Grantham Research Institute on Climate Change and the Environment

Compliance, efficiency and instrument choice: evidence from air pollution control in China

Thomas Stoerk

August 2017

Grantham Research Institute on Climate Change and the Environment
Working Paper No. 273

ISSN 2515-5717 (Online)

The Grantham Foundation for the Protection of the Environment

LSE The London School of Economics and Political Science

Compliance, Efficiency, and Instrument Choice: Evidence from Air Pollution Control in China*

Thomas Stoerk[1,2†]
[1]Environmental Defense Fund [2]Grantham Research Institute, LSE

August 21, 2017

Abstract

This research evaluates China's main air pollution control policy. In 2005, China decided on a 10% SO_2 emissions reduction goal as part of the 11th Five-Year Plan (2006-2010). I study the effect of this policy on pollution outcomes, using both the official, misreporting-prone indicator and independent NASA SO_2 satellite data in a differences-in-differences strategy that exploits variation in target stringency at the province level. I find that results from the official and the satellite data differ initially when the Chinese government lacked the ability to effectively monitor SO_2 pollution. Ultimately, however, the policy worked and reduced air pollution by 11%. The regulated provincial governments react through rhetorical compliance, measured by a unique dataset of quantified political statements, and by shutting down small, inefficient thermal units. Rhetorical compliance increases, especially before the government gained the ability to monitor SO_2 in 2008. Real compliance sets in through the shutdown of small, inefficient thermal units. Next, I compute detailed marginal abatement cost curves for SO_2 for each province in China, thus illustrating the large heterogeneity in abatement cost across provinces. I use those curves to construct the counterfactual cost-efficient allocation of SO_2 reduction targets across provinces. Using this benchmark, I find that the cost-efficient allocation would increase efficiency by 49% at the margin, by lowering marginal abatement cost from 658€/tSO_2 to 338€/tSO_2. This finding is robust to inclusion of a back-of-the-envelope measure for the marginal benefits of abatement. I conclude that a market-based allocation of SO_2 reduction targets would have doubled the efficiency of China's main air pollution control policy. Contrary to the US experience, I find that a mandate on scrubbers would reap most of those efficiency gains.

JEL Codes: Q52, Q53, D73, H11

Keywords: Air Pollution, Regulation, China, Marginal Abatement Cost, Instrument Choice

*I am indebted to Antoine Dechezleprêtre and Humberto Llavador for their continued advice and encouragement. I thank Geoff Barrows, Tim Besley, Robin Burgess, François Cohen, Denny Ellerman, Ruben Enikolopov, Roger Fouquet, Friedrich Geiecke, Matthieu Glachant, Ashley Gorst, Stephen Hansen, Martin Heger, Matthew Kahn, Felix Koenig, Gianmarco León, Kyle Meng, Niclas Moneke, Gerard Padró i Miquel, Steve Pischke, Giacomo Ponzetto, Alessandro Tavoni, Jeroen van den Bergh, Guo Xu and Yimei Zou for valuable comments and conversations. I thank Janusz Cofala and Robert Sander at IIASA as well as Nickolay Krotkov, Can Li and Chris McLinden at NASA for generous help in data access. I have further benefitted from comments by audiences at UPF, CREI, LSE, EAERE 2015, EAERE-FEEM-VIU Summer School 2015, China Economics Summer Institute 2016, EEA 2016 and the EDP Jamborees in 2015 and 2016. First version: 8th June 2014. All errors are my own.

†Contact: tstoerk@edf.org

1 Introduction

The effective design and implementation of environmental regulation is crucial for correcting environmental externalities. Traditionally, economists have analyzed environmental regulation in developed countries where technical expertise, appropriate monitoring of pollution and rule of law often allowed successful implementation of regulation. Recently, attention turned to environmental regulation in developing countries and how compliance with regulation interacts with imperfect institutions (Duflo et al., 2013; Oliva, 2015). At the same time, developing countries are often more severely affected by the most important environmental externalities such as air pollution. According to the latest WHO estimates, air pollution is responsible for one in eight of total global deaths, which makes 7 million deaths a year (WHO, 2014). One country which is particularly struck by air pollution is China. As development has soared, so has air pollution. Recent air pollution levels in Northern China are as severe as those in London at the height of the Industrial Revolution[1].

This study provides the first empirical evaluation of China's total emissions control target in the 11th Five-Year Plan (FYP) from 2006 to 2010. In an effort to bring down air pollution, the Chinese government decided to limit the total emissions of sulphur dioxide (SO_2) by 10% relative to 2005 baseline levels. The national limit was later split into widely varying reduction targets for each province. This study evaluates the SO_2 reduction policy along three margins: First, did the policy improve SO_2 pollution outcomes? Second, how did the regulated provincial governments comply? Third, how efficient was the policy compared to a counterfactual market-based policy instrument?

The 11th FYP marks a turning point in environmental policy-making in China; it is considered 'the most environmentally ambitious document in the history of the Communist Party' (Watts, 2011). When the policy was passed in 2005, China's regulatory agency was the weak State Environmental Protection Administration (SEPA). SEPA did not have access to reliable SO_2 pollution data in 2005 and had to implement the regulation based on limited information from SO_2 emission statistics. This situation changed drastically in 2008, when the central government upgraded SEPA to become the Ministry of Environmental Protection (MEP), gave it high-level political backing and allowed it to track SO_2 pollution at the source and independent of provincial governments (State Council, 2007).

This empirical setting is insightful for several reasons. While data on the implementation of the SO_2 reduction policy is only available at the province level, my setting is unique because of the availability of real pollution data in the period before the Chinese government could monitor it. This is due to coincidence: in late 2004, just before the start

[1] The level of total suspended particles (TSP) in London in 1890 was just beyond 600 $\mu g/m^3$ (see Figure 3 in Fouquet, 2011). Similarly, TSP levels North of China's Huai River were just beyond 600 $\mu g/m^3$ up until 1980 (see Figure 2 in Chen, Ebenstein, et al., 2013).

of the policy, NASA launched the *EOS-Aura* satellite that provides an independent and reliable data source for SO_2 pollution in China. Additionally, rich outcome variables are available for China to track the behaviour of provincial governments along multiple margins, including quantified political statements about current and past political priorities at the province-year level. This allows me to study whether regulated agents react differently when monitored appropriately. Lastly, further advantage to studying air pollution regulation in China is the availability of micro-level data on the cost of SO_2 abatement in each province.

Firstly, I evaluate whether the SO_2 control policy improved pollution outcomes despite the lack of regulatory capacity at the start of the 11th FYP. Exploiting variation across provinces in a differences-in-differences (DID) specification, I recover the causal effect of the SO_2 control policy on both reported pollution, using the official indicator, and on real pollution, measured through NASA satellite data. I then study whether the effect of the SO_2 reduction target differs at the county level according to the initial distribution of pollution within the province. Finally, I investigate the timing of the effects of the policy to study compliance over time.

I find that the policy was a success according to the official SO_2 emissions indicator: a one-standard deviation increase in the stringency of the reduction target leads to a statistically significant 6% decrease in SO_2 emissions. The effect is similar based on the SO_2 NASA satellite data at an effect size of 11% per standard deviation. Withing a province, the estimated effect is stronger for counties that were initially more polluted. However, the timing of the estimated differs when comparing the official and the satellite data: according to the official data, provinces with a higher reduction target started to decrease their SO_2 emissions immediately and monotonically after the start of the 11th FYP. The satellite data, by contrast, paint a different picture: according to the satellite data, SO_2 pollution in provinces with higher targets only started to decrease in 2008, after the government gained the ability to monitor SO_2 pollution. A plausible reason for these differences in timing is misreporting due to moral hazard. The evidence, however, is only suggestive because these differences are not statistically different.

Next, I evaluate the mechanisms through which the regulated agents comply with the policy. I study their reactions along two margins: rhetorical and real compliance. Given the initial period of moral hazard when the Chinese government could not monitor SO_2, I first ask whether provincial governments reacted to the SO_2 reduction targets by changing their political rhetoric. I construct a unique province-year dataset of political statements by the governments of each province in China. These statements discuss the political priorities of each provincial government, both for the preceding year and the year to come. This allows me to study not only the content of the provincial governments' rhetoric but also the timing. I use these statements to quantify the importance of air pollution in a province's political rhetoric. To study real compliance, I collect data on the two most economic measures to reduce SO_2 at the province level. These are the installation of desulfurization equipment (scrubbers) in power plants and industrial facilities, and the shutdown of small, inefficient power plants.

I find that provincial governments strongly adjusted their rhetoric to the political goals from the center: a one-standard deviation increase in the province's SO_2 reduction target leads to a 30% increase in political statements on air pollution, mainly driven by mentions of sulfur. A placebo test using other goals from the 11th FYP shows that provinces specifically adapt their rhetoric to the SO_2 targets. Regarding the timing of the political statements, I find that the increase in political statements is most pronounced in statements that praise past work on air pollution. This is particularly visible for the period just before the government could monitor air pollution. In light of the difference between the official and the satellite SO_2 data, this provides further suggestive evidence for strategic misreporting. Ultimately, however, the SO_2 reduction policy worked, and compliance became real. Among the abatement measures available to provinces, I find that real compliance for provinces with higher SO_2 reduction targets relied less on the installation of desulfurization equipment, or on an earlier installation of this equipment. Instead, I find that the main driver for bringing down SO_2 pollution was the shutdown of small, inefficient power plants.

To study the efficiency of the policy design and quantify the gains from trade across different policy instruments, I compute detailed SO_2 marginal abatement cost (MAC) curves for each province in China. These curves show the large heterogeneity in SO_2 abatement cost across the provinces of China. Based on the MAC curves, I find that the Chinese government did not equate marginal abatement cost across space. Instead, the reduction targets favoured coastal provinces in the East where abatement costs are higher. This is consistent with a tale-of-two-cities story, in which China would develop amenity-based consumer cities along the coast, while maintaining a base of polluting manufacturing in its interior (Kahn, 2006; Zheng and Kahn, 2013). Using the MAC curves, I construct the counterfactual cost-efficient allocation of SO_2 reduction targets across provinces needed to achieve the 10% SO_2 reduction target. This allows me to study the gains from trade from moving from a command-and-control regulation to a market-based allocation of SO_2 reduction targets. I find that moving to a cost-efficient allocation would increase efficiency by 49% at the margin, lowering marginal abatement cost from 658€/tSO_2 to 338€/tSO_2.

As a robustness check, I quantify how sensitive these estimates are to the inclusion of a measure for the marginal benefits of reducing air pollution. I use the back-of-the-envelope method employed by Oliva (2015) to construct a measure for the marginal abatement benefits of reducing SO_2 emissions at the province level. This method proceeds in 3 steps: (i) how does the SO_2 control policy change pollutant concentrations?, (ii) what health effects do the changes in pollutant concentrations cause?, and (iii) what is the monetary value of those health effects? I find that the gains from trade from moving to the cost-efficient allocation are robust to including the benefits of air pollution abatement: at the margin, welfare increases are 45%, a similar magnitude as the cost-based efficiency increases.

Overall, my empirical findings show that China's flagship air pollution control policy worked. SO_2 pollution, as measured through independent satellite data, decreased by more than 10% as a result of the policy. This sucess, however, appears to hinge on the Chinese government's ability to fully monitor SO_2 pollution. Before the 2008 changes in monitoring, compliance is mostly rhetorical. This finding provides support for a theoretical literature pioneered by Jean-Jacques Laffont suggesting the importance of regulatory capacity for the working of even simple policies (for an overview, see Laffont, 2005). Assessing the efficiency of the air pollution control policy, I find that the Chinese government fails to exploit the large heterogeneities in abatement cost across provinces. A more market-based allocation of SO_2 reduction targets, such as through an emissions trading scheme, would double allocative efficiency.

My analysis thus suggests that China currently faces problems in its air pollution control that are similar to the problems that today's developed countries faced in the past. Air pollution control is initially plagued by monitoring difficulties, and works in general, though it is not initially cost-efficient. Compared to environmental regulations in other emerging countries such as India and Mexico that have been evaluated in the literature (Duflo et al., 2013; Oliva, 2015), it is likely that China's better pollution monitoring capacities towards the end of the implementation explain the success of the policy in reducing air pollution.

The paper is organised as follows: Section 2 discusses the related literature. Section 3 describes the setting and discusses environmental governance in China as well as the 2008 changes in the SO_2 monitoring. Section 4 explains my data sources, while Section 5 contains the empirical analysis. Section 6 constructs detailed marginal abatement cost curves at the province level and, in combination with a marginal benefit measure, computes the gains from trade across different policy instruments. Section 7 concludes.

2 Related literature

The two main reasons to study air pollution in China are the magnitude of the problem and the possibility to alleviate it through policy. Air pollution in China causes enormous health costs as shown by literatures in both economics and health. Chen, Ebenstein, et al. (2013), for instance, use a natural experiment to find that one coal-subsidy alone led to the loss of 2.5 billion life-years in Northern China. Epidemiological studies summarized in Yang et al. (2013) give the same sense of magnitude: they find air pollution to be the fourth most important health burden in China. In monetized terms, the health cost amount to 1.2 to 3.8% of GDP (World Bank and State Environmental Protection Administration, 2007). Air pollution furthermore induces direct productivity losses Chang et al. (2016). At the same time, Jia (2014) has shown in a convincing causal setting that pollution is a side effect of political incentives. A large literature in urban economics, political economy and environmental law backs this conclusion (Almond et al., 2009; Wang, 2013; Zheng and Kahn, 2013; Zheng, Sun, et al., 2014). It therefore seems possible to mitigate air pollution through the right combination of policies and incentives.

