municipalities with coverage

[11] See http://www.cdc.gov/nchs/fastats/accidental-injury.htm accessed on June 6, 2015.

[12] See http://www.iihs.org/iihs/topics/t/general-statistics/fatalityfacts/state-by-state-overview/2005 accessed June 6, 2015.

[13] See http://www.bls.gov/iif/osh_nwrl.htm#cfoi accessed June 6, 2015

for each province-year cell for the nine provinces where we were at least partially successful in obtaining city-level data. In Table 1, Panel B, we presents the municipality-year deaths data by category.

To examine whether meeting death targets affected officials' promotions, we collected information on the identities of mayors and city-level party secretaries going back as far as 1990, using two sources. The first is the Local Leader Database , maintained by People.cn, the official website of the *People's Daily*.[14] This database was used to obtain current city-level information on high-ranking officers.. The second source is the set of Provincial Yearbooks, published by each province annually, which provides information on former local leaders. After obtaining the names of these high-level city officials, we use baike.Baidu.com to obtain the detailed resume for each leader.[15]. Based on these resumes, for each official we are able to code the positions held both before and after serving as a municipal leader. For each employment spell we record the city or province where the official was stationed, the official title and rank, and the years the position was held. We additionally collected data on age to account for retirements unrelated to performance.

For each official-year observation, we code $Promotion = 1$ if the official's next position is of a higher rank, based on (i) a higher title within the municipal or provincial bureaucracy; (ii) a move to a comparable position in a higher-ranked municipality. To illustrate, consider the case of Mr. Jianye Ji, who became mayor of Yangzhou City in 2002, where he also acted as party vice-secretary. (In almost all cities in China, the mayor serves as party vice-secretary.) Two years later, in 2004, he was promoted to party secretary of Yangzhou. In 2010, he then became mayor and party vice-secretary of Nanjing, which serves as the capital of Jiangsu Province and is officially designated as a vice-province-level city (Fu Shengji Chengshiin in Chinese).[16] We record two promotions for Mr. Jianye. First, following the oft-repeated principle of "The Party leads everything", we code his 2004 transition from Yangzhou City mayor to Yangzhou City party secretary as a promotion.[17] Second, in his move to Nanjing, as mayor of a vice-province-level city, Mr. Jianye Ji held a rank equivalent to vice-minister, a significant promotion relative to party secretary of a lesser municipality.

Of the 1,369 officials who, at some point during 2004 - 2014, held the position of municipal mayor or party secretary (a total of 2,172 official$\times$position observations), nearly half (48.3 percent) obtained a promotion in their next position. By far the highest promotion rate is among deputy city mayors. First, considering officials in non-vice-province level cities, of the

[14]See http://ldzl.people.com.cn/dfzlk/front/firstPage.htm (accessed June 6, 2015)

[15]Baidu Baike is similiar to Wikipedia in terms of function, and is a division of Baidu, China's leading search engine.

[16]There are 15 such cities in China: Shenyang, Dalian, Changchun, Ha'erbin, Nanjing, Hangzhou, Ningbo, Xiamen, Jinan, Qingdao, Wuhan, Chengdu, Xi'an, Guangzhou, Shenzhen.

[17]See http://news.xinhuanet.com/ziliao/2002-11/18/content_633225.htm for full text of the Chinese Constitution for a direct reference. Accessed on June 6, 2015.

575 deputy city mayor observations (26.5 percent of the sample), 517 (89.9 percent) received promotions. There are 738 city mayor observations (34.0 percent), of whom 379 (51.4 percent) were promoted. Promotion is lower for city party secretaries: of the 663 observations (30.5 percent of the sample), 80 (12.1 percent) were promoted. The promotion rates are mostly lower for officials in vice-province cities (82, 22.2, and 21.6 percent for deputy mayor, mayor, and party secretary respectively).

In order to merge our official×position data into our province×year data, we proceed as follows: if a mayor or party secretary is promoted at the end of his current term, we set $Promotion = 1$ for all years during his current term. So, in our example above, Mr. Jianye Ji would have $Promotion = 1$ for each year during 2002 - 2003, *and* 2004 - 2009. If, at the end of his term, an official is not promoted, we define the official-year observation as having $Promotion = 0$ for all years of that term. We then aggregate these promotion data at the province-year level to construct $FractionPromoted_{py}$, the proportion of high-level municipal officials in office in province p in year y for whom $Promotion = 1$.

For most accidental death, 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. 1998 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.[18] Figure 2 provides a pair of histograms showing, at the province-year level, the distribution of $\%\triangle DeathsIMCT_{py}$, the annual percentage change in workplace-related deaths in the IMCT category, for pre- and post-2004. The distribution shows a much lower dispersion for the post-2004 data, mainly as a result of the small number of observations with positive growth in the later period. In fact, the 25th percentiles of the two distributions are identical (-0.13), while the 75th percentile is far higher for the pre-2004 period (0.15 versus -0.01 for the post-2004 period). This is a point we return to in greater detail in Section 3.4.

From the Guoyan database, we further obtain province-quarter level sales ($Sales\_IMCT_{cpy}$) and cost ($Costs\_IMCT_{py}$) information for IMCT industries, for the years 2001 - 2012.[19] These data may be matched to IMCT death data to examine whether death ceilings have any impact on industrial production.

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

[18] 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.

[19] http://www.drcnet.com.cn/, maintained by the Development Research Center of the State Council of China. Accessed on June 6, 2015.

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.[20] We combine the GDP growth targets with realized real GDP data ($GDPGrowth_{py}$) 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).

Table 1, Panel C, provides summary statistics for the province-year level variables $GDPGrowth$, $FractionPromoted$, $GDP/GDPTarget$, and $Population$, as well as for production cost and revenue data for for the IMCT category.

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.[21]

In Figure 3A, we first 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 used by McCrary's density test. A discontinuity is clearly discernable at one. Figure 3B presents the McrCrary 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 Figures 4A - 4H, 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

[20] See, for example, the 2015 report for Hunan Province is published on the official website of Hunan Province (http://www.hunan.gov.cn/zwgk/hnyw/zwdt/201502/t20150205_1209622.html). Accessed June 10, 2015.

[21] This 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 6, 2015.

or concealment [22]

In Appendix Figure A1 we show histograms for each category of the distribution of $DeathsQ3_{cpy}/Ceiling_{cpy}$, 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 already are close to the year-end threshold by the end of Q3.) Thus, a 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. Before comparing accidental death rates in $NSNP = 0$ versus $NSNP = 1$ provinces, we first assess whether the local targets provided through NSNP affected the distribution of accidental deaths at the municipal level. This analysis comes with the strong caveat that we may only observe these data for municipalities that *self-select* into making these data public, possibly because they are most capable of meeting their targets. With this qualification in mind, Figures 5A - 5D present a sequence of McCrary density tests for discontinuities in the municipality-level data deaths-to-ceiling ratio. For the first three figures (overall; road; IMCT), we observe clearly discernible discontinuities (significant at the 1 percent level) in the distribution of $Deaths_{cpy}/Ceiling_{cpy}$ around a cutoff of one; for coal the discontinuity, while negative, is insignificant, though we only have a total of 51 province-year observations, as compared to 181 for overall deaths, 136 for road, and 99 for IMCT.

To examine whether these local effects aggregate to better province-level responsiveness to safety targets, we compare pre- versus post-NSNP passage death rates and probabilities of meeting safety targets. Owing to the relative sparseness of the data, we perform this analysis for the pooled sample of all accidental death categories.[23] 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:

[22] There are certainly examples that have come to light of attempted suppression of news of high-fatality 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.

[23] For 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.

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

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 fact that the categories are non-independent.

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 (3) we limit the analysis to provinces that implemented *NSNP* incentives during 2005 - 2011, which we refer to as the $Var(NSNP) > 0$ sample. The magnitude of 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.)[24]

In Figure 6 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

[24]If 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.

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 to the extent that it reduced the likelihood of exceeding safety targets - conditional on meeting the target, the two distributions are quite similar.