My study is the first to provide a causal empirical evaluation of the total emissions control policy in the 11th FYP. Despite the huge costs of air pollution in China, there has been no empirical evaluation of China's flagship air pollution control policy. Evaluation so far has come in one of two guises: through detailed narrative accounts of the changes (Hao et al., 2007; Schreifels, Fu, and Wilson, 2012) or through model-based studies in atmospheric science (Lu et al., 2010; Wang, Jang, et al., 2010a,b). Additionally, I am amongst the first to evaluate any environmental policy in China. The main other study I am aware of is Kahn, Li, and Zhao (2015), who analyze water pollution regulation. Their findings complement the empirical part of my research and allow for a richer interpretation of the role of pollution monitoring.

This research also contributes to two distinct literatures in environmental economics. The first is a nascent empirical literature that evaluates the design and implementation of environmental policies in imperfect settings, while the second literature studies the efficiency of the design of regulation through comparison of different policy instrument.

In the first literature, Duflo et al. (2013) study how the emissions of regulated industrial firms in India respond to changes in effective monitoring. Using a field experiment, they find that improved auditing leads regulated firms to reduce pollution, most notably through water pollution. Relatedly, Oliva (2015) provides a static snapshot of an air pollution regulation gone wrong. She documents the ineffectiveness of emissions regulation for private cars in Mexico City. Using a statistical test to detect misreporting, she shows that car owners often fail to comply and instead pay a bribe to circumvent regulation, leading to a private benefit to car owners at great damage to the public. Hansman, Hjort, and Leon (2015), by contrast, study the design of a reform aimed at preserving fish stock in Peru. They show that the piecemeal nature of the reform led to unforeseen side effects for air pollution outcomes.

This paper sheds light on whether a simple command-and-control policy can be effective in a setting in which the government has a very low regulatory capacity initially. Due to the availability of independent satellite data on pollution throughout, I can study the behaviour of regulated agents in terms of likely misreporting, pollution outcomes and rhetorical compliance both before and after the Chinese government gains the ability to monitor pollution in 2008. Understanding how the affected agents comply with an environmental regulation under moral hazard is central to policy implementation and thus of great importance for contexts beyond China. The analysis of a command-and-control regulation in an emerging country such as China contributes furthermore to the existing literature on the effectiveness of air pollution control, which has focused almost exclusively on developed countries such as the US (Auffhammer and Kellogg, 2011; Chay and Greenstone, 2005; Chay, Dobkin, and Greenstone, 2003; Henderson, 1996)[2]. Additionally, while research on air pollution regulation has led to clear findings for pollutants such as TSP, the results for SO_2 are more mixed (Greenstone, 2004; Greenstone and Hanna, 2014).

The second literature to which I contribute studies the design of environmental regulation and the use of different policy instruments. This line of research asks a normative question - what would be the efficient allocation of SO_2 reduction targets? - and uses this benchmark to assess the efficiency losses of actual policy. Most contributions come from analyses of air pollution control regulation in the US, most notably through the US SO_2 mitigation efforts (for an overview, see Schmalensee et al., 1998 and Stavins, 1998), but a small number of studies exist from other countries such as Chile (Montero, Sanchez, and Katz, 2000). My work relates most closely to Carlson et al. (2000) and Oates, Portney, and McGartland (1989).

Carlson et al. (2000) compute marginal abatement cost curves for SO_2 for the electricity sector in the US to quantify the welfare gains from trade of moving from command-and-control regulation to SO_2 emissions trading. While those cost savings are large, they are surprisingly lower than anticipated ex ante. Another closely related paper is Oates, Portney, and McGartland (1989), who compare the efficiency of incentive-based regulation against command-and-control regulation to control air pollution in Baltimore. They find that a well designed command-and-control regulation can deliver pollution reductions at a welfare cost that can be lower than a comparable incentive-based regulation. This conclusion, however, only applies when that command-and-control regulation is informed by the marginal cost of abatement rather than by political considerations. Taken together, these studies show that while moving from a command-and-control regulation to a market-based regulation is generally seen as increasing the efficiency of the regulation (Schultze, 1977),

[2]One exception are Greenstone and Hanna (2014), who investigate the effect of regulation on air and water pollution in India.

whether this is so is an empirical question that depends on the particular case of the regulation under consideration. In particular, it will depend on whether the cost of abatement was taken into account when designing the the command-and-control allocation of targets[3].

I estimate detailed marginal abatement cost curves for SO_2 for each province in China in 2005, contributing to the few studies that estimate full marginal abatement cost curves in environmental economics in general (Gollop and Roberts, 1985; Carlson et al., 2000 and Abito, 2012)[4]. In particular, this study is the first to derive comprehensive marginal abatement cost curves at the province level in China. Two earlier contributions by Tu and Shen (2014) and Li, Wu, and Zhang (2015) provide interesting analysis in this direction but both studies rely on modelling in addition to microdata and only compute partial snapshots rather than full MAC curves. I use the marginal abatement cost curves to predict the counterfactual allocation of SO_2 reduction targets across provinces in China. This allows me to quantify the efficiency gains from trade from moving from the actual command-and-control allocation of SO_2 reduction targets to the cost-efficient allocation. Furthermore, I investigate the robustness of this quantification to the inclusion of a measure for the marginal abatement benefits. Overall, I contribute by studying whether the design of air pollution control regulation in a developing country such as China differs from the experience of developed countries like the US, and show possible differences in detail.

3 The SO_2 reduction policy in context

This section provides the context around the 11th Five-Year Plan (FYP) and China's flagship air pollution control policy that is the focus of this research. Environmental governance in China has undergone a rapid transformation in the last two decades. Until 2005, economic growth was the defining development paradigm. Environmental policies, where they existed, were paper tigers: they lacked political support from the central government and were rarely enforced. 2005 marks the turning point with a Five-Year Plan that 'was the most environmentally ambitious document in the history of the Communist Party' (Watts, 2011). The following paragraphs sketch how this change can best be seen as a change in political will rather than a change in formal laws.

[3]Further related studies are Newell and Stavins (2003), who propose a simple way of calculating the efficiency gains from market-based policies in situations with limited data, and show an application to NOx control in the US, and Antle et al. (2003), who investigate spatial heterogeneity in the abatement cost for carbon sequestration in the US and show that different policy instruments vary in efficiency by a factor of 5.

[4]Partial estimates of marginal abatement cost curves for compliance with air pollution control regulation are further reported in Hartman, Wheeler, and Singh (1997), Becker and Henderson (2001), Keller and Levinson (2002) and Becker (2005).

Before the policy Laws regulating SO_2 emissions have existed in China since 1998, when the State Council approved the establishment of the 'Two Control Zones', a policy to address acid rain and SO_2 emissions (McElwee, 2011). Enforcement of this policy was intensified in 2000, but has remained constant since. Implementation of SO_2 policies, however, still encountered great difficulties (Gao et al., 2009), as China's overall development strategy remained firmly rooted in economic growth. Existing environmental policies overall, for instance on energy efficiency, were left underfunded by the central government (Gao et al., 2009; Lin, 2007). And while the 10th Five-Year Plan included a nationwide goal to reduce SO_2 emissions by 10%, it did not have political backing, and failed to induce SO_2 emissions reductions (Schreifels, Fu, and Wilson, 2012), possibly because of a lack of incentives for meeting the targets (Wang, 2013).

First change in the political outlook of the central government came in 2003, when President Hu - a hydroengineer - and Prime Minister Wen - a geologist - took power. The 'scientific development' paradigm, which emphasized environmental protection alongside economic growth, started to substitute for economic growth. Environmental governance, however, was still weak.

2005: The 11th Five-Year Plan (2006-2010) Amidst the increasing political will to implement and enforce environmental policies, the general directions of the 11th FYP started being discussed as early as mid-2003, and probably ended by 2004 (Xu, 2011). During the National People's Congress in March 2006, the 11th FYP was presented in its final form and included emissions control targets for air pollution (SO_2) and water pollution (chemical oxygen demand, or COD) as well as a target on energy efficiency. Concurrently, environmental governance started being taken seriously, when the once powerless SEPA successfully stopped hundreds of billions of Yuan of industrial investment on environmental grounds at the beginning of of 2005. This radical action came as a surprise to Chinese society (Gao et al., 2009). In March 2008, the SEPA's new political authority was formalized when SEPA obtained full rank in the State Council and received ministry status as the MEP (McElwee, 2011). This is the period that I analyze in this research.

The air pollution control target consisted of a 10% SO_2 emissions control target for China as a whole. This reduction target was handed down to the provinces in May 2006 at the latest, when SEPA - with high-level political backing - signed formal, binding reduction targets with the provincial governments (Gao et al., 2009; Xu, 2011)[5]. These reduction targets were given the highest political priority, paralleled only by mandates on growth, social stability and the one-child policy (Wang, 2013). Compliance with environmental targets entered the performance evaluation of local leaders (Moser, 2013). Provincial SO_2

[5]The 31 province-level administrative units include the four municipalities directly under the central government (Beijing, Chongqing, Shanghai and Tianjin) and the five autonomous regions (Guangxi, Neimongol, Ningxia, Tibet and Xinjiang).

reduction targets, in particular, were made a veto target: failure to comply would nullify all other performance achievements of a provincial leader (Kahn, Li, and Zhao, 2015; Xu, 2011). The reduction targets specified reductions in SO_2 emissions from 2006 to 2010 with 2005 as the baseline. These are the reduction targets that I use in this study. Figure 1 shows that these targets vary considerably, mandating reductions from 0% to more than 25%.

[INSERT FIGURE 1 ABOUT HERE]

SO_2 emissions data in China in 2005 were of bad quality, and misreporting-prone. A province's SO_2 emissions were calculated by aggregating the physical quantity of coal used in a province in a given year, and by then multiplying this quantity with SO_2 emissions factors depending on the sulfur-intensity of the type of coal used. The political authority to compile these data rested with the provincial governments. While this procedure yields only coarse aggregate data at best (Guan et al., 2012), it can also be corrupted to produce the desired data. Anecdotal evidence and extensive field work in several provinces confirm that misreporting of SO_2 emissions data was prevalent in the first 2 years of the SO_2 control policy (Song et al., 2015). The likely reason why the central government decided to use emissions rather than other pollution data is that the network of *in situ* measurement stations was not of sufficiently high quality and quantity in 2005 to track SO_2 pollution outcomes on the ground.

Up to 2005, provincial governments in China had little incentive to control air pollution. A province in China is governed by a pair of provincial leaders, the governor and the party secretary. Provincial leaders are career officials who are appointed in a top-down manner. As career officials, they are often positioned outside their native provinces and move frequently as a results of promotion or demotion. Cadre regulations stipulate a maximum term length of 5 years for both governors and party secretaries (Kahn, Li, and Zhao, 2015). This rule is mostly respected: on average, provincial leaders have held office for just two years, with only a small number exceeding a tenure of 5 years. As a consequence of these frequent changes, provincial leaders are mainly concerned with furthering their own career in the short period that they hold office. Following the regulations in the past, this meant relying on pollution-intensive, quick-and-dirty GDP growth (Chen, Li, and Zhou, 2005; Jia, 2014; Li and Zhou, 2005). Incentives for provincial leaders to fix air pollution are further weakened by vested interests in polluting enterprises, of which local governments are often major shareholders (Gao et al., 2009). As provincial leaders got promoted to posts elsewhere, air pollution therefore remained unfixed.

2008: Changes to SO_2 monitoring In 2007, the State Council passed a law that fundamentally changed the regulator's capacity to monitor SO_2 pollution (the law is known as the 'Reduction of the Three Ways'). The core of the law was twofold: to change the politics of SO_2 data collection to avoid tampering with data at the political level, and to install appropriate monitoring equipment and ensure frequent statistical inspections on the ground (State Council, 2007).

On the political side, reporting was taken from provincial governments and put directly under the political control of the MEP. The MEP, in turn, directly reports to the State Council. On the ground, SO_2 measurement stations were build in pollution hotspots and the number of environmental monitoring officials was increased by 17% (Song et al., 2015). Key industrial polluters for each prefecture had their SO_2 emissions tracked on site. By May 2008, uninterrupted automatic monitoring devices with data feeds directly into the local environmental agency were used for this purpose (McElwee, 2011; MEP, 2008a,b). All changes became effective in July 2008 at the latest (Song et al., 2015). These changes mark a landslide for environmental governance in China: a nearly powerless agency was upgraded into a powerful ministry. A ministry given the tools to ensure it was going to be a force to be reckoned with.

The 2008 changes naturally divide the analysis into three periods: (i) before the policy; (ii) SO_2 reduction targets, but little monitoring; (iii) SO_2 reduction targets, and comprehensive monitoring. These periods will guide my subsequent empirical estimations to evaluate the policy.

4 Data

I compile a unique dataset that allows me to study the effect of the SO_2 control policy along two margins: Firstly, to evaluate the effect of the policy on pollution outcomes, I use two different data sources: (i) the official, misreporting prone SO_2emissions indicator and (ii) independent satellite data from NASA. Secondly, I evaluate the reactions by the regulated provincial governments. I divide these reactions into rhetorical compliance and real compliance. To measure rhetorical compliance, I build a novel dataset of political reports, which I quantify. Behaviour related to real compliance is based on official data sources on SO_2 abatement measures. I firstly describe the pollution data, and then the data related to the behaviour of the regulated provincial governments.

4.1 Data on pollution outcomes

Official SO_2 data The SO_2 control policy relies on SO_2 *emissions* as the official indicator. Note that this is a proxy indicator for the ultimate goal of reducing pollutant *concentrations*. It is likely that the central government chose this indicator due to a combination of a legacy of state-planning that focused on total emissions control and a lack of suitable SO_2 concentrations data from *in situ* measurement stations. I use the data reported in the *China Energy Databook* (Fridley, Romankiewicz, and Fino-Chen, 2013), who compile the official SO_2 emissions data from the statistical yearbooks.

The official SO_2 emissions data gives at best a noisy picture of true SO_2 emissions (Guan et al., 2012), and anecdotal evidence and research based on fieldwork suggest that the incentives before the 2008 policy changes led to severe misreporting (Song et al., 2015). The literature more generally has also noted misreporting of air pollution data (Andrews, 2008; Chen, Jin, et al., 2012; Ghanem and Zhang, 2014). However, this kind of misreporting consists of relabeling around the politically sensitive *Blue Sky Day* threshold in air pollution index data. It is thus tangential to misreporting in SO_2 emissions data, affects the mean of the data very little, and there are signs that this particular kind of misreporting has come to an end with the introduction of $PM_{2.5}$ measurements from 2012 onwards (Stoerk, 2016). The potential source of misreporting was therefore removed after the MEP started collecting SO_2 pollution data directly in 2008.

SO_2 satellite data To measure real SO_2 pollution, I make use of the uniqueness of my empirical setting. In August 2004 - just before the start of the SO2 control policy - NASA launched a satellite with the Ozone Monitoring Instrument (OMI). In lay terms, OMI is an instrument that captures images of Earth from space at different wavelengths. Post-processing trough extraction algorithms produces SO_2 vertical columns of high precision that became available in 2014[6]. It is unlikely that the Chinese government would have had access to the satellite data during the policy, an these data were not used in the official evaluation of the policy. Dates with cloud cover can lead to missing values over individual pixels, but this is not generally considered a first-order problem (Krotkov et al., 2016)[7]. All in all, the NASA SO_2 satellite data are a very good proxy for ground-level SO_2 emissions. To illustrate, Figure 2 shows the cross-section in January 2006.

[INSERT FIGURE 2 ABOUT HERE]

Relation between both data sources SO_2 emissions decay in a span of 4-36 hours (Fioletov et al., 2015; He, 2012). Since the satellite data capture Earth daily, they represent a snapshot of SO_2 pollution on that day. Given the quick decay, this prevents leakage from confounding the outcome and allows me to capture local rather than transported SO_2 emissions. NASA's retrieval algorithm produces four different data products, each of which corresponds to SO_2 pollution at different levels of altitude in the atmosphere

[6]The new algorithm significantly improved the precision of the extracted vertical column densities and removed a number of biases compared to the earlier data product that was available from OMI (Krotkov et al., 2016). Note further that the OMI data itself have a detection threshold that is two magnitudes smaller than earlier satellite data and can thus enable the detection of SO_2 pollution from human activity in the lowest part of the atmosphere (NASA, 2014).

[7]Furthermore, I aggregate daily values to at least monthly, and pixels of not bigger than 0.25 degrees by 0.25 degrees latitude-longitude to counties and provinces, further reducing the magnitude of the potential problem due to clouds.

(NASA, 2014). For this analysis, I use the lowest level at an altitude of 900m above ground, for two reasons: firstly, because this is the best proxy for anthropogenic emissions sources, and secondly, because a lower altitude further minimizes transportation. Secondly, the precision of the satellite images is high enough to identify individual sources of pollution that produce as little as 30kt of SO_2 anually (Fioletov et al., 2015). I aggregate these daily cross-sections to the province-month and the province-year levels. Annual changes in the SO_2 satellite data can therefore be expected to be mimicked closely in the official statistics[8].

4.2 Data on reactions by the regulated provincial governments

This section describes the data sources that I use to study the reactions of the provincial governments along two margins: (i) rhetorical compliance and (ii) real compliance.

Rhetorical compliance I build a comprehensive dataset of political statements by each provincial government to study whether the regulated provincial governments respond to the SO_2 reduction targets by mimicking the central government's rhetoric. In China, each provincial government has the obligation to issue a government work report every year. This report is publicly delivered by one of the two highest ranked officials in the province, the party secretary or the governor. Each report contains information on the provincial government's activities and achievements. These reports are divided into two parts: Part 1 discusses the provincial government's work and achievements in the preceding period, while Part 2 discusses the work in the period to come. This unique setting allows me not only to investigate political rhetoric in general, but to also specifically investigate rhetorical responses relating to past and future achievements in response to the SO_2 reduction targets. Figure 6 shows an example of a government work report from Liaoning province in 2003, delivered by then-governor Bo Xilai.

[INSERT FIGURE 6 ABOUT HERE]

To measure the provincial government's political attention towards air pollution for a given province-year, I scan each province's government work report for the years 2002-2010 and construct a variable that is equal to the number of occurrences of keywords related to air pollution[9]. Keywords were chosen from a technical document on urban air pollution in developing countries (GTZ, 2009) as well as from China-specific air pollution articles, webpage entries and blog posts in March 2014 from *China Daily*, *Global Times*, *Beijing*

[8]A cross-check of known ground-level emission sources with SO_2 OMI satellite data revealed a correlation of 0.91 (Fioletov et al., 2015).

[9]8 reports are missing in 2002 and 1 report is missing in 2003.

Review and *Jinyang Yangcheng Evening News*. To further rule out cherry-picking of keywords, I have defined the list of keywords as widely as possible. Figure 7 contains the raw count of keywords over time, offering two take-aways. Firstly, there is a distinct increase in air pollution related keywords: Mentions of air pollution increase by more than 400% during the 11th Five-Year Plan (2006-2010). Secondly, one keyword drives this increase: the specific keyword 'sulfur'. This keyword is directly related to the provincial SO_2 reduction targets from the 11th Five-Year Plan. The outcome variable therefore appears to capture relevant political statements[10].

[INSERT FIGURE 7 ABOUT HERE]

Real compliance Based on my own calculations of the marginal cost of SO_2 emissions abatement (see Section 6), the installation of desulfurization devices in existing industrial and power plants (scrubbers), fuel-switching to better quality coal and the shutdown of small, inefficient thermal units are the main margins by which the provincial governments could reduce SO_2 emissions over the relatively short time horizon of 5 years. I collect data on both the timing and the quantity of installation of desulfurization devices at the province level during the 11th Five-Year Plan (2006-2010). Furthermore, I compute the number and capacity of small thermal units that were shut down by 2010 in each province. The latter are based on a planning document from 18th January 2008, in which the MEP asked the provincial governments to submit a concrete proposal for the thermal units to be shut down over the following two years. Anecdotal reports suggest that these shutdowns did indeed take place. All data were compiled from sources available through the data center of the MEP (datacenter.mep.gov.cn). I do not analyze the sulfur content of coal used at the province-year level since reliable data is unavailable.

5 Empirical analysis

5.1 Empirical strategy

Baseline specification I use the following difference-in-differences model to investigate whether a higher SO_2 reduction target led to a relatively stronger decrease in SO_2 emissions:

$$y_{pt} = \beta_0 + \beta_1 Reductiontarget_p \times D(Post)_t + \beta_2 Reductiontarget_p + \sum_{t=1}^{T} \beta_{3t}\gamma_t + \alpha_p + u_{pt} \quad (1)$$

[10]One valid concern with the raw count of air pollution keywords is that the length of the government work reports might have changed over time or across provinces. This could introduce a bias in my measure. To show that this is not the case, Figure A.4 in the Appendix divides the number of keywords by the length of the report and finds a very similar pattern. I prefer to use raw data because of its clearer interpretation. All regressions results are robust to using the adjusted number of keywords. These results are available upon request.

The outcome y_{pt} for province p at time period t is either the official SO_2 emissions data that was used by the central government to assess the policy or the independent satellite SO_2 data. The variable $Reductiontarget_p$ is the provincial SO_2 reduction target and captures the cross-sectional variation shown in Figure 1. Policy variation over time is captured in the indicator $D(Post)$ that takes on the value 1 from 2006 onwards. α_p are province fixed effects. The estimate for β_1 gives the causal effect of an increase in the target stringency by one unit for the whole period of the 11th Five-Year Plan, from 2006 to 2010.

Identification The DID specficiation exploits cross-province variation in the stringency of the SO_2 reduction target to estimate the causal effect of the pollution control policy. Identification relies on a combination of three factors: (i) common trends in the outcome variables prior to the SO_2 control policy, (ii) a sharp deviation from those trends following the policy changes, and (iii) the absence of forward-looking considerations that would also explain SO_2 abatement efforts by the provincial governments in 2006.

Provincial governments can rely on three main channels to bring down pollution. All of which are quick to implement, in particular for a country like China: fuel-switching to higher quality coal with a lower sulfur content, installation of desulfurisation devices, and the shutdown of small, inefficient thermal units. It is therefore reasonable to expect an immediate effect of the SO_2 reduction policy on pollution outcomes. Lu et al. (2010) provide evidence that abatement measures started immediately: power plants already started to switch to better quality coal in 2005 (with a sulfur content reduced by about 20%compared to the preceding year), and flue-gas desulfurization technology doubled from below 10% to more than 20% of all operating power plant capacity.

Common pre-trends and the timing of effects are empirically testable, and I show below that these conditions are fulfilled. Consideration (iii) is not directly testable, but supporting evidence shows that it is likely fulfilled. While the exact algorithm used by the Chinese government to allocate the targets is unknown, and it is unlikely that the targets were distributed randomly, random allocation is not needed for my identification. Instead, I only need that the allocation of SO_2 reduction targets across provinces was independent of forward-looking considerations that would explain SO_2 pollution reductions by a province independent of the SO_2 reduction targets.

The official statement by the State Council on how the target distribution would have taken place mentions a whole array of factors that were used to determine the allocation of targets for a province (State Council, 2006): (i) environmental quality and environmental capacity, (ii) current pollution levels, (iii) level of economic development, (iv) SO_2 mitigation capabilities and (v) regional differentiation (Eastern, Central, Western). Xu (2011) finds that the allocation of targets does not correlate with either of those factors (barring a correlation of non-power SO_2 emissions divided by the area of a province). Therefore, there is no evidence that any of these factors drove the target allocation, which suggests that the allocation of targets followed guidelines orthogonal to changes in SO_2 emissions at the turn of the 11th FYP. I also rule out the possibility that the allocation of SO_2 reduction targets followed the cost of abatement, both by itself and net of benefits. If that were the case, my empirical strategy would pick up the compound effect of an SO_2 reduction target and a cost advantage. In Section 6, I derive detailed marginal abatement cost (MAC) curves and combine those with a measure of marginal abatement benefits. I find that neither the MAC nor the marginal welfare impacts correlate with the target allocation, lending further credibility to my empirical strategy (shown in Figures 19 and A.5). Furthermore, I empirically test for the influence of both factors in a robustness check below.

Inference I follow Bertrand, Duflo, and Mullainathan (2004) and compute standard errors that are heteroskedasticity-robust and clustered at the province level. Statistical inference based on these standard errors, however, could still be incorrect if the number of clusters is too small, as the required asymptotic results might not apply. China has 31 provinces, yielding significantly less than 50 clusters, the usual rule of thumb. While reporting the clustered standard errors, I therefore base my statistical inference on p-values derived from the wild bootstrap method described in Cameron, Gelbach, and Miller (2008). This is common practice in applied research on China (see, for instance, Martinez-Bravo et al., 2013).

5.2 The effect of the SO_2 control policy on SO_2 pollution

Baseline results Table 1 provides the summary statistics for my sample. The first two columns of Table 2 show the results from estimation Equation (1) for the effect of the policy for the full 11th FYP. A first glance reveals that the policy had a different effect depending on the indicator used for evaluation. According to the official SO_2 emissions data, the policy was a success: SO_2 emissions decrease in response to the target, with the estimated magnitude being a 5.8% decrease for a one-standard deviation increase in target stringency. The satellite data, by contrast, do not show a significant relationship between the targets and the SO_2 pollution outcomes. The sign of the estimated coefficient is in the same direction as with the official indicator, and the estimated magnitude is even higher, but it is not statistically different from zero.

To improve the precision of the estimates, I increase statistical power by focussing on polluted cities only. Given the same amount of noise in the data, a higher absolute effect can be expected to be more easily detected in this sample. The sample of polluted cities is built by taking the location of each *in situ* measurement stations run by the MEP[11].

[11] I select the 25 nearest pixels up to a distance of at most 25 kilometers around the city centroids, and compute the province-level observations only based on those pixels. Results from this sample are robust to changes in the 25 kilometers cutoff, as most pixels are closer than 15 kilometers from the city centroid.

The point estimate for the effects of a higher reduction target is nearly identical to the overall sample in percentage terms, but the coefficient turns significant because the higher absolute effect in the polluted sample increases statistical power (Table 2). All baseline estimates taken together, I find that that the SO_2 control policy was effective in reducing air pollution over the whole sample period.