The imposition of NSNP regulations also allows us to better isolate the effect of safety incentives on promotions. This relationship might be obscured in the data overall, as bureaucrats face a classic multitask problem: during our sample period production and economic growth targets remained in place, and to the extent that output-safety tradeoffs had to be made, they might still favor increased output (though it is not obvious that this is the case, given the correlations reported in Table 2). Further, because the 2004 death ceiling program was imposed uniformly (and we further have a paucity of pre-period data), we have no benchmark against which to compare the effect of accidental deaths on promotion under the death ceiling system. We will control in what follows for $GDP_{py}/GDPTarget_{py}$, to account for whether provinces met their targeted GDP growth rates, but as these are self-imposed they may not fully capture the role of economic performance. We therefore exploit NSNP regulations to impose an additional layer of differences, comparing the correlation between accidental deaths and promotion in $NSNP = 1$ provinces versus those that had yet to impose NSNP incentives, which presumes that safety played a greater role in promotion decisions in NSNP provinces. We use the following specification:

$$\begin{aligned} FractionPromoted_{py} = \quad & \beta_1 I(Deaths_{cpy} > Ceiling_{cpy}) + \beta_2 NSNP_{py} \\ & + \beta_3 I(Deaths_{cpy} > Ceiling_{cpy}) * NSNP_{py} \\ & + Controls_{cpy} + \mu_{cy} + \upsilon_{cp} + \epsilon_{cyp} \end{aligned} \tag{2}$$

In column (1) of Table 3 we first present results including only the direct effect of exceeding the ceiling, controlling for NSNP status and including year, province, and category fixed effects. There is no average effect of NSNP or $I(Deaths_{cpy} > Ceiling_{cpy})$ on promotions of local officials. In fact, though β_1 is very small in magnitude, its point estimate goes in the "wrong" direction. In column (2) we add the interaction term in specification 2. Its coefficient is negative and significant at the 5 percent level (p-value = 0.011), indicating a greater negative impact of exceeding death ceilings on fraction promoted when $NSNP = 1$. A test of $\beta_1 + \beta_3 = 0$ indicates that the sum (the overall effect of missing safety targets when $NSNP = 1$) is significantly less than zero at the 5 percent level (p-value = 0.023). In column (3) we limit the sample to $Var(NSNP) > 0$ provinces, and in column (4) we further include category-province and category-year fixed effects. In column (5) we add $GDP_{py}/GDPTarget_{py}$ and $GDP_{py}/GDPTarget_{py} * I(Deaths_{cpy} > Ceiling_{cpy})$ as controls. The coefficients of interest remain largely unchanged. In column (6) we replace $I(Deaths_{cpy} > Ceiling_{cpy})$ with the death-to-ceiling ratio, $Deaths_{cpy}/Ceiling_{cpy}$. The coefficient on the interaction term $I(Deaths_{cpy} > Ceiling_{cpy}) * NSNP_{py}$ no longer approaches significance, suggesting a distinct effect of meeting safety targets on promotion, rather than simply promoting safety more broadly.[25]

3.3 Real effects of safety incentives

The very sharp discontinuities we observe in Section 3.1, particularly given the lack of any discernable discontinuity prior to the final quarter of the year, are most readily reconciled with manipulation of death data. It is plausible that aggressive safety efforts might produce the same result, but that would likely come with at least some cost. (If it were costless, it is natural to ask why accidental deaths failed to decline previously.) The category that lends itself most readily to examining real efforts at safety provision is industrial production (IMCT) where, we conjecture, real efforts at reducing workplace fatalities should result in either higher production costs or lower output.

It is difficult to make general statements about the relationship between the deaths-to-ceiling ratio at the end of the third quarter, $Deaths_IMCT_Q3_{py}/Ceiling_IMCT_{py}$and production slowdowns: this will depend on the incremental cost of effort, and the likelihood that it will be decisive in bringing the province under its ceiling.[26]

We can make somewhat sharper predictions about fourth quarter effort as a function of whether, ex post, the *year-end* deaths-to-ceiling ratio is close to one. Intuitively, among the many observations clustered just below one (31.5 percent of all observations lie in the range [0.95,1], and 19.0 percent are in [0.98,1]) we would anticipate that in a relatively high fraction of cases manipulation and/or effort would have been required to keep the ratio below one. We further expect that the required manipulation or effort would have been greater (or with higher probability) if, at the end of the third quarter, the deaths-to-ceiling ratio were already close to one. In Figure 7A, we show a scatterplot of the year-end deaths-to-ceiling ratio, $Deaths\_IMCT_{py}/Ceiling\_IMCT_{py}$, and the fraction of sales that took place during the fourth quarter ($\Delta Sales\_IMCT\_Q4_{py}/Sales\_IMCT_{py}$), along with a locally weighted regression line of best fit. If production were slowed to keep fatalities under the ceiling,

[25] To conserve space we do not include other specifications involving $Deaths_{cpy}/Target_{cpy}$, but in no case does β_c approach significance. When both $Deaths_{cpy}/Target_{cpy}$ and I($Deaths_{cpy} > Target_{cpy}$) are included in the same specification, the sign on $(Deaths_{cpy}/Target_{cpy}) * NSNP_{py}$ actually becomes significant and positive. Given the colinearity of these two variables, however, one may not impose any interpretation on this fact.

[26] Interestingly, we observe a discontinuity in the year-end deaths-to-ceiling ratios for all four quartile subsamples of $Deaths\_IMCT\_Q3_{py}/Ceiling\_IMCT_{py}$. That is, regardless of a province's "distance" to its ceiling at the end of Q3, we observe a discontinuity by year's end. This indicates that (a) there is some probability of exceeding the ceiling regardless of one's position at the beginning of the fourth quarter, and hence a potential need to reduce reported deaths in Q4; and (b) as we discuss in the next section, some provinces with very low third quarter deaths may be motivated to *increase* fourth quarter deaths to avoid too harsh a ceiling the following year. These considerations highlight the difficulty in generating sharp predictions relating third quarter deaths-to-ceiling rates to the returns to effort in the fourth quarter.

we would anticipate that sales would, on average, be relatively low for observations with a deaths-to-ceiling ratio just under one. While we observe a clustering of observations just below one - essentially the bunching we observed in the histogram in Figure 3 - there is no discernable relationship between the death-to-ceiling ratio and fourth quarter sales. In Figure 7B, we replace the fourth quarter sales rate with the fourth quarter cost-to-sales ratio ($\Delta Costs\_IMCT\_Q4_{py}/\Delta Sales\_IMCT\_Q4_{py}$). If production costs increased in order to remain under the ceiling, we would expect observations with deaths-to-ceiling ratios just below one to have relatively high cost-to-sales ratios. But again we observe no distinct pattern for observations just below one. Overall, we find no evidence that production decisions were affected by safety targets.

We may make a sharper prediction for the subset of observations with a death-to-ceiling ratio just below one at year's end: among these, we would anticipate that the most effort (or greatest output reduction) would be required for provinces with reported deaths *already* close to their ceilings at the end of the third quarter. Figures 8A and 8B show the relationship between $Deaths\_IMCT\_Q3_{py}/Ceiling\_IMCT_{py}$ and the fourth quarter sales fraction and cost ratio, respectively, for observations where the province came close to its ceiling by year's end ($0.95 \leq Deaths\_IMCT\_Q4_{py}/Ceiling\_IMCT_{py} \leq 1$). We again observe no discernable relationship - the fourth quarter rate of output and cost-to-sales ratios for provinces that ended up very close to their death ceilings were unaffected by the "distance" to the ceiling at the beginning of the fourth quarter.

There are several major caveats to interpreting these results as evidence against any "real" effect of safety targets. First, if death reporting is subject to manipulation, the same might be true of production data, and if the official output figures we use are essentially noise we may not observe any statistical regularities. Second, safety measures may have been taken that are not captured in cost or output data. This might include greater attentiveness to safety on the part of managers - that did not distract them from other output-enhancing or cost-reducing efforts - or greater vigilance by workers that similarly did not affect costs or output.

As we observe in the preceding section, however, there is also reason to expect that output would be unaffected by safety targeting, and in this sense the lack of connection between safety targeting and output decisions could plausibly be so small as to be indetectable. Officials face a multitasking problem and, despite the advent of safety incentives, economic growth remained a primary concern of the State Council. (Wu et.al. (2013), for example, show that the central importance of economic growth for provincial officials' promotion continued even after the central government began emphasizing the importance of improving environmental quality. Government proclamations notwithstanding, provincial officials' promotion still depended in large part on GDP growth, while environmental quality had no effect on promotion.) In the absence of a "free lunch" in safety reduction, the sharp death reductions we document around the death ceiling could then be attributed in large part to manipulation.