Dynamic treatment effects To take a closer look at what might explain these differing findings, I estimate yearly versions of the above DID specification. Instead of collapsing the time periods into before and after the policy as in Equation (1), I interact a dummy for each time period with the reduction target. This yields the following equation:

$$y_{pt} = \beta_0 + \sum_{t=1}^{T} \beta_{1t} Reductiontarget_p \times \gamma_t + \beta_2 Reductiontarget_p + \sum_{t=1}^{T} \beta_{3t}\gamma_t + \alpha_p + u_{pt} \quad (2)$$

As before, y_{pt} are the SO_2 pollution outcomes for province p in time period t. In this specification, the SO_2 reduction targets are interacted with each time period γ_t to estimate the differential trajectory of SO_2 pollution for provinces with different SO_2 reduction targets. These effects are captured in the point estimates for β_{1t} for periods 2 through T, where T is the last observation for 2010.

Figures 3 and 4 plot the estimate of the interaction coefficients $\sum_{t=1}^{T} \beta_{1t}$ in equation (2) and show that the identifying assumption of common pre-trends is satisfied. Both indicators are on a common trend before the start of the policy in 2006. However, there is only one year of satellite data before the start of the policy, because NASA only launched the satellite in late 2004. Figure 5 zooms in on the satellite data and plots the estimates for $\sum_{t=1}^{T} \beta_{1t}$ estimated on monthly satellite data. While there is more noise in the monthly data, there are common pre-trends before the start of the policy, as illustrated through the vertical grey lines.

[INSERT FIGURES 3, 4 and 5 ABOUT HERE]

The official SO_2 emissions data also show a clear deviation from the pre-trend at the start of the policy, lending further support to the identification. SO_2 emissions decrease immediately and linearly until the end of the 11th FYP. On the other hand, a look at the SO_2 satellite data paints a different picture: SO_2 pollution does not go down for provinces with a higher reduction target when whe central government cannot monitor SO_2. After 2008, the satellite data show a distinct drop in SO_2 pollution. Taken together, I find the following: (i) there is possible misreporting of SO_2 emissions during the first two years of the 11th FYP when the Chinese government lacks the ability to properly monitor SO_2 emissions. (ii) once the government gains the ability to monitor, rhetorical compliance turns into real compliance and SO_2 pollution is reduced. Ultimately, the policy worked.

Heterogeneous treatment effects based on initial pollution levels Where do the improvements in air pollution take place? Figure 2 shows a large heterogeneity in the location of SO_2 pollution across provinces. Sichuan, for instance, suffers from a pollution hotspot towards the East, and enjoys comparatively lesser pollution in the Western part of the province. The NASA satellite data allow me to exploit this heterogeneity to ask whether the effect of the SO_2 pollution reduction target in the province differs depending on the initial level of pollution. To answer this question, I map the SO_2 satellite data to the county-level, yiedling 2,638 cross-sectional units. For each county, I measure its mean SO_2 pollution for 2005 relative to all other counties within the same province. This information is captured in a variable that takes the value of 1 for counties in the lowest quartile of initial pollution within their province, up to a value of 4 for counties that have the highest initial pollution. I then re-estimate the DID specification from Equation (1) at the county level on each subsample along the distribution of initial pollution. As above, statistical inference relies on heteroskedasticity standard errors at the province-level and t-statistics from a wild bootstrap procedure with 1000 repetitions. I expect to find no effects for the lowest quartile, because initial pollution in these counties is below 0.2 Dobson Units on average, making further air quality improvements unlikely. Furthermore, for the same precision of data, a nominally smaller effect is harder to detect statistically, making it more likely to find a significant effect the higher the initial level of pollution.

The results in Table 4 confirm that this is the case. As expected, SO_2 reductions in response to the targets mainly take place in the higher two quartiles of the initial distribution, though only the most polluted quartile is statistically significant. This reproduces the baseline results, where the effect only turned statistically significant for the sample of polluted cities due to noise. The effect size for the two highest quartiles is a 9.4% and 9.9% decrease in SO_2 pollution per one standard-deviation increase in the stringency of the reduction target, respectively.

Robustness checks for identification Firstly, I show that the neither the cost nor the welfare impact at the margin correlates with the SO_2 reduction targets. To test this, I use the following empirical specification:

$$\begin{aligned} y_{pt} = {} & \beta_0 + \beta_1 Reductiontarget_p \times D(Post)_t + \beta_2 X_p \times D(Post)_t \\ & + \beta_3 Reductiontarget_p + \beta_4 X_p + D(Post)_t + \alpha_p + u_{pt} \end{aligned} \tag{3}$$

where X_p is either a province-level measure of (i) the marginal abatement cost or (ii) the ratio of marginal abatement benefits to marginal abatement cost. Specific details for the construction of those measures are provided in Section 6. In a nutshell, my approach is this: I compute detailed marginal abatement cost curves for SO_2 for each province in China, based on a reliable set of micro data on the cost and abatement potential of fine-grained polluting activities. The marginal benefits from reducing air pollution are based on a back-of-the-envelope calculation that follows Oliva (2015) and evaluates health improvements based on the value of a statistical life.

Table 5 shows that the estimates for β_2 are 0 and that the estimates for β_1 are nearly identical to those in Table 2. Neither the marginal abatement cost by themselves nor the marginal abatement benefits change the estimated effect of an SO_2 reduction target.

Secondly, I find that there is no relationship between the economic downturn in the Great Recession and the stringency of a province's SO_2 reduction targets. If this were the case, my results could be confounded because a slowdown in economic activity could go hand in hand with a reduction in SO_2 pollution. It is important to note that even though China experienced a slowdown in growth in 2009, this was only a decrease in growth rates. Even in 2009, there are only 3 provinces with a growth rate below 5%, and the growth rate for those provinces is still positive. To check whether the 2009 downturn is correlated with the SO_2 reduction targets, I compute the magnitude of the downturn as the deviation in the average of the growth rates for 2002-2007 for each province. I find that there is no relationship between the magnitude of the recession and the reduction targets in a province. This is because different provinces with similar reduction targets experienced rather different deviations from their long term growth rates in 2009, and overall there is no statistically significant correlation between SO_2 reduction targets and the downturn. A linear regression with standard errors clustered at the province level finds a best fit with a p-value of 0.31 (these results are shown in Figure A.1 in the Appendix). The Great Recession can therefore not explain the decrease in SO_2 pollution from 2008 onwards.

5.3 The effect of monitoring

This subsection provides suggestive evidence that the difference between the official SO_2 emissions data and the SO_2 satellite data before 2008 is driven by the government's newly-gained ability to monitor SO_2 pollution. Firstly, it is important to note that the difference between the NASA satellite data and the official SO2 data starts growing in 2006 and spikes in 2007. From 2008 on, the difference shrinks again to less than the initial levels and stays there by 2010. Secondly, this spike is related to the SO2 reduction targets: the interaction between the year dummy for 2007 and the reduction target of a province from estimating Equation (2) is positive, albeit only significant at p=0.08 for the cluster-robust standard errors or p=0.13 for the wild bootstrap procedure. Both analyses are shown in Figures A.2 and A.3 in the Appendix.

I estimate heterogeneous treatment effects by splitting the sample into provinces for which misreporting before 2008 is more likely than for others. Factors that influence the propensity of regulated provincial governments to take advantage of lacking monitoring and to misreport untrue pollution improvements initially can be grouped into: (i) technological factors that determine the cost of abatement; (ii) socioeconomic factors, such as population density that determine the benefits of abatement; and (iii) the career concerns of provincial leaders themselves. I provide evidence for all three categories.

I construct detailed marginal abatement cost curves (Section 6) and split provinces into those where the marginal abatement cost is high given the actual reduction targets. Marginal abatement costs range from 170€/tSO_2 (Gansu) to 660€/tSO_2 (Heilongjiang). My hypothesis is that those provinces with higher abatement costs will have a higher incentive to misreport in the first period. Next, I compute the marginal benefits for each province and split the sample into two groups according to their welfare impacts. Finally, I use data on party secretaries from Persson and Zhuravskaya (2016) to identify provinces led by politicians with career concerns. I hypothesize that provinces with leaders in the last term in office or with more local ties are less career concerned will engange in less misreporting when the government cannot monitor real compliance[12].

The following triple-differences specification interactions the reduction target with an indicator variable that takes on the value of one for the subsample with more likely misreporting according to the hypotheses discussed above. In this way, I can test whether the likely reason that explains why compliance becomes real in 2008 is in fact the government's ability to monitor SO_2 pollution outcomes. I restrict the sample to the years until 2008 since this test concerns only the subperiod during which the government cannot monitor SO_2 pollution.

$$\begin{aligned} y_{pt} =& \beta_0 + \beta_1 Reductiontarget_p \times D(Post)_t + \beta_2 Reductiontarget_p + \beta_3 D(Post)_t \\ &+ \beta_4 H_p + \beta_5 Reductiontarget_p \times H_p + \beta_6 H_p \times D(Post)_t \\ &+ \beta_7 Reductiontarget_p \times D(Post)_t \times H_p + \alpha_p + u_{pt} \end{aligned} \tag{4}$$

In this specification, H_p is an indicator that takes the value 1 for those provinces more likely to engage in misreporting when the government cannot monitor compliance. The interpretation for the causal effect of a higher reduction target is the estimate for β_1 for the subgroup unlikely to misreport, and the estimate for $\beta_1 + \beta_7$ for the subgroup likely to misreport.

I estimate this specification for both outcome variables and for the normalized differ-

[12] According to Persson and Zhuravskaya (2016), politicians with local ties are more prone to contributing to local public goods.

ence between both outcome variables. This yields 6 different point estimates for the effect of a higher reduction target (3 for the different outcomes, 2 for each subgroup). Table 6 illustrates this. My hypothesis is that for the subsample of provinces more likely to engage in misreporting, the drop in *official SO_2 emissions* is greater, which implies that $\hat{\beta}_1 > \hat{\beta}_1 + \hat{\beta}_7$, which is equivalent to $\hat{\beta}_7 < 0$. For real compliance measured by *SO_2 satellite data*, my hypothesis is that the provinces more likely to misreport will reduce pollution less, implying $\hat{\beta}_1 < \hat{\beta}_1 + \hat{\beta}_7$, which is equivalent to $\hat{\beta}_7 > 0$.

I repeat this procedure for all different ways of splitting the sample. The results are shown in Table 7 and provide only mixed evidence that provinces with a higher incentive to misreport do so. Judging from the signs of the estimated coefficients, it is possible that provinces with leaders who are subject to career concerns report higher SO_2 emission reductions, but lower real SO_2 improvements. The evidence for misreporting based on the cost and benefits from abatement is more mixed: SO_2 emission changes are higher for provinces with higher abatement cost, but lower for provinces with a higher marginal welfare impact. The normalized difference between both data sources is positive for all three hypotheses of misreporting. It is important to note, however, that the interaction coefficient β_5 is never significant due to the low number of cross-sectional units, making these results suggestive rather than definite. Subsection 5.5 below therefore compares my empirical findings to recent research on water pollution control in China to better understand the effect of monitoring on pollution outcomes.

5.4 Reactions by the regulated agents in response to the targets

This subsection provides evidence into the responses by the regulated agents. Two facts stand out from the empirical analysis above. On the one hand, regulated provincial governments likely misreport the desired emissions data when the central government cannot monitor them. On the other hand, my analysis shows that the SO_2 reduction control policy was ultimately successful in reducing air pollution. Provincial governments thus reacted along two margins. They appeared to bring down air pollution through *rhetorical compliance* initially, and switched to *real compliance* towards the end of the policy.

Rhetorical compliance As shown above in Table 3, the effect of the reduction targets has a different sign when comparing the official SO_2 emissions data and the NASA SO_2 satellite data before 2008. Given that the SO_2 satellite data move closely with real SO_2 emissions (see Section 4), this suggests initial misreporting of data. This subsection investigates whether this misreporting was accompanied by further political rhetoric by the provincial governments. In particular, I test whether the provincial governments changed their political rhetoric in response to the SO_2 reduction control targets. To do so, I estimate versions of Equation (2), using different counts of keywords as outcome variables.

Firstly, I use the overall number of keywords related to air pollution in each government work report. Figure 8 plots the estimates for the yearly interaction coefficients β_{1t}. It shows that provinces that received a higher SO_2 reduction target show a distinct increase in their political rhetoric on air pollution.

[INSERT FIGURE 8 ABOUT HERE]

Next, I zoom in and split political statements into those on work done in the preceding year and those related to projects for the year to come. Figure 9 shows that the increase is most pronounced in the period before the central government had the capacity to monitor. Provincial governments that received a higher SO_2 reduction target claim past work on projects related to air pollution in 2007, thus exploiting the central government's inability to monitor SO_2 initially. That is, a higher SO_2 reduction target induces provinces to claim work on air pollution in 2006. Political statements relative to future work on air pollution also increase with the SO_2 reduction targets, although the evidence is less stark (Figure 10). Finally, Table 8 summarizes these results based on estimating Equation (1) on the political attention variables to find that: (i) political attention to air pollution increases with target stringency, and (ii) statements about past work on air pollution peak in the period in which the government could not monitor SO_2[13]. As shown through the comparison of results from the official SO_2 emissions indicator and NASA satellite data, however, SO_2 pollution during that period did not improve.

[INSERT FIGURES 9 and 10 ABOUT HERE]

Taken together, my empirical findings suggest that the provincial governments exploit the central government's inability to monitor along two margins: they likely misreport the desired SO_2 emissions data, and they back up the misreporting through rhetorical compliance in their public political statements.