3.4 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 are not 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 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 9, 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 4, 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 effects of death ceilings may lead to manipulation of reported deaths downward (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. [27] 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 would 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 10, 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, whereas 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 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 the introduction of a high-stakes incentive program that aimed to improve safety practices in China. The sharp discontinuities we observe in the deaths-to-ceiling distribution across all accident categories are indicative of a strong response to safety incentives. We argue that three patterns, when taken together, strongly suggest that manipulation played a dominant role in accidental death reduction: (a) the sharpness of the year-end deaths-to-ceiling discontinuities; (b) the absence of discontinuities in this ratio at the end of the third quarter; (c) the absence of any discernable impact on 'real' outcomes.

Our findings may be seen first and foremost as a cautionary tale on the use of incentives based on performance metrics in bureaucracies. In the absence data on actual accidental deaths, though, 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). This limits our ability to make a general assessment of the welfare effects of the program. Nor can we necessarily 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

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

its objective, from the perspective of the implementing authority - 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] 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).

[5] Dixit, Avinash. "Power of incentives in private versus public organizations." The American Economic Review (1997): 378-382.

[6] Easley, David, and Maureen O'Hara. "The economic role of the nonprofit firm." The Bell Journal of Economics (1983): 531-538.

[7] Eterno, John A., and Eli B. Silverman. The crime numbers game: Management by manipulation. CRC Press, 2012.

[8] Goodhart, Charles. Problems of monetary management: the UK experience. 1975.

[9] Holmstrom, Bengt. "Moral hazard and observability." The Bell Journal of Economics (1979): 74-91.

[10] Holmstrom, Bengt, and Paul Milgrom. "Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design." Journal of Law, Economics, & Organization (1991): 24-52.

[11] Hao, Jiawei. "Decoding Death Targets in China." Southern Weekly (May 4th, 2007).

[12] Mookherjee, Dilip. "Decentralization, hierarchies, and incentives: A mechanism design perspective." Journal of Economic Literature (2006): 367-390.

[13] Mookherjee, Dilip. "Political Decentralization." Forthcoming, Annual Review of Economics (2015).

[14] McCrary, Justin. "Manipulation of the running variable in the regression discontinuity design: A density test." Journal of Econometrics 142.2 (2008): 698-714.

[15] Oyer, Paul. "Fiscal year ends and nonlinear incentive contracts: The effect on business seasonality." Quarterly Journal of Economics (1998): 149-185.

[16] Rivkin, Jan W., Michael Roberto, and Ranjay Gulati. "Federal Bureau of Investigation, 2009." (2009).

[17] Su Lin, “Death Ceilings and the dilemma of local safety regulations in China.” Working Paper (2011).

[18] 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.

[19] World at Work, Incentive Pay Practices Survey: Non-Profit/ Government Organizations, 2014, mimeo

[20] Wu, Jing, et al. Incentives and outcomes: China’s environmental policy. No. w18754. National Bureau of Economic Research, 2013.

Table 1: Province-year summary statistics, 2005-2012

	Mean	StdDev	Min	Max	Obs	Mean	StdDev	Min	Max	Obs
Panel A: Deaths at the province-year level, by category										
			$Deaths_{cpy}$					$Deaths_{cpy}/Ceiling_{cpy}$		
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
Panel B: Deaths at the municipal-year level, by category										
			$Deaths_{cmy}$					$Deaths_{cmy}/Ceiling_{cmy}$		
All	120.68	120.11	0	1233	1406	.48	.3	0	2.13	1271
IMCT	24.52	38.12	0	323	582	.44	.4	0	2.41	528
Road	113.81	124.16	0	1131	578	.53	.31	0	1.58	552
Coal	21.37	37.58	0	299	351	.69	1.09	0	8.67	269
Panel C: Additional province-year level variables										
$FractionPromoted_{py}$	.34	.15	0	1	216					
Real GDP growth	-98.87	.02	-98.95	-98.76	248					
Population (millions)	43.74	26.3	5.43	105.94	240					
GDP/GDPTarget	1.01	.03	.91	1.07	248					
Sales (Year End; millions of RMB)	1765.87	2161.56	43.15	11745.28	240					
Costs (Year End; millions of RMB)	1507.09	1874.12	31.4	10209.12	240					

Notes: The first set of columns in Panel A 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. Panel B provides summary statistics for municipal-level accidental deaths for the subset of observations listed in Apendix Table A2. 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). Panel C provides summary statistics for additional variables employed in our analysis, at the province-year level. *FractionPromoted* is the fraction of local officials in province p in year y who receive a promotion in their next posting. See text for details of this variable's construction.

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

	(1)	(2)	(3)	(4)
Dependent Variable	$I(Deaths_{cpy} > Ceiling_{cpy})$			
$NSNP_{py}$	-0.051	-0.090	-0.098*	-0.104*
	(0.040)	(0.053)	(0.053)	(0.057)
$GDP_{py}/GDPTarget_{py}$			-0.800*	-0.903*
			(0.388)	(0.439)
Category and Year FEs	Yes	Yes	Yes	No
Province FEs	Yes	Yes	Yes	No
Cat. × Year FEs	No	No	No	Yes
Cat. × 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: Accidental deaths and promotion of municipal leaders

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable			$FractionPromoted_{py}$			
$I(Deaths_{cpy} > Ceiling_{cpy})$	0.010	0.016	0.019*	0.023*	0.023*	
	(0.008)	(0.010)	(0.009)	(0.011)	(0.011)	
$NSNP_{py}$	-0.002	0.003	0.003	0.007	0.507	-0.003
	(0.030)	(0.030)	(0.037)	(0.040)	(0.571)	(0.045)
$I(Deaths_{cpy} > Ceiling_{cpy}) * NSNP_{py}$		-0.041**	-0.042**	-0.054**	-0.052**	
		(0.015)	(0.015)	(0.022)	(0.022)	
$Deaths_{cpy}/Ceiling_{cpy}$						0.005
						(0.010)
$Deaths_{cpy}/Ceiling_{cpy} * NSNP_{py}$						-0.001
						(0.019)
$GDP_{py}/GDPTarget_{py} * NSNP_{py}$					-0.492	
					(0.570)	
$GDP_{py}/GDPTarget_{py}$					0.385	
					(0.490)	
Category and Year FEs	Yes	Yes	Yes	No	No	Yes
Province FEs	Yes	Yes	Yes	No	No	Yes
Cat. × Year FEs	No	No	No	Yes	Yes	No
Cat. × Province FEs	No	No	No	Yes	Yes	No
Full Sample	Yes	Yes	No	No	No	No
$Var(NSNP) > 0$ Sample	No	No	Yes	Yes	Yes	Yes
Observations	1540	1540	799	799	799	799
Adjusted R-Squared	.459	.46	.517	.431	.434	.511

Notes: Standard Errors clustered at the province level. All analyses use province-year level data for the sample period 2005-2012. The dependent variable, *FractionPromoted* is the fraction of local officials in province p in year y who receive a promotion in their next posting. See text for details of this variable's construction. $Deaths_{cpy}/Ceiling_{cpy}$ is the ratio of reported deaths to the government-mandated death ceiling, and $I(Deaths_{cpy} > Ceiling_{cpi})$ is an indicator variable denoting whether this ratio 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 4: Determinants of death ceilings

	(1)	(2)	(3)	(4)
Dependent Variable		$log(Ceiling_{cpy})$		
$log(Deaths_{cpy-1})$	0.562***		0.549***	0.545***
	(0.068)		(0.071)	(0.083)
$log(Ceiling_{cpy-1})$		0.210***	0.059	0.065
		(0.069)	(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 × Year FEs	Yes	Yes	Yes	Yes
Province × Year FEs	Yes	Yes	Yes	Yes
Province × 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: Accidental Deaths by Year, National Aggregate 2005-2012

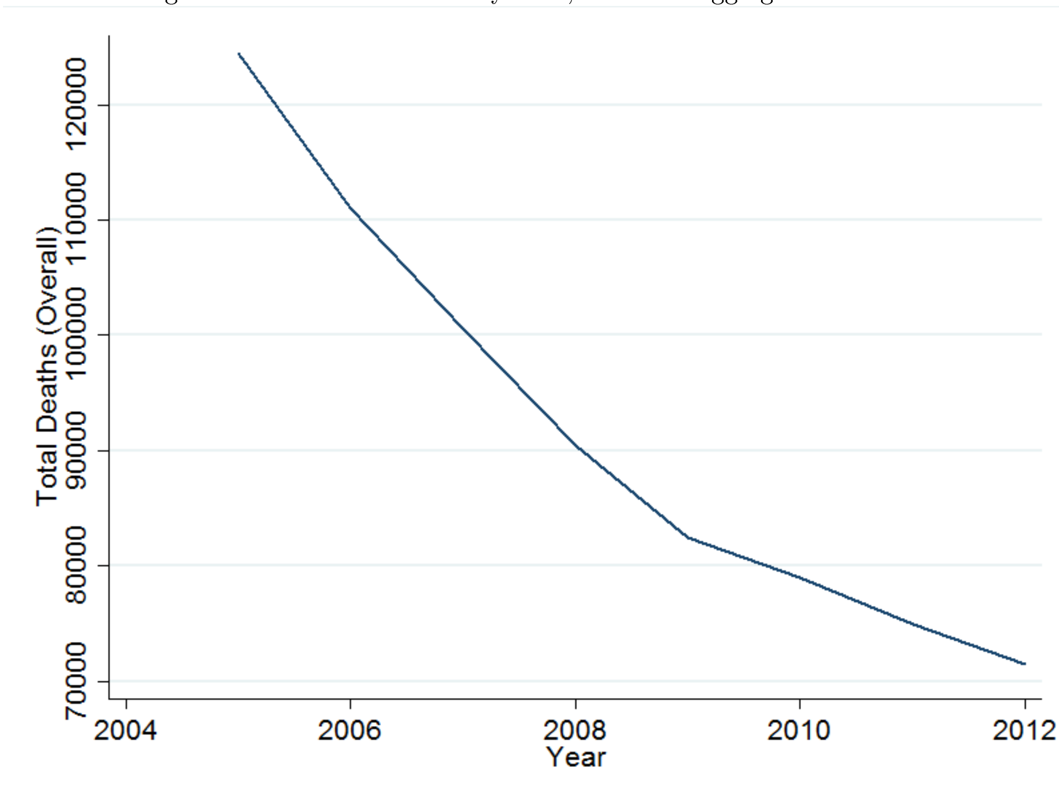

Figure 2: Annual change in industrial deaths (IMCT), pre- versus post-2004

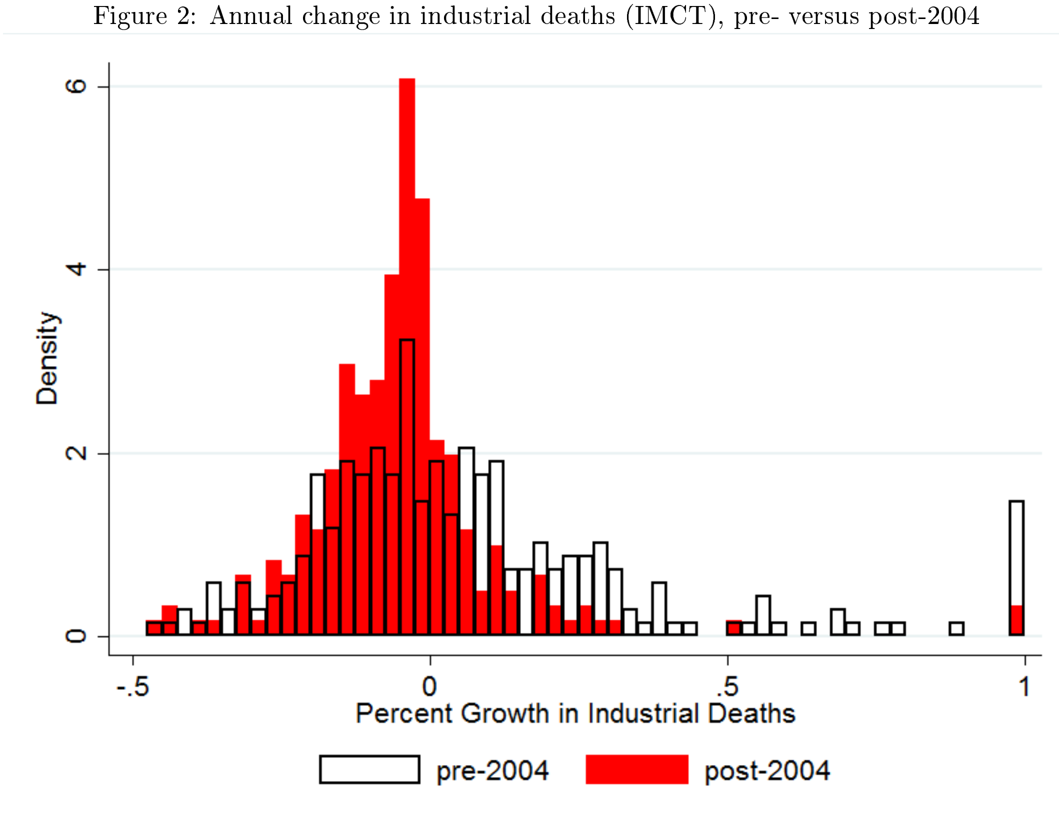

Notes: Histograms show the annual change in accidental deaths, at the province-level, for industrial, non-coal mining, construction, and trade (IMCT) $\%\Delta DeathsIMCT_{py}$ for the years 1993-1997, 1999-2003, and 2005-2012. We split the sample by 2004, the year of implementation of death ceilings. See Section 2 for details on the variable definition and sources.

Figure 3: Histogram and McCrary density test for $Deaths_{cpy}/Ceiling_{cpy}$ for overall province-year level accidental deaths