Real compliance As shown above, the SO_2 reduction targets worked in reducing SO_2

[13]Table 8 includes a robustness check to further show that my estimates on rhetorical compliance are meaningful. This test involves estimating the effect of the SO_2 reduction targets on closely related, yet different placebo outcomes that should not be affected by the SO_2 reduction targets. The 11th Five-Year Plan (2006-2010) included goals to increase China's forest cover from 18.2% to 20% and to extend the coverage of rural medical care from 23.5% to 80%. To measure the political attention towards these policies, I use the count of 'forest' and of 'medical care' in the government work reports as dependent variables for the falsification test. Results based on these outcomes using the specification in Equation (1) are reported in the columns 'Placebo Outcomes' in Table 8. As in the case of the keywords related to air pollution, both keywords are mentioned more often during the 11th Five-Year Plan (2006-2010) than before. Provinces that received a higher SO_2 reduction target, however, do not talk more about either topic. This strongly suggests that governments of provinces with higher reduction targets do not mimic the Central government's political agenda in their own statements in general. Instead, they specifically change their political communications in response to the SO_2 reduction targets.

pollution significantly from 2008 onwards. Based on my own calculations of the marginal cost of SO_2 emissions abatement (see Section 6), the installation of desulfurization devices in existing industrial and power plants, fuel-switching to better quality coal, and the shutdown of small, inefficient thermal units are the main margins by which the provincial governments could reduce SO_2 emissions over relatively short time horizon of 5 years. I provide evidence on the last two channels.

The data on the *installation of scrubbers* include both the capacity of the desulfurization equipment as well as the timing of its installation in each province. Figure 11 shows that while provinces with a higher SO_2 reduction targets installed more desulfurization devices on average, that effect is not statistically different from zero. Figure 12 computes the skewness of the timing of the installation of the scrubbers between the years 2006 and 2010 and correlates it to the SO_2 reduction targets at the province level. As can be seen, provinces with a higher SO_2 reduction target did not install scrubbers earlier than other provinces. The data on the *shutdown of small, inefficient thermal units* shows a much clearer picture. Figure 13 shows that the higher the reduction target, the higher the capacity of small thermal units shut down by 2010. Shanghai and Beijing are the exception to this rule, most likely because they already shut down inefficient plants in the past.

[INSERT FIGURES 11, 12 and 13 ABOUT HERE]

5.5 The empirical findings in context

My findings suggest the importance of the government's monitoring to achieve compliance with the pollution control regulation. A comparison to findings from contamporaneous research strengthens this conclusion: Kahn, Li, and Zhao (2015) study water pollution control in China during the 11th Five-Year Plan (2006-2010). That is, they estimate the effect of a pollution control policy for the same country, time period and institutional setting. Their setting differs in one crucial aspect, however: the central government had access to high-quality monitoring data for water pollution from the start. A comparison to their findings can thus shed light onto the effect of monitoring on compliance.

In Table 4 (Panel A), Kahn, Li, and Zhao (2015) estimate yearly treatment effects for the effect of the water pollution control policy. Their findings show a pattern that is very different from my findings along two characteristics: (i) real compliance is immediate and statistically different from zero for all years (excepting 2009) and (ii) the yearly effects are of a similar magnitude throughout. Given that the salient difference to my setting is the government's ability to monitor pollution already in 2006 and 2007, this comparison provides further evidence that information is the likely channel that explains the behaviour by the provincial governments until 2008 and the lack of real compliance[14].

6 Instrument choice and gains from trade

In this section I address two questions on the evaluation of policy design: what is the welfare-optimizing allocation of targets across provinces? And: what are the welfare gains from trade associated with moving from the actual allocation to the welfare-optimizing allocation?

As discussed in the Introduction, both types of question have been studied for the US (Carlson et al., 2000; Oates, Portney, and McGartland, 1989). This research has shown that the efficiency of different policy instruments is an empirical question. It will depend, among other things, on whether the command-and-control regulation is designed in an enlightened way that takes the cost of pollution abatement into account. Whether China's flagship air pollution control regulation has been cost-efficient is therefore a question of great interest. It is also of immediate policy-relevance: The Economist hypothesized that China relies too much on command-and-control regulation (The Economist, 2013).

In this section, I construct detailed marginal abatement cost (MAC) curves for SO_2 at the province level for China. I use these SO_2 MAC curves to predict the counterfactual cost-efficient allocation. In this way, I can assess whether the actual command-and-control regulation was enlightened and took into account the cost of abatement. Furthermore, I can quantify the efficiency gains from trade from moving from the actual allocation of reduction targets compared to the cost-efficient allocation. As noted by Stavins (2003), cost might be the best measure to assess policy efficiency. Finally, I compute a back-of-the-envelope measure for the marginal benefits of SO_2 abatement at the province level to study whether the gains from trade based on cost alone are driven by omitting the benefits of reducing air pollution.

6.1 Construction of the marginal abatement cost curves

Data I use a rich set of micro data on SO_2 emissions and abatement costs in each province in China at a very detailed level. These data are compiled by the IIASA research institute for use as input into their GAINS model for China (on the model, see IIASA, 2010a, IIASA, 2010b, Klimont et al., 2009). These data rely on a variety of sources of two kinds: *common data*, that are used across different countries and rely on the assumption of free international markets in abatement equipment, and *country-specific data*. Common data include the unit investment cost for technologies, fixed costs of operation, and the amount of input factors needed for some of the variable cost components. These data have

[14]Relatedly, Stavins (1998) mentions that 'the SO_2 program has also brought home the importance of monitoring and enforcement provisions' in the US experience of SO_2 pollution control, lending further support to my empirical findings.

been compiled and updated by IIASA for several decades on the basis of expert meetings at the UN (AP EnvEcon, 2010).

Country-specific data, on the other hand, include a detailed breakdown of China's industrial structure: the type and size of polluting installations, the facility operating conditions, national fuel consumption data, prices of inputs (labour, electricity, fuel cost) as well as unabated emission factors and removal efficiencies (AP EnvEcon, 2010). The local data are compiled by IIASA experts in collaboration with local experts from the Chinese Energy Research Institute in Beijing and Tsinghua (Purohit et al., 2010). These data are combined into unit cost estimates per technology, as well as abatement potential[15].

The data give a detailed breakdown of all SO_2 emission sources for each province, split into different sectors (such as the combustion of coal) that use different fuels (such as gas or low-sulfur coal) and each of which has different abatement technologies at its disposal (such as limestone injection). Each of the abatement technologies is characterised by a unit cost of abatement (in €/tSO_2) and its abatement potential (i.e. how much would the emissions factor of the current sector-fuel combination be lowered when switching to the abatement technology). To illustrate, one abatement technology would be the use of limestone-injection (*abatement technology*) in a modern coal-fired power plant (*fuel* and *sector*) at the cost of 515.38€/tSO_2 (*unitcost of abatement*).

Construction of the MAC curves To construct the MAC curve for one province, I follow a two-step procedure. First, I rank each abatement option by the unit cost of abatement within each sector and fuel. Second, I abate SO_2 emissions within each sector and fuel in increasing cost order across the Province. Figure 14 shows two examples, Beijing and Sichuan, to illustrate the large heterogeneity in marginal abatement cost across provinces (Appendix A.2 contains the MAC curves for each province). In contrast to many MAC studies, I do not rely on a top-down model but base the cost estimates entirely on data, and I am the first to provide complete MAC curves on Chinese provinces in this way.

[INSERT FIGURE 14 ABOUT HERE]

Consistency checks To check that my data represent the official Chinese data well, Figure 15 shows the average ratio between the GAINS SO_2 emissions and the official MEP SO_2 emissions as a function of the province's level of emissions: the overall fit be-

[15]The data on the cost of abatement are thus not purely engineering estimates, as they rely on the real cost of abatement in different countries. Rather than being a lower bound on the cost of abatement, they are more likely to be an upper bound: cheaper local inputs and China's capacity for scale will likely allow for cheaper abatement. The use of MAC estimates from revealed-preference settings such as Meng (Forthcoming) and Gosnell, List, and Metcalfe (2016) would be desirable, but no such estimates are available for China or at the required level of comprehensiveness.

tween my data and the MEP data is good (correlation: 85.9%) and fairly stable across provinces, except for two outliers with very low emissions: Tibet and Hainan. Overall, the GAINS data report higher emissions for all provinces than do the official data. This is in line with the literature in atmospheric science that has found that GAINS data tend to be more comprehensive and slightly overpredict official SO_2 data sources, which is possibly due to differing assumptions on the distribution of fuel consumption across sectors (Klimont et al., 2009) and the fact that official MEP statistics lack rural pollution sources and biofuels (Lu et al., 2010).

[INSERT FIGURE 15 ABOUT HERE]

Additionally, there is a drawback in using detailed SO_2 emissions data at the microlevel: the detailed breakdown across sectors does not allow me to compute the cost of moving activity across sectors (such as moving electricity generation from coal-fired power plants to wind turbines). In other words, these data limitations make it necessary to assume scrappage cost and imperfect substitution across sectors in the short run. Given the 5-year horizon of the policy, however, I feel that this is a reasonable assumption to make.

6.2 The cost-efficient counterfactual allocation of SO_2 reduction targets

In this subsection, I compute the counterfactual allocation of SO_2 reduction targets across provinces for the cost-efficient policy. I pool the province-level microdata and construct the SO_2 marginal abatement cost curve at the national level for China. As shown in Figure 16, marginal abatement cost is relatively flat until an abatement level of 35% of 2005 SO_2 emissions. Up until this abatement level, the cost of SO_2 abatement at the margin is less than 500€/tSO_2. Beyond ca. 45% of 2005 SO_2 emissions, however, the picture changes and the cost of abatement rises rapidly. Beyond 50%, abatement becomes prohibitively costly.

[INSERT FIGURE 16 ABOUT HERE]

Figure 16 also shows the counterfactual marginal abatement cost for the national SO_2 reduction target of 10%, given by the intersection between the vertical line at 10% and the MAC curve. At an abatement level of 10%, the cost-efficient allocation of SO_2 reduction targets would have led to a marginal abatement cost of 338€/tSO_2. This figure is the counterfactual marginal cost of the Chinese government's SO_2 emissions control strategy had the central government distributed the provincial reduction targets in a cost-optimal way.

The intersection of the marginal abatement cost curve for China as a whole with the

10% national abatement level also produces the cost-efficient allocation of SO_2 reduction targets across provinces. Figure 17 shows that this allocation is far more skewed than the distribution of targets that was actually used in the 11th FYP. The actual allocation already ranges from targets of 0% to 25.9%, but is approximately uniformly distributed within this range. The cost-efficient allocation, by contrast, is more unequal. A small number of provinces would bear most of the reductions. These provinces are Sichuan, Shandong, and Zhejiang. The industrial structure of these provinces allows for comparatively cheap installation of wet flue-gas desulfurisation in industry and power plants and the use of more efficient combustion processes in refineries and steel sintering.

[INSERT FIGURE 17 ABOUT HERE]

Based on the MAC curves, I find that the Chinese government did not equate marginal abatement cost across space. Instead, the reduction targets favoured coastal provinces in the East even though abatement coasts are higher at the margin. These results are shown in Table 9. Figure 18 depicts the same data on a map that colours the difference between the actual and the cost-efficient allocation. The actual allocation is consistent with a tale-of-two-cities story, in which China would develop amenity-based consumer cities along the coast, while maintaining a base of polluting manufacturing in its interior (Kahn, 2006; Zheng and Kahn, 2013). Shanghai is a prime example for this: under the cost-efficient allocation, Shanghai would have received an SO_2 reduction target of 6.2% on 2005 levels. The actual allocation, however, gave Shanghai a reduction target of 25.9%, or more than four times the cost-efficient target.

[INSERT FIGURE 18 ABOUT HERE]

My findings show that actual command-and-control regulation that China used to control SO_2 pollution in the 11th FYP was not cost-efficient. As suggested by Oates, Portney, and McGartland (1989), command-and-control regulation will only be efficient if it is designed in an enlightened fashion by keeping an eye on the cost of abatement. Figure 19 shows that this was not done by the Chinese government in 2005. There is no statistically significant relationship between the SO_2 reduction target a province received under the 11th FYP and its abatement cost at the margin.

[INSERT FIGURE 19 ABOUT HERE]

6.3 A measure for marginal abatement benefits

My final aim is to quantify the gains from trade from moving from a command-and-control regulation to a market-based allocation based on abatement cost. Because China's provinces differ markedly with respect to income levels, population densities and initial pollution, it is important to not rely exclusively on a measure of cost. I use the method employed by Oliva (2015) to construct a back-of-the-envelope measure for the marginal abatement benefits of reducing SO_2 pollution at the province level. This method proceeds in 3 steps: (i) how does the SO_2 control policy change pollutant concentrations?, (ii) what health effects do the changes in pollutant concentrations cause? and (iii) what is the monetary value of those health effects?

(i) Changes in pollutant concentrations I use the results from Wang, Jang, et al. (2010b), who use the CMAP modelling system (maintained by the US Environmental Protection Agency) to simulate the ex ante effects of the SO_2 reduction policy on concentrations of SO_2 and $PM_{2.5}$. Their measurements allow me to attribute changes in pollutant concentrations to the SO_2 emissions reduction target of each province.