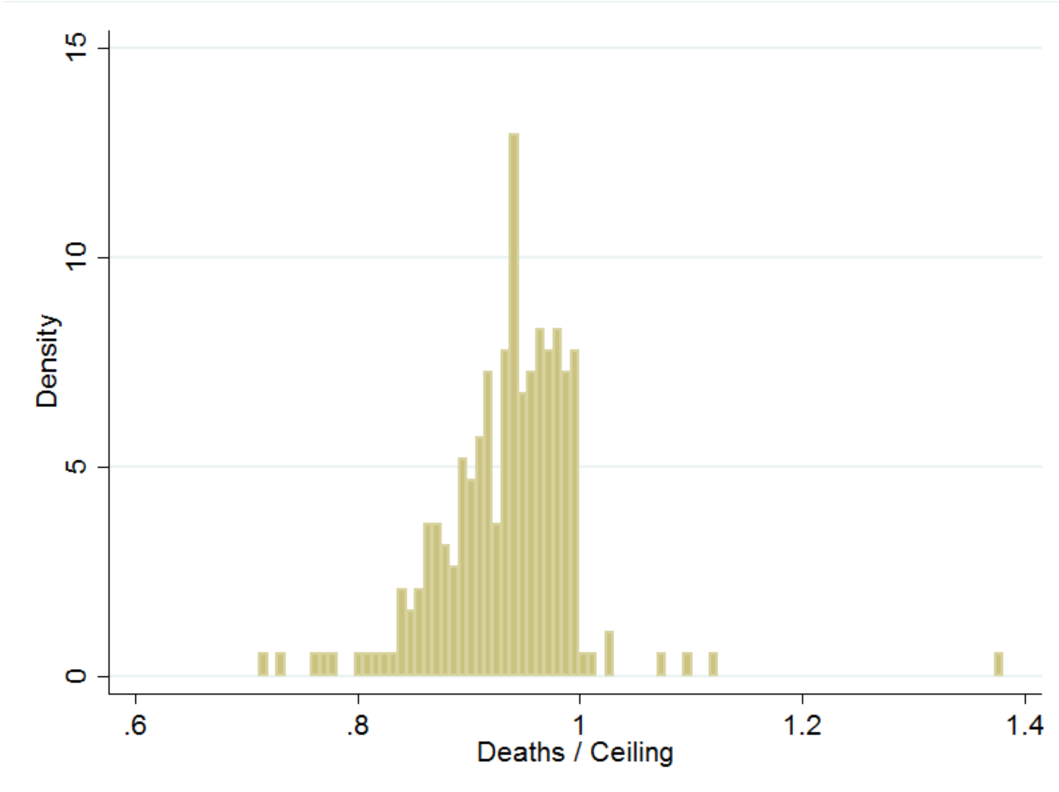

Panel A: Histogram

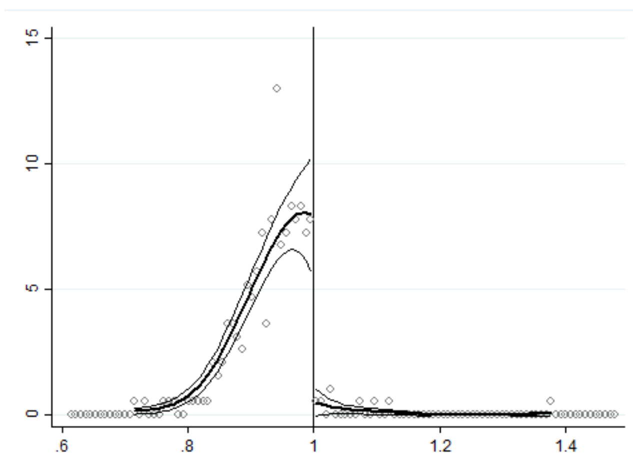

Panel B: McCrary Density Test (t-stat: -4.252)

Notes: Panel A shows the histogram of the ratio of overall reported accidental deaths to the government-mandated ceiling, for province-year observations during 2005-2012. Panel B shows McCrary's density test for discontinuity in the distribution at 1. The bold line shows the local linear regression estimated on either side of 1, and the lighter lines show the 90 percent confidence interval. The t-statistic for the significance of the discontinuity at 1 is listed below Panel B.

Figure 4: McCrary density tests for $Deaths_{cpy}/Ceiling_{cpy}$ for province-year level accidental deaths, by category

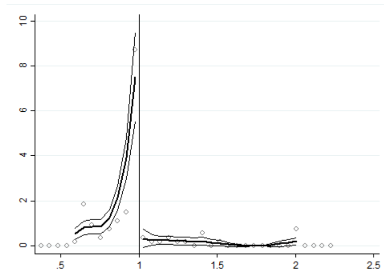

Panel A: Agriculture
McCrary Density Test t-stat: -3.995

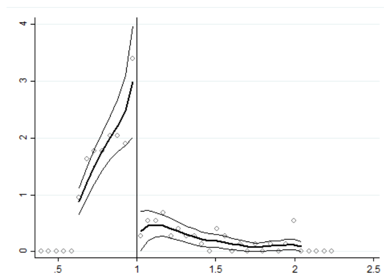

Panel B: Coal
McCrary Density Test t-stat: -3.426

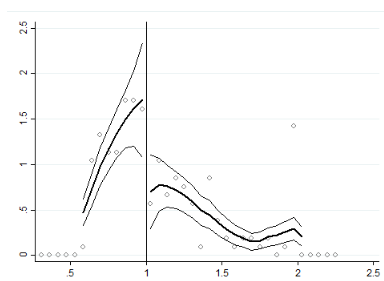

Panel C: Fire
McCrary Density Test t-stat: -2.138

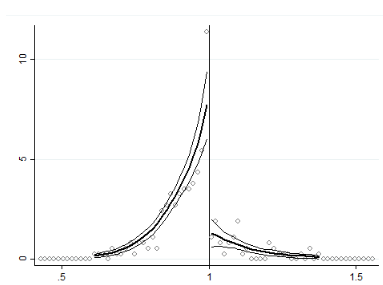

Panel D: IMCT
McCrary Density Test t-stat: -5.76

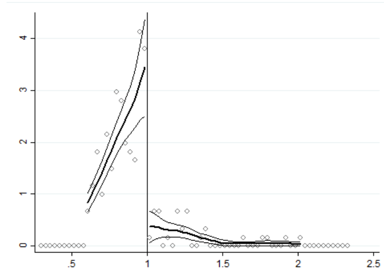

Panel E: Railways
McCrary Density Test t-stat: -4.557

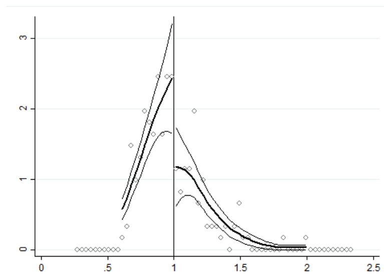

Panel F: Severe Accidents
McCrary Density Test t-stat: -2.321

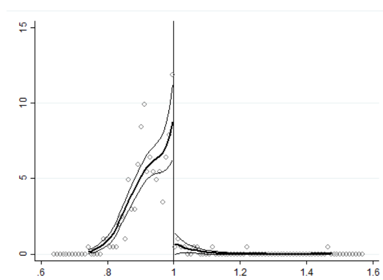

Panel G: Road
McCrary Density Test t-stat: -4.359

Notes: Each panel shows a McCrary density test for discontinuity in the distribution of $Deaths_{cpy}/Ceiling_{cpy}$ at 1 for a separate category of accidental deaths. In each case, the bold line shows the local linear regression estimated on either side of 1, and the lighter lines show the 90 percent confidence interval. The t-statistic for the significance of the discontinuity at 1 is listed below each panel.

Figure 5: McCrary density tests for $Deaths_{cpy}/Ceiling_{cpy}$ for municipal-year level accidental deaths, by category

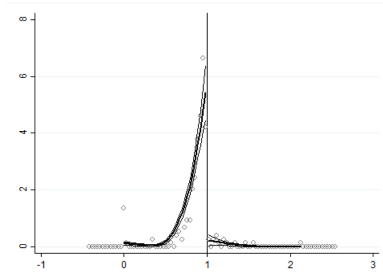

Panel A: All; McCrary Density Test t-stat: -6.546

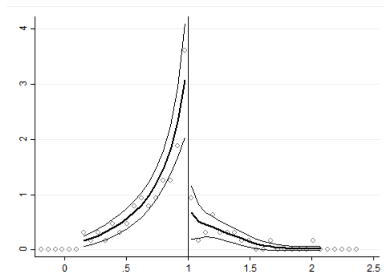

Panel B: IMCT; McCrary Density Test t-stat: -3.318

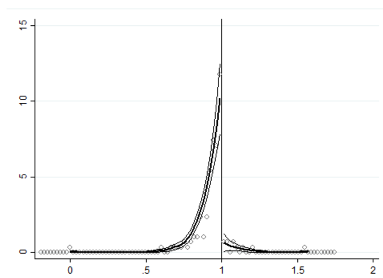

Panel C: Roads; McCrary Density Test t-stat: -5.242

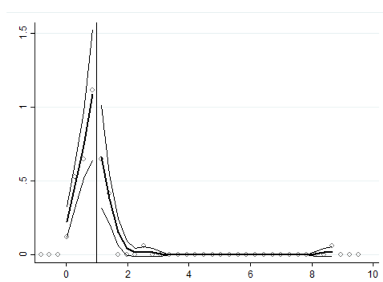

Panel D: Coal; McCrary Density Test t-stat: -1.059

Notes: The t-statistic for the significance of the discontinuity at 1 is listed below each panel.