(ii) Health effects To convert the changes pollutant concentrations to changes in health outcomes, I use dose-response estimates for $PM_{2.5}$ and SO_2 from Bombardini and Li (2016). They estimate an elasticity of infant mortality rates of 0.9 to SO_2 and of 2.2 for $PM_{2.5}$. I combine those estimates with data on SO_2 and $PM_{2.5}$ levels in 2005 from MEP and the *China Energy Databook* (Fridley, Romankiewicz, and Fino-Chen, 2013) to approximate a linear dose-response function for each pollutant.

I use the estimates from Bombardini and Li (2016) for two reasons: firstly, they use an instrumental strategy approach to estimate a dose-response function for the health effects of air pollution, thus correcting downward bias from OLS estimates (due to migration, income effects and avoidance behaviour). Additionally, their study is from China, from a recent period, and includes consistent estimates for both SO_2 and $PM_{2.5}$, which are the main pollutants that are affected by the SO_2 reduction policy. The downside is that this restricts my focus on infant mortality when calculating the benefits from reducing air pollution[16].

(iii) Valuation To convert the health damages into monetary values, I use a baseline value of a statistical life (VSL) of 1 million yuan from World Bank and State Environmental Protection Administration (2007). This value is a midpoint between the VSL estimates from the reviewed studies ranging from 0.24 to 1.7 million yuan[17]. Following Hammitt and

[16] Infant mortality, however, is likely to capture a first-order welfare effect. Matus et al. (2012) calculate that 71.4% of all air pollution costs in China are health costs, and that mortality captures over 85% of those health costs. Chen, Ebenstein, et al. (2013), in turn, show that mortality impacts from TSP are strongest in infants. Evidence from Indonesian wildfires in 1997 also points to large infant mortality effects from exposure to particulates, mostly driven by prenatal exposure (Jayachandran, 2009). Greenstone and Hanna (2014) also focus on infant mortality to evaluate the effect of pollution.

[17] VSL estimates for China are available for more recent periods (Ito and Zhang, 2016), but to evaluate a policy from 2005 I prefer to use estimates from that period.

Robinson (2011), I account for the income heterogeneity across Chinese provinces by adjusting the central VSL estimate according to the income level in each province using an income elasticity of VSL of one. This yields VSL estimates from 360,000 yuan (Guizhou) to 3,000,000 yuan (Shanghai).

I multiply these VSL estimates by the mortality numbers to compute the benefit of reducing SO_2 emissions by 1 kt for each province. Since I employ a linear approximation, the marginal benefits of abatement is constant[18].

Caveats I rely on an overly conservative measure of benefits by including only infant mortality which may underestimate the true benefits and lead to an underestimation of the welfare-optimizing level of SO_2 abatement for China. However, this concern is muted in practice since even my lower bound benefit measure suggests SO_2 abatement up until the prohibitive marginal abatement cost ranges. Additional abatement benefits at the margin would thus only have a negligible effect on the welfare-optimizing abatement level. The advantage is that the dose-response functions used are more precise for infants because low migration translates into better knowledge of lifetime exposure to pollution, improving the consistency of estimates *across* provinces.

The welfare-optimal SO_2 target allocation In a similar vein to the MAC data, I use the combined data on both marginal cost and benefit to construct marginal welfare impact curves. I combine the marginal abatement benefit data with the marginal abatement cost data by dividing the marginal abatement cost by the marginal abatement benefit to obtain a measure of the marginal welfare impact of abatement. When this ratio is below one, benefits are larger than cost. Once this ratio exceeds one, costs are higher than benefits. Figure 20 pools the province-level data to obtain the marginal welfare impacts for China as a whole. The vertical dashed line at 0.1 marks the 10% national SO_2 emissions control target of the 11th FYP. The marginal benefits of abatement exceed the marginal cost of abatement by more than 14 times at this abatement level[19]. Again, there is no correlation between a province's SO_2 reduction target and the marginal welfare impact of reducing SO_2 emissions (shown in Figure A.5 in the Appendix).

[18] For those provinces with a reduction target of zero, I find that the benefits are 0 or nearly 0, too. This is due to low initial pollution, low VSL estimates, and low population numbers. The exception is Gansu, for which I compute the marginal benefits as the average of its 5 nearest neighbours with respect to initial pollution, VSL estimates, and population numbers.

[19] It is also possible to predict the welfare-optimal allocation of SO_2 reduction targets across provinces. These results are shown in Table 9 for comparison with the counterfactual cost-efficient allocation. The distribution of targets is skewed similar to the cost-efficient allocation. The provinces that would receive a higher reduction targets are different, however. Tianjin, Guangdong, Zhejiang and Hubei would receive the highest reduction targets. The exception is Shanghai, which would maintain its 6.2% reduction target from the cost-efficient allocation.

6.4 Gains from trade

Finally, I study the gains from trade of the two different policy instruments. Gains from trade will be possible if the actual command-and-control regulation used to allocate the SO_2 reduction targets in the 11th FYP was not done efficiently. As I have shown above, this is the case. This subsection quantifies these efficiency gains in terms of cost-efficiency. To further take into account the heterogeneities across Chinese provinces, I construct a back-of-the-envelope measure for the marginal benefits of air pollution abatement to study the robustness of the findings based on pure cost-efficiency.

Table 9 shows that the cost-efficient and the optimal allocation differ strongly for a number of provinces, most notably for Sichuan, Chongqing, Tianjin and Guangdong (right panel). This shows that moving from the actual to the cost-efficient allocation does not necessarily present an efficiency gain ex ante along both efficiency measure. Instead, the correlation between the cost-efficient and the welfare-optimising allocation is slightly negative, while the correlation between the actual and the cost-efficient and optimal allocations is similar (29% vs 28%)[20].

Table 10 shows the gains from trade that are possible. Firstly, I follow the literature Stavins (2003) and assess the gains from trade using the cost efficiency measure. I find that moving from the actual allocation to the cost-efficient allocation would decrease abatement cost by 49% at the margin (from 658€/tSO_2 to 338€/tSO_2). Secondly, I find that this conclusion is robust to taking into account the benefit side. Adding the benefit measure, I find efficiency improvements of 45% (from a welfare ratio of 2.19 to a welfare ratio of 1.2).

Summarizing the findings, I find that the command-and-control regulation on SO_2 pollution control in the 11th FYP does not take into account cost of abatement. Secondly, I find that a market-based allocation would improve both cost-and welfare-efficiency by 49%. Taken at face value, this finding seems to differ from Carlson et al. (2000) for the US. They find that while an emissions trading scheme led to appreciable, but lower than expected efficiency gains in the US context, these efficiency gains would have been twice as high had the alternative command-and-control policy been the forced adoption of scrubbers. In China, those two allocations are nearly identical: the cost-efficient allocation of SO2 reduction targets relies heavily on the use of wet flue-gas desulfurization in the electricity sector. Low-sulfur coal, by contrast, only becomes cost-efficient after about 18% of 2005 SO_2 emissions in China are abated, making scrubbers the most economical abatement technology to reduce SO_2 emission levels by 10%. Overall, my findings are therefore surprising compared to Carlson et al. (2000)'s findings from the US: while a market-based policy instrument such as an SO_2 emissions trading scheme with a national cap would

[20] It is further possible to calculate other allocations such as achieving the 10% reduction through the mandated use of scrubbers or with a 10% target for each provinces. Detailed results from these calculations are available upon request.

double allocative efficiency in China and, at the margin, induce cost savings and increase welfare by 50%, most of those efficiency gains could be reaped by forced installation of scrubbers.

7 Concluding remarks

This research evaluates China's main air pollution control policy. In 2005, China decided on a 10% SO_2 emissions reduction goal as part of the 11th Five-Year Plan (2006-2010). To study the effect of this policy on pollution outcomes, I use both official, misreporting-prone indicator and independent NASA SO_2 satellite data in a differences-in-differences strategy that exploits variation in target stringency at the province level. Overall, my empirical findings show that SO_2 pollution control in the 11th Five-Year Plan (2006-2010) was a success. SO_2 pollution, as measured by independent satellite data, decreased by more than 10% as a result of the policy, most notably through the shutdown of inefficient and small thermal units. Compliance, however, only started when the central government upgraded its pollution monitoring capacity in 2008. I provide suggestive evidence that the increase in monitoring capacity is the likely explanation for this compliance. Before the changes in monitoring, the behaviour by the provincial governments fits the old Chinese adage 'Heaven is High and the Emperor far away': misreporting of pollution data seems pervasive and backed up by misleading rhetorical compliance. This finding is in line with the experiences of environmental regulation summarized in Stavins (2003), who cautions about the importance of monitoring and enforcement. My findings also provide support for a theoretical literature pioneered by Jean-Jacques Laffont suggesting the importance of regulatory capacity for the working of even simple policies (for an overview, see Laffont, 2005).

To analyze the efficiency of the chosen command-and-control regulation, I construct detailed marginal abatement cost curves for each province in China and show the large heterogeneity across provinces in marginal abatement cost. Based on these cost curves, I can quantify the efficiency gain of moving from the actual command-and-control regulation to a counterfactual market-based policy instruments such as a cost-based SO_2 emissions trading scheme. I find that efficiency would have increased by 95% at the margin, lowering marginal abatement cost from 658€/tSO_2 to 338€/tSO_2. This finding is robust to the inclusion of a back-of-the-envelope measure of marginal abatement benefits.

My analysis, then, suggests that China faces similar problems in its air pollution control as more developed countries in the past. Air pollution control is initially plagued by monitoring difficulties, and works in general, though it is not initially cost-efficient. In analyzing environmental policy in China during the 11th Five-Year Plan (2006-2010), I zoom in on a period in recent Chinese history that saw the turning point in environmental governance with the creation of the Ministry of Environmental Protection in 2008. Figure 21 shows that the change in policy coincides with a change in China's dynamics of pollution and economic growth. My research can thus be seen as a detailed account of how China achieved the turning point on its environmental Kuznets curve (Grossman and Krueger, 1995).

Understanding China's ongoing transition from the planet's biggest dirty industrial powerhouse to a cleaner consumer economy is crucial for China, but also for the planet as a whole. The study of China adds to the empirical knowledge about the working of regulation in developing countries (Duflo et al., 2013; Oliva, 2015) and shows how appropriate regulatory capacity of the government is a pre-requisite for effective regulation.

References

Abito, J. M. (2012): "Welfare Gains from Optimal Pollution Regulation", Job Market Paper.

Almond, D., Y. Chen, M. Greenstone, and H. Li (2009): "Winter Heating or Clean Air? Unintended Impacts of China's Huai River Policy", *American Economic Review (Papers and Proceedings)* 99(2), pp. 184–190.

Andrews, S. Q. (2008): "Inconsistencies in air quality metrics: 'Blue Sky' days and PM10 concentrations in Beijing", *Environmental Research Letters* 3, pp. 1–14.

Antle, J., S. Capalbo, S. Mooney, E. Elliott, and K. Paustian (2003): "Spatial heterogeneity, contract design, and the efficiency of carbon sequestration policies for agriculture", *Journal of Environmental Economics and Management* 46(2), pp. 231–250.

AP EnvEcon (2010): *GAINS Abatement Cost Guide. Version 1: A synthesis of information relating to the estimation, handling and use of abatement cost variables in the GAINS model framework.*

Auffhammer, M. and R. Kellogg (2011): "Clearing the Air? The Effects of Gasoline Content Regulation on Air Quality", *American Economic Review* 101, pp. 2687–2722.

Becker, R. A. (2005): "Air pollution abatement costs under the Clean Air Act: evidence from the PACE survey", *Journal of Environmental Economics and Management* 50(1), pp. 144–169.

Becker, R. A. and J. V. Henderson (2001): "Cost of Air Quality Regulation", *Behavioral and Distributional Effects of Environmental Policy*, ed. by C. Carraro and G. E. Metcalf, Chicago: University of Chicago Press, pp. 159–186.

Bertrand, M., E. Duflo, and S. Mullainathan (2004): "How Much Should We Trust Differences-in-Differences Estimates?", *Quarterly Journal of Economics* 199(1), pp. 249–275.

Bombardini, M. and B. Li (2016): "Trade, Pollution and Mortality in China", Mimeograph.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008): "Bootstrap-Based Improvements for Inference with Clustered Errors", *Review of Economics and Statistics* 90(3), pp. 414–427.

Carlson, C., D. Burtraw, M. Cropper, and K. L. Palmer (2000): "Sulfur Dioxide Control by Electric Utilities: What Are the Gains from Trade?", *Journal of Political Economy* 108(6), pp. 1292–1326.

Chang, T., J. Graff Zivin, T. Gross, and M. Neidell (2016): "The Effect of Pollution on Worker Productivity: Evidence from Call-Center Workers in China", NBER Working Paper No. 22328.

Chay, K. Y. and M. Greenstone (2005): "Does Air Quality Matter? Evidence from the Housing Market", *Journal of Political Economy* 113(2), pp. 376–424.

Chay, K., C. Dobkin, and M. Greenstone (2003): "The Clean Air Act of 1970 and adult mortality", *Journal of Risk and Uncertainty* 27(3), pp. 279–300.

Chen, Y., H. Li, and L.-A. Zhou (2005): "Relative performance evaluation and the turnover of provincial leaders in China", *Economics Letters* 88(3), pp. 421–425.