Figure 6: Pre vs Post-NSNP Cumulative density function of $Deaths_{cpy}/Ceiling_{cpy}$

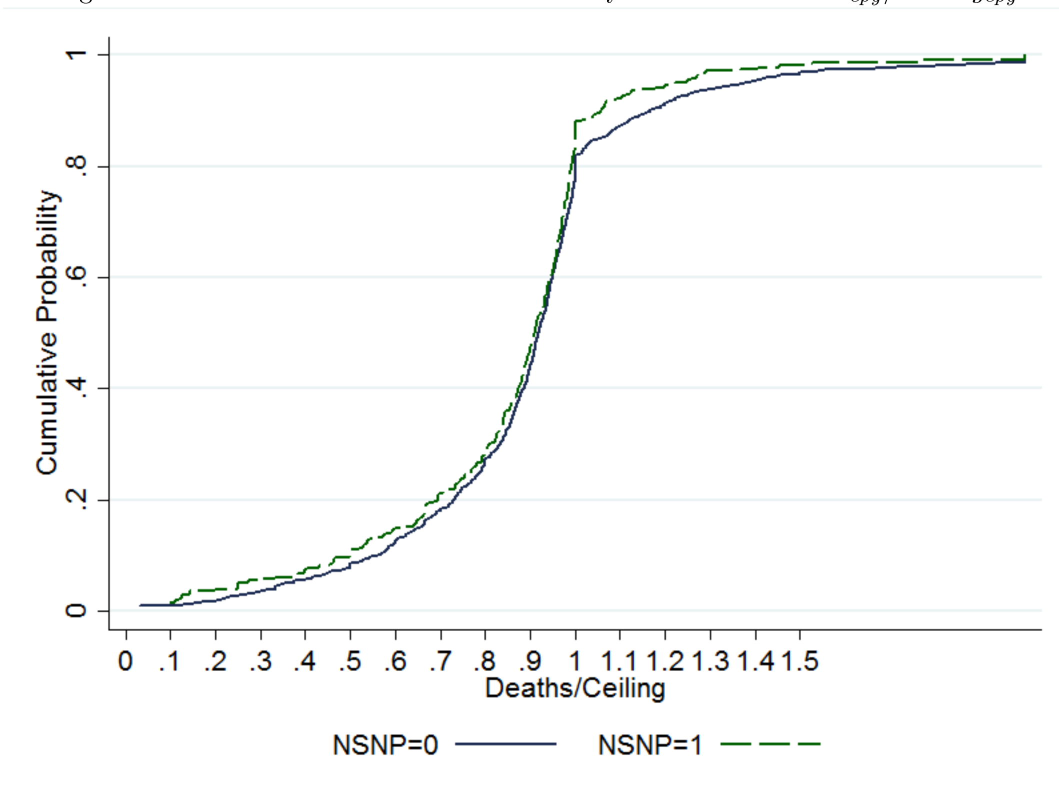

Notes: Each line shows a cumulative density function for the ratio of reported deaths to government-mandated ceiling, $Deaths_{cpy}/Ceiling_{cpy}$. The solid line employs data from province-year observations where "No safety, no promotion" legislation had not been passed (NSNP=0), while the dashed line employs observations where such legislation was in place (NSNP=1).

Figure 7: Relationship between deaths-to-ceiling ratio and production costs and revenues

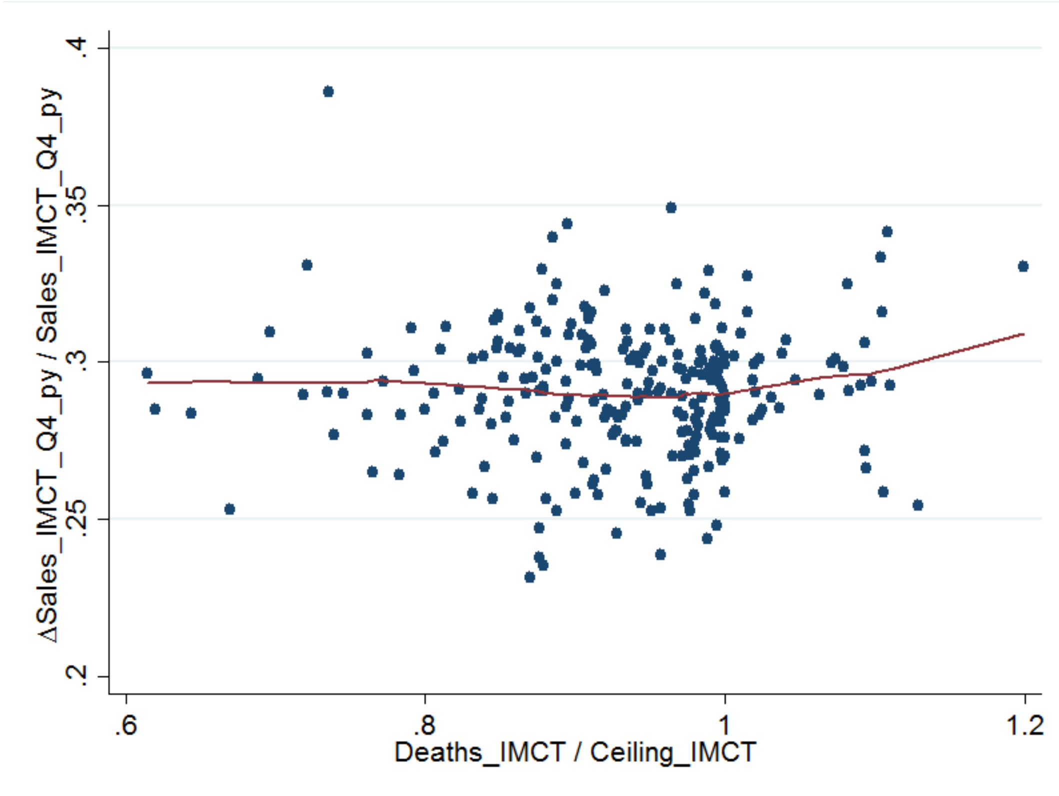

Figure 7A: Year-end deaths-to-ceiling ratio and fourth quarter sales

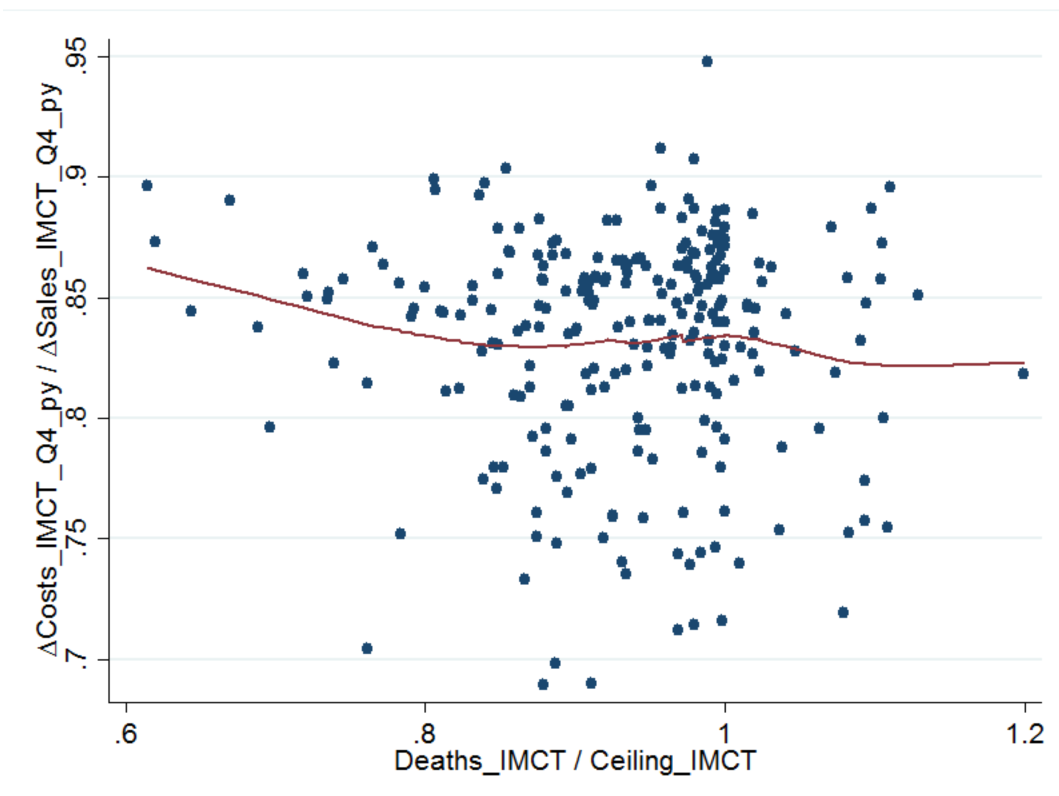

Figure 7B: Year-end deaths-to-ceiling ratio and fourth quarter cost/sales

Notes: Both panels use data from the category IMCT (industrials, non-coal mining, commercial, and trade), for all province-year observations during 2005-2012. Panel A shows the relationship between the year-end deaths-to-ceiling ratio and the fraction of annual revenues that were generated during the fourth quarter of the year. Panel B shows the relationship between the year-end deaths-to-ceiling ratio and the cost-to-sales ratio for the fourth quarter. The line in both panels shows a locally weighted smoothed regression.

Figure 8: Relationship betwen third quarter deaths-to-ceiling ratio and production costs and revenues for provinces that nearly exceed their ceilings

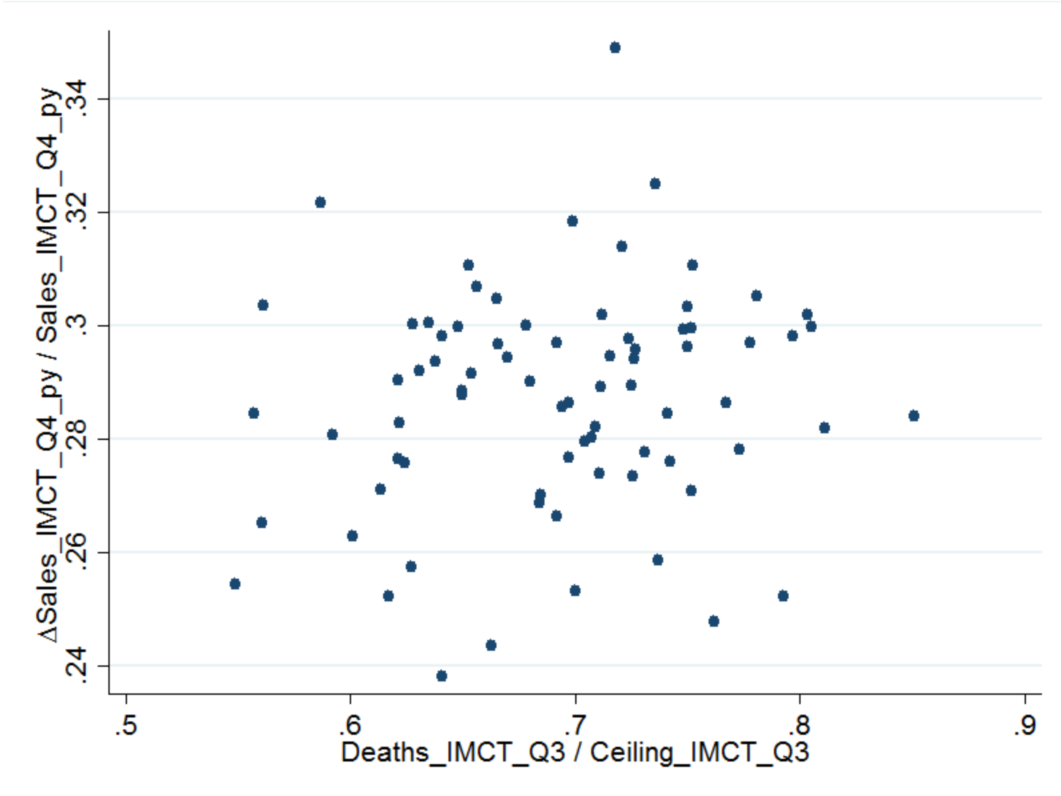

Figure 8A: Third quarter deaths-to-ceiling ratio and fourth quarter sales

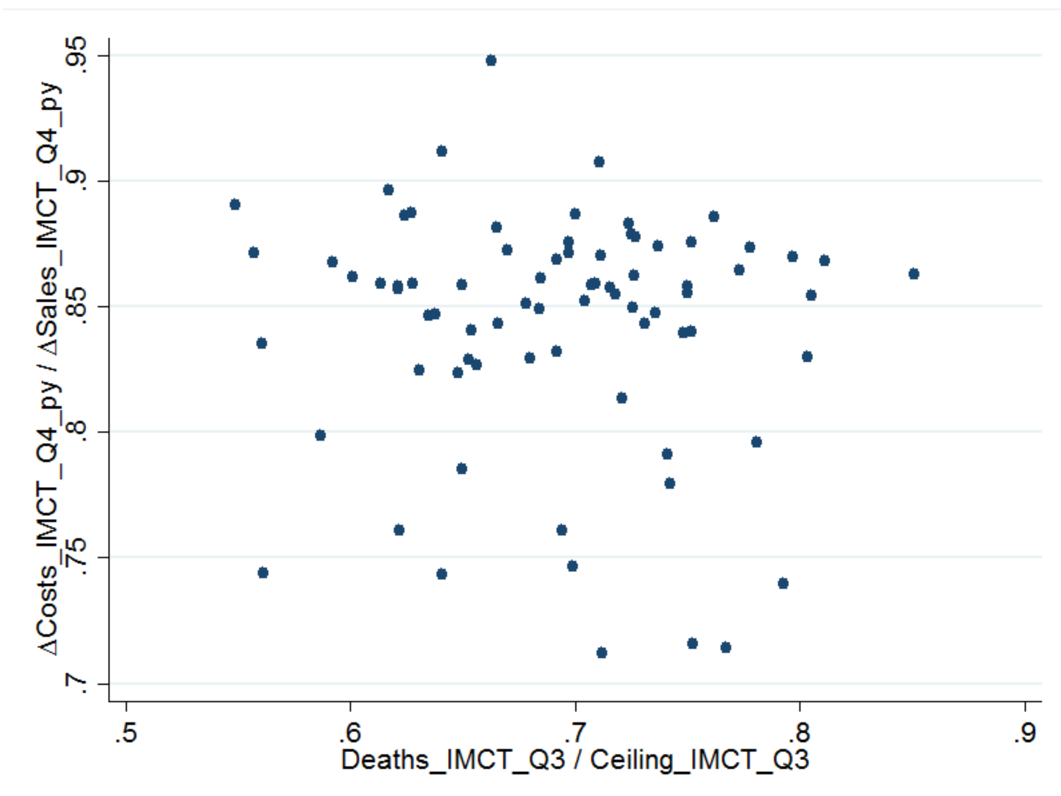

Figure 8B: Third quarter deaths-to-ceiling ratio and fourth quarter costs/sales

Notes: Both panels use data from the category IMCT (industrials, non-coal mining, commercial, and trade), for province-year observations during 2005-2012 where at year's end, $0.9 \leq Deaths/Ceiling \leq 1$. Panel A shows the relationship between the ratio of third quarter reported deaths to the year-end ceiling, and the fraction of annual revenues that were generated during the fourth quarter of the year. Panel B shows the relationship between the ratio of third quarter reported deaths to the year-end ceiling and the cost-to-sales ratio for the fourth quarter.

Figure 9: Relationship between lagged reported deaths and current death ceiling

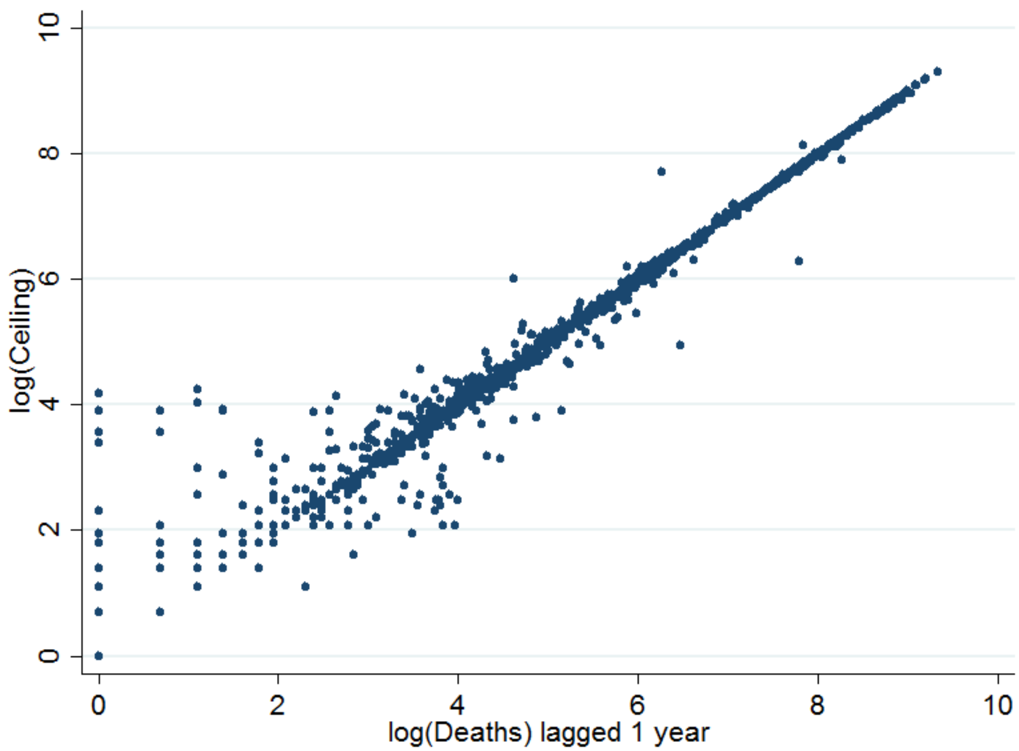

Notes: The graph uses data from all category-year-province observations during 2005-2012. It shows the relationship between the natural logarithm of reported deaths in year $y-1$ and the natural logarithm of the death ceiling in year y .