Chen, Y., A. Ebenstein, M. Greenstone, and H. Li (2013): "Evidence on the impact of sustained exposure to air pollution on life expectancy from China's Huai River policy", *Proceedings of the National Academy of Sciences* 110(32), pp. 12936–12941.

Chen, Y., G. Z. Jin, N. Kumar, and G. Shi (2012): "Gaming in Air Pollution Data? Lessons from China", *B.E. Journal of Economic Analysis & Policy* 12(3).

Duflo, E., M. Greenstone, R. Pande, and N. Ryan (2013): "Truth-telling by Third-party Auditors and the Response of Polluting Firms: Experimental Evidence from India", *Quarterly Journal of Economics* 128(4), pp. 1499–1545.

Fioletov, V. E., C. A. McLinden, N. Krotkov, and C. Li (2015): "Lifetimes and emissions of SO2 from point sources estimated from OMI", *Geophysical Research Letters* 42(6), pp. 1969–1976.

Fouquet, R. (2011): "Long run trends in energy-related external costs", *Ecological Economics* 70(12), pp. 2380–2389.

Fridley, D., J. Romankiewicz, and C. Fino-Chen (2013): *China Energy Databook Version 8.0*, Lawrence Berkeley National Laboratory.

Gao, C., H. Yin, N. Ai, and Z. Huang (2009): "Historical Analysis of SO2 Pollution Control Policies in China", *Environmental Management* 43, pp. 447–457.

Ghanem, D. and J. Zhang (2014): "'Effortless Perfection:' Do Chinese cities manipulate air pollution data?", *Journal of Environmental Economics and Management* 68(2), pp. 203–225.

Gollop, F. M. and M. J. Roberts (1985): "Cost-Minimizing Regulation of Sulfur Emissions: Regional Gains in Electric Power", *Review of Economics and Statistics* 67(1), p. 81.

Gosnell, G. K., J. A. List, and R. Metcalfe (2016): "A New Approach to an Age-Old Problem: Solving Externalities by Incenting Workers Directly", NBER Working Paper No. 22316.

Greenstone, M. (2004): "Did the Clean Air Act cause the remarkable decline in sulfur dioxide concentrations?", *Journal of Environmental Economics and Management* 47(3), pp. 585–611.

Greenstone, M. and R. Hanna (2014): "Environmental Regulations, Air and Water Pollution, and Infant Mortality in India", *American Economic Review* 104(10), pp. 3038–3072.

Grossman, G. and A. Krueger (1995): "Economic Growth and the Environment", *Quarterly Journal of Economics* 110(2), pp. 353–377.

GTZ (2009): "Air Quality Management", *Sustainable Transport: A Sourcebook for Policy-Makers in Developing Cities*, Eschborn, Germany: Gesellschaft fuer Technische Zusammenarbeit.

Guan, D., Z. Liu, Y. Geng, S. Lindner, and K. Hubacek (2012): "The gigatonne gap in China's carbon dioxide inventories", *Nature Climate Change* 2(9), pp. 672–675.

Hammitt, J. K. and L. A. Robinson (2011): "The Income Elasticity of the Value per Statistical Life: Transferring Estimates between High and Low Income Populations", *Journal of Benefit-Cost Analysis* 2(1), pp. 1–27.

Hansman, C., J. Hjort, and G. Leon (2015): "Interlinked Firms and the Consequences of Piecemeal Regulation", Mimeograph.

Hao, J., K. He, L. Duan, J. Li, and L. Wang (2007): "Air pollution and its control in China", *Frontiers of Environmental Science & Engineering in China* 1(2), pp. 129–142.

Hartman, R. S., D. Wheeler, and M. Singh (1997): "The cost of air pollution abatement", *Applied Economics* 29(6), pp. 759–774.

He, H. (2012): "Air Pollutant Concentrations and Trends over the Eastern U.S. and China: Aircraft Measurements and Numerical Simulations", Doctoral Dissertation, University of Maryland, College Park.

Henderson, V. (1996): "Effects of Air Quality Regulation", *American Economic Review* 86(4), pp. 789–813.

IIASA (2010a): "GAINS Asia: A Tool to Combat Air Pollution and Climate Change Simultaneously".

— (2010b): "GAINS Asia: Scenarios for Cost-Effective Control of Air Pollution and Greenhouse Gases".

Ito, K. and S. Zhang (2016): “Willingness to Pay for Clean Air: Evidence from Air Purifier Markets”, NBER Working Paper No. 22367.

Jayachandran, S. (2009): “Air quality and early-life mortality evidence from Indonesia’s wildfires”, *Journal of Human Resources* 44(4), pp. 916–954.

Jia, R. (2014): “Pollution for Promotion”, Job Market Paper.

Kahn, M. E. (2006): *Green Cities: Urban Growth and the Environment*, Washington D.C.: Brookings Institution Press.

Kahn, M. E., P. Li, and D. Zhao (2015): “Water Pollution Progress at Borders: The Role of Changes in China’s Political Promotion Incentives”, *American Economic Journal: Economic Policy* 7(4), pp. 223–242.

Keller, W. and A. Levinson (2002): “Pollution abatement costs and foreign direct investment inflows to US states”, *Review of Economics and Statistics* 84(4), pp. 691–703.

Klimont, Z. et al. (2009): “Projections of SO2 , NOx and carbonaceous aerosols emissions in Asia”, *Tellus B* 61(4), pp. 602–617.

Krotkov, N. A. et al. (2016): “Aura OMI observations of regional SO2 and NO2 pollution changes from 2005 to 2015”, *Atmospheric Chemistry and Physics* 16(7), pp. 4605–4629.

Laffont, J.-J. (2005): *Regulation and Development*, Cambridge, UK: Cambridge University Press.

Li, H. and L.-A. Zhou (2005): “Political turnover and economic performance: the incentive role of personnel control in China”, *Journal of Public Economics* 89(9-10), pp. 1743–1762.

Li, X., X. Wu, and F. Zhang (2015): “A method for analyzing pollution control policies: Application to SO2 emissions in China”, *Energy Economics* 49, pp. 451–459.

Lin, J. (2007): “Energy conservation investments: A comparison between China and the US”, *Energy Policy* 35, pp. 916–924.

Lu, Z., D. G. Streets, Q. Zhang, S. Wang, G. R. Carmichael, Y. F. Cheng, C. Wei, M. Chin, T. Diehl, and Q. Tan (2010): “Sulfur dioxide emissions in China and sulfur trends in East Asia since 2000”, *Atmospheric Chemistry and Physics* 10(13), pp. 6311–6331.

Martinez-Bravo, M., G. Padró i Miquel, N. Qian, and Y. Yao (2013): “Political Reform in China: Elections, Public Goods and Income Distribution”, Mimeograph.

Matus, K., K.-M. Nam, N. E. Selin, L. N. Lamsal, J. M. Reilly, and S. Paltsev (2012): “Health damages from air pollution in China”, *Global Environmental Change* 22(1), pp. 55–66.

McElwee, C. R. (2011): *Environmental Law in China*, Oxford; New York: Oxford University Press.

Meng, K. C. (Forthcoming): “Using a Free Permit Rule to Forecast the Marginal Abatement Cost of Proposed Climate Policy”, *American Economic Review.*

MEP (2008a): “Issuance of Automatic Pollution Source Monitoring”, Original in Chinese.

— (2008b): “Management Measures for the Operation of Automatic Monitoring Installations for Pollution Sources”, Original in Chinese.

Montero, J.-P., J. M. Sanchez, and R. Katz (2000): “A Market-Based Environmental Policy Experiment in Chile”, MIT CEEPR Working Paper 00-005.

Moser, A. (2013): “The Yangzonghai case: struggling for environmental justice”, *China and the Environment. The Green Revolution*, ed. by S. Geal, London: Zed Books, pp. 96–135.

NASA (2014): “OMSO2 Readme File v1.2.0 Released Feb 26, 2008 Updated: September 26, 2014”.

Newell, R. G. and R. N. Stavins (2003): “Cost heterogeneity and the potential savings from market-based policies”, *Journal of Regulatory Economics* 23(1), pp. 43–59.

Oates, W. E., P. R. Portney, and A. M. McGartland (1989): "The net benefits of incentive-based regulation: a case study of environmental standard setting", *American Economic Review* 79(5), pp. 1233–1242.

Oliva, P. (2015): "Environmental Regulations and Corruption: Automobile Emissions in Mexico City", *Journal of Political Economy* 123(3), pp. 686–724.

Persson, P. and E. Zhuravskaya (2016): "The Limits of Career Concerns in Federalism: Evidence from China", *Journal of the European Economic Association* 14(2), pp. 338–374.

Purohit, P., M. Amann, R. Mathur, I. Gupta, S. Marwah, V. Verma, I. Bertok, J. Borken-Kleefeld, A. Chambers, J. Cofala, et al. (2010): "Gains Asia: Scenarios for cost-effective control of air pollution and greenhouse gases in India".

Schmalensee, R., P. L. Joskow, A. D. Ellerman, J. P. Montero, and E. M. Bailey (1998): "An Interim Evaluation of Sulfur Dioxide Emissions Trading", *Journal of Economic Perspectives* 12(3), pp. 53–68.

Schreifels, J. J., Y. Fu, and E. J. Wilson (2012): "Sulfur dioxide control in China", *Energy Policy* 48(2012), pp. 779–789.

Schultze, C. R. (1977): *The Public Use of Private Interest*, Washington D.C.: The Brookings Institution.

Song, X.-l., L.-y. Zhang, M.-j. Gong, X. Wang, and Y. Qi (2015): "China's Emission Data Accounting System of the Total Emission Control During the 11th Five-Year Plan Period", *China Population, Resources and Environment* 25(1), pp. 46–54.

State Council (2006): "Major Pollutant Emissions Control in the 11th Five-Year Plan Period", Original in Chinese.

— (2007): "State Council approved the transfer of energy - saving emission reduction statistics. Monitoring and assessment of the implementation of programs and methods of notification", Original in Chinese.

Stavins, R. N. (1998): "What Can We Lean from the Grand Policy Experiment? Lessons from SO2 Allowance Trading", *Journal of Economic Perspectives* 12(3), pp. 69–88.

— (2003): "Experience With Market-Based Environmental Policy Instruments", *Handbook of Environmental Economics*, ed. by K.-G. MÃler and J. Vincent, vol. Volume 1, Elsevier Science B.V.

Stoerk, T. (2016): "Statistical corruption in Beijing's air quality data has likely ended in 2012", *Atmospheric Environment* 127, pp. 365–371.

The Economist (2013): "The East is grey", *The Economist* (Aug 10th 2013).

Tu, Z. and R. Shen (2014): "Can China's Industrial SO2 Emissions Trading Pilot Scheme Reduce Pollution Abatement Costs?", *Sustainability* 6(11), pp. 7621–7645.

Wang, A. L. (2013): "The search for sustainable legitimacy: Environmental law and bureaucracy in China", *Harvard Environmental Law Review* 37, pp. 367–440.

Wang, L., C. Jang, et al. (2010a): "Assessment of air quality benefits from national air pollution control policies in China. Part I: Background, emission scenarios and evaluation of meteorological predictions", *Atmospheric Environment* 44(28), pp. 3442–3448.

— (2010b): "Assessment of air quality benefits from national air pollution control policies in China. Part II: Evaluation of air quality predictions and air quality benefits assessment", *Atmospheric Environment* 44(28), pp. 3449–3457.

Watts, J. (2011): *When a Billion Chinese Jump*, London: Faber and Faber Ltd.

WHO (2014): *7 million premature deaths annually linked to air pollution.*

World Bank and State Environmental Protection Administration (2007): *Cost of Pollution in China: Economic Estimates of Physical Damages*, Conference Edition.

Xu, Y. (2011): "The use of a goal for SO2 mitigation planning and management in China's 11th Five-Year Plan", *Journal of Environmental Planning and Management* 54(6), pp. 769–783.

Yang, G., Y. Wang, Y. Zeng, G. F. Gao, X. Liang, M. Zhou, X. Wan, S. Yu, Y. Jiang, M. Naghavi, et al. (2013): "Rapid health transition in China, 1990-2010: findings from the Global Burden of Disease Study 2010", *The Lancet* 381(9882), pp. 1987–2015.

Zheng, S. and M. E. Kahn (2013): "Understanding China's Urban Pollution Dynamics", *Journal of Economic Literature* 51(3), pp. 731–772.

Zheng, S., C. Sun, Y. Qi, and M. E. Kahn (2014): "The Evolving Geography of China's Industrial Production: Implications for Pollution Dynamics and Urban Quality of Life", *Journal of Economic Surveys* 28(4), pp. 709–724.