Figure 10: Pre vs. post-2004 cumulative distribution of % Δ Deaths for IMCT category

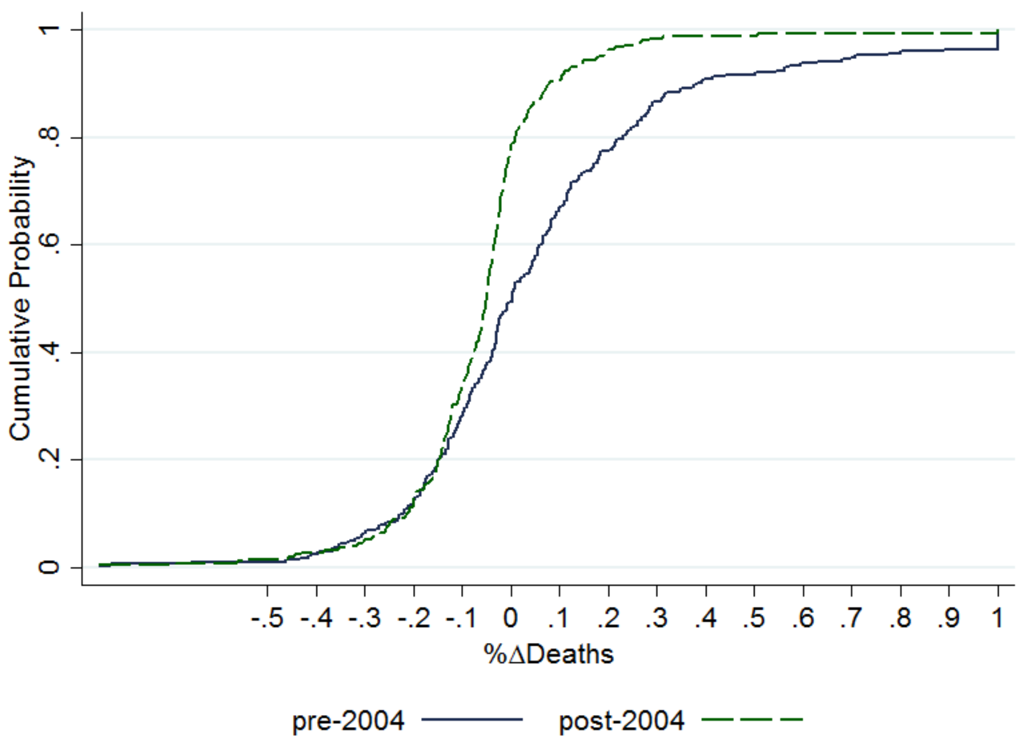

Notes: The graph depicts cumulative density functions for the annual change in accidental deaths, at the province-level, for industrial, non-coal mining, construction, and trade (IMCT) for the years 1993-1997, 1999-2003, and 2005-2012. We split the sample by 2004, the year of implementation of death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

Table A1: Instatement of "no safety, no promotion" laws by province

Province	Effective Date	Passage Date	Population in 2005 (unit: 10,000)	GDP per capita in 2005 (unit: 1 RMB)
Guangdong	01-Feb-2005	01-Feb-2005	9194	24435.02
Heilongjiang	01-Jan-2007	28-Nov-2006	3820	14434.06
Guizhou	01-Jun-2006	01-Jun-2006	3730	5051.96
Ningxia	01-Nov-2006	27-Sep-2006	596	10239
Shanxi	29-Jan-2008	29-Jan-2008	3355	12495
Yunnan	18-Sep-2008	18-Sep-2008	4450	7835
Guangxi	01-Jan-2009	27-Dec-2008	4660	8787.729
Jilin	20-Jul-2009	08-Jan-2009	2716	13348
Hainan	02-Jun-2010	02-Jun-2010	828	10871
Jiangsu	23-Nov-2010	23-Nov-2010	7475	24560
Hebei	01-Jan-2011	08-Nov-2010	6851	14782.26
Henan	03-Jan-2011	03-Jan-2011	9380	11346.5
Hunan	22-Mar-2011	22-Mar-2011	6326	10426
Inner Mongolia	04-Mar-2011	04-Mar-2011	2386	16330.82
Sichuan	11-Apr-2011	11-Apr-2011	8212	9060
Chongqing	04-Jul-2011	04-Jul-2011	2798	10982
Jiangxi	03-May-2012	03-May-2012	4311	9440
Liaoning	01-Feb-2012	30-Nov-2011	4221	18983.2
Shandong	06-Apr-2012	06-Apr-2012	9248	20096.46
Shaanxi	04-Jan-2012	04-Jan-2012	3720	9899
Hubei	No safety quota	No safety quota	5710	11431
Qinghai	No safety quota	No safety quota	543	10044.74
Shanghai	No safety quota	No safety quota	1778	51474
Tianjin	No safety quota	No safety quota	1043	35783.19
Tibet	No safety quota	No safety quota	277	9114
Xinjiang	No safety quota	No safety quota	2010	13108
Zhejiang	No safety quota	No safety quota	4898	27702.69
Anhui	No safety quota	No safety quota	6120	8675.145
Beijing	No safety quota	No safety quota	1538	45443.69
Fujian	No safety quota	No safety quota	3535	18645.84
Gansu	No safety quota	No safety quota	2594	7476.529

Notes: In response to the 2004 death ceiling policy, provinces began adopting "no safety, no promotion" (NSNP) policies that made promotion of safety regulators and other local government officials contingent on meeting the safety production target set for their province by the provincial government. The first two columns list the dates of NSNP legislation passage and implementation respectively. Both population and income figures are obtained from CSMAR Regioual Economy Database.

Table A2: Availability of death data by category and year

Category	2005	2006	2007	2008	2009	2010	2011	2012	Total
All	0	0	31	30	26	23	25	29	164
Agriculture	31	31	31	31	31	31	31	31	248
Coal	27	27	28	27	25	25	24	25	208
Fire	31	31	31	31	31	30	31	28	244
IMCT	31	31	31	31	31	31	31	31	248
Railway	29	31	30	31	30	29	31	30	241
Severe Accidents	0	0	31	30	31	31	31	31	185
Transportation	31	31	31	31	31	31	0	31	217
Total	180	182	244	242	236	231	204	236	1755

Notes: Each cell lists the number of provinces (out of a possible 31) with data available in a category-year cell.

Table A3: Coverage of municipal death data: number of observations by province-year

Province	2005	2006	2007	2008	2009	2010	2011	2012	2013	Total
Guang Dong	0	0	0	0	126	21	66	66	22	301
Gui Zhou	99	81	87	0	108	0	18	0	0	393
Hei Long Jiang	0	17	0	17	17	68	0	38	19	176
Hu Nan	0	0	0	0	0	0	14	14	27	55
Inner Mongolia	0	13	0	26	39	51	0	0	0	129
Liao Ning	0	0	14	56	28	0	0	0	0	98
Shaan Xi	0	0	11	11	0	0	0	22	0	44
Shan Xi	0	0	0	0	11	44	0	0	0	55
Si Chuan	0	0	0	0	0	0	0	22	0	22
Total	99	111	112	110	329	184	98	162	68	1273

Notes: The nine provinces in this table made public accidental deaths and death ceilings at the municipal level in each of the categories Overall, Coal, IMCT, and Road. The entry in each province-year cell indicates the number of city-category observations in that cell.

Appendix Figure A1: Histograms of Death/Ceiling Ratios by Category