Figures

Figure 1: The SO_2 emission control targets

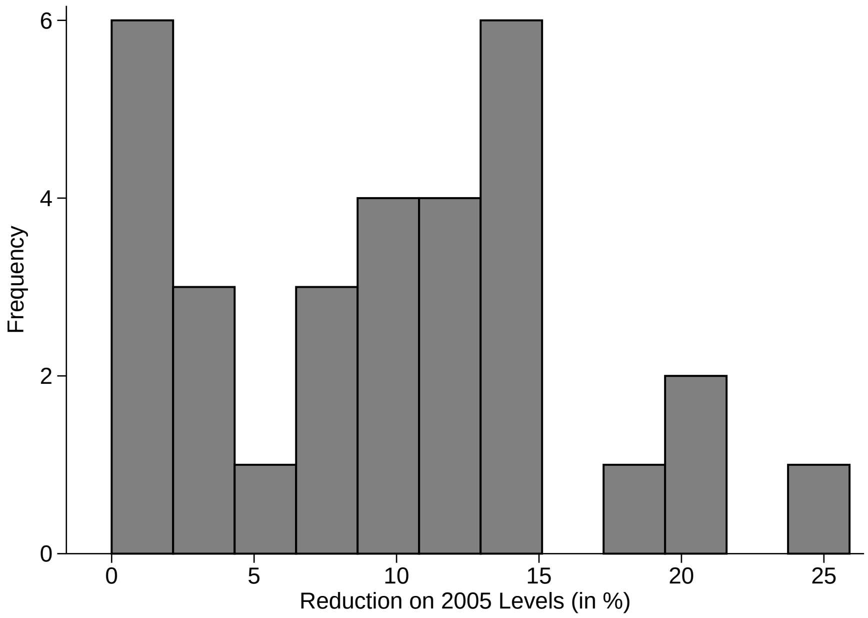

Notes: This figure shows the variation in the SO_2 emissions reduction targets across the 31 provinces. Targets are shown as percentage reduction on 2005 SO_2 emission baselines. The mean of the distribution is 9.4% and the standard deviation is 6.8 percentage points. Data source: State Council (2006).

Figure 2: SO_2 POLLUTION IN CHINA IN JANUARY 2006 BASED ON NASA SATELLITE DATA

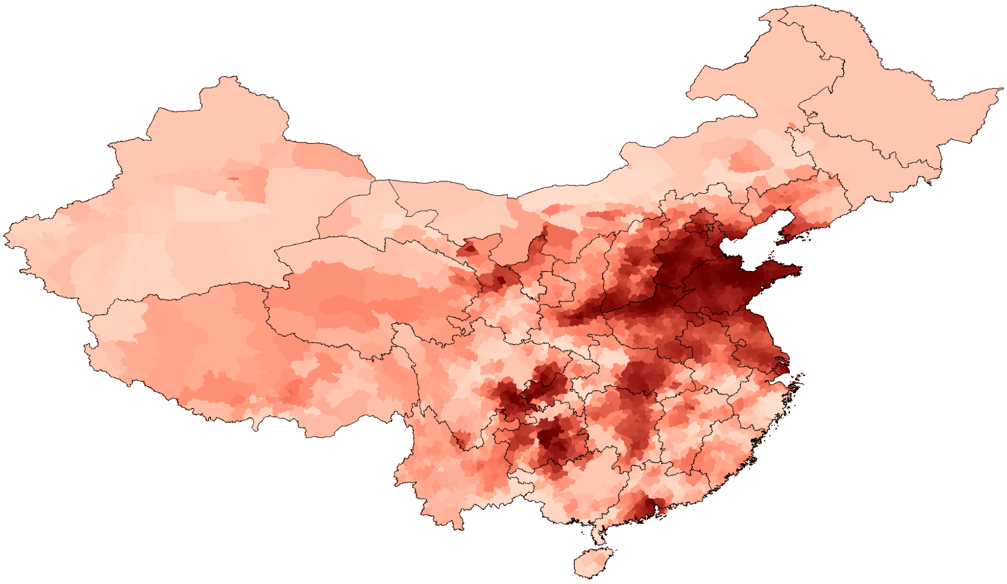

Notes: This figure shows the cross-section of the SO_2 satellite data based on the NASA OMI SO_2 data product for January 2006 mapped to the county-level.

Figure 3: Dynamic treatment effects: Official SO_2 emissions

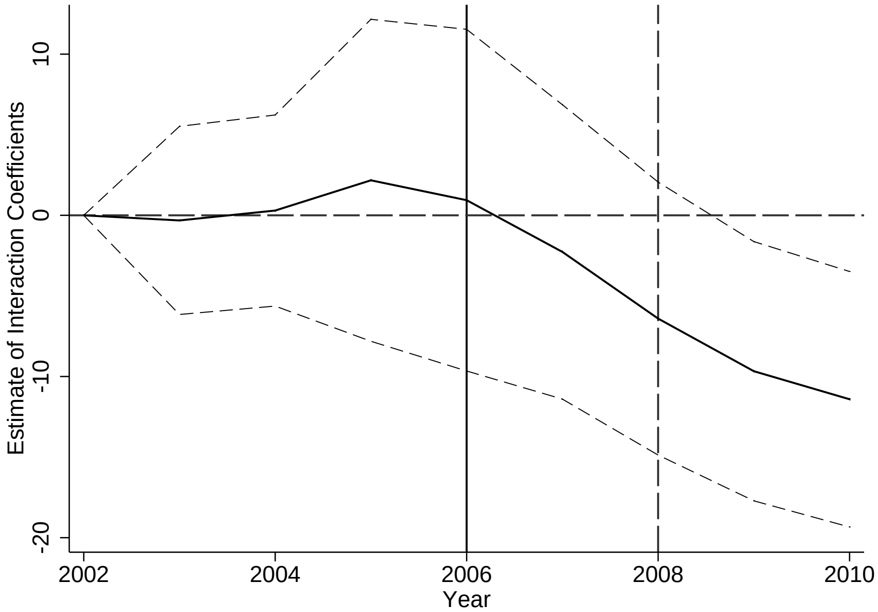

Notes: The solid line plots the point estimate for yearly coefficient estimates of the interaction coefficients $\sum_{t=1}^{\tau} \beta_{1t}$ in Equation (2) for the official SO_2 emissions data. The dashed lines show the 95% confidence bands, based on heteroskedasticity-robust standard errors clustered at the province-level. The excluded time period $t = 1$ is the year 2002. The vertical line before 2006 marks the start of the policy and the vertical line before 2008 marks the start of SO_2 monitoring.

Figure 4: DYNAMIC TREATMENT EFFECTS: SO_2 SATELLITE DATA

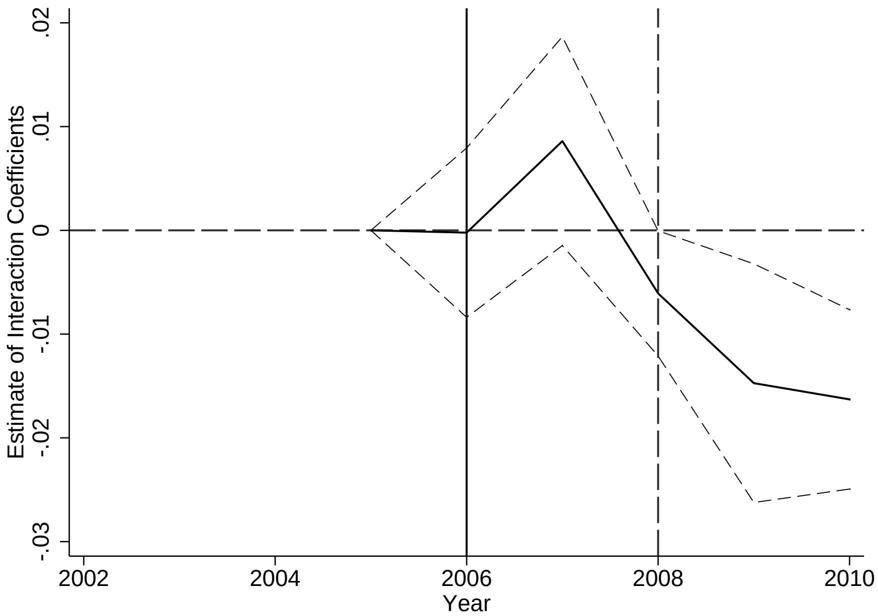

Notes: The solid line plots the point estimate for yearly coefficient estimates of the interaction coefficients $\sum_{t=1}^{\tau} \beta_{1t}$ in Equation (2) for the NASA SO_2 satellite data. The dashed lines show the 95% confidence bands, based on heteroskedasticity-robust standard errors clustered at the province-level. The excluded time period $t = 1$ is the year 2005. The vertical line at 2006 marks the start of the policy and the vertical line at 2008 marks the start of SO_2 monitoring.

Figure 5: Dynamic treatment effects: SO_2 satellite data (Monthly)

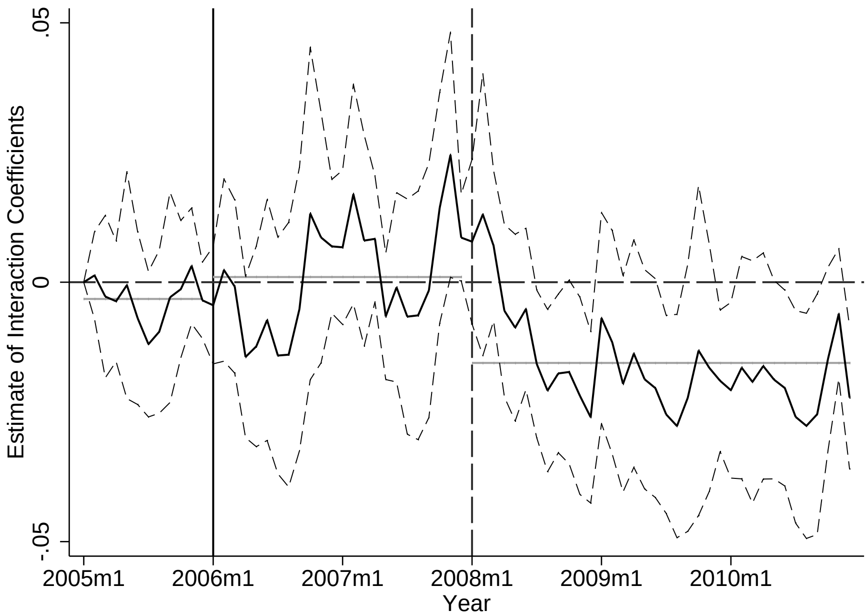

Notes: The solid line plots the point estimate for monthly coefficient estimates of the interaction coefficients $\sum_{t=1}^{\tau} \beta_{1t}$ in Equation (2) for the NASA SO_2 satellite data. The dashed lines show the 95% confidence bands, based on heteroskedasticity-robust standard errors clustered at the province-level. The horizontal grey lines plot the average of the point estimates of the interaction coefficient in the pre-period (2005), the period without monitoring (2006-2007) and the period with monitoring (2008-2010). The excluded time period $t = 1$ is January 2005. The vertical line at 2006 marks the start of the policy and the vertical line at 2008 marks the start of SO_2 monitoring.

Figure 6: EXAMPLE OF A GOVERNMENT WORK REPORT

2003年省政府工作报告

——2003年1月20日在辽宁省第十届人民代表大会第一次会议上

省长　薄熙来

各位代表：

现在，我代表省人民政府，向大会报告工作，请予审议，并请省政协委员和列席同志提出意见。

一、本届政府五年工作的回顾

过去的五年，正处世纪之交。全省各族人民在辽宁省委的领导下，高举邓小平理论伟大旗帜，认真贯彻“三个代表”重要思想，同心协力，艰苦奋斗，开拓进取，克服了种种困难，改革开放和现代化建设取得了显著成就。

——国民经济取得新进展。国内生产总值2002年完成5458亿元，五年年均增长8.9%。社会消费品零售总额2002年实现2258亿元，五年年均增长9.3%；地方财政一般预算收入2002年完成398.8亿元，年均增长12%。固定资产投资五年共完成6449亿元。

Notes: The graph shows the beginning of the government work report for Liaoning province in 2003. It was delivered by then-governor Bo Xilai.

Figure 7: Political attention to air pollution over time

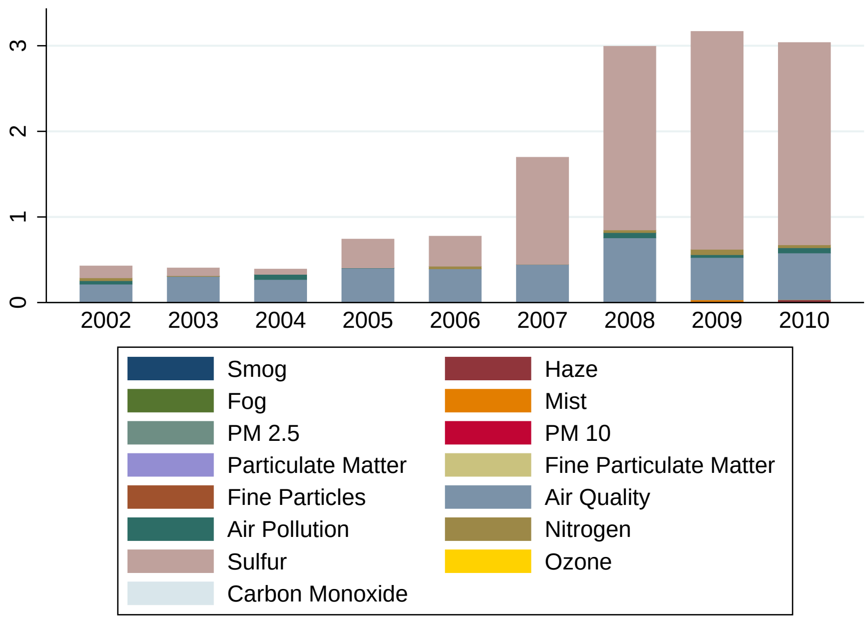

Notes: The graph shows the mean count of keywords related to air pollution for all provincial government work reports in a given year from 2002 to 2010.

Figure 8: Dynamic treatment effects: rhetorical compliance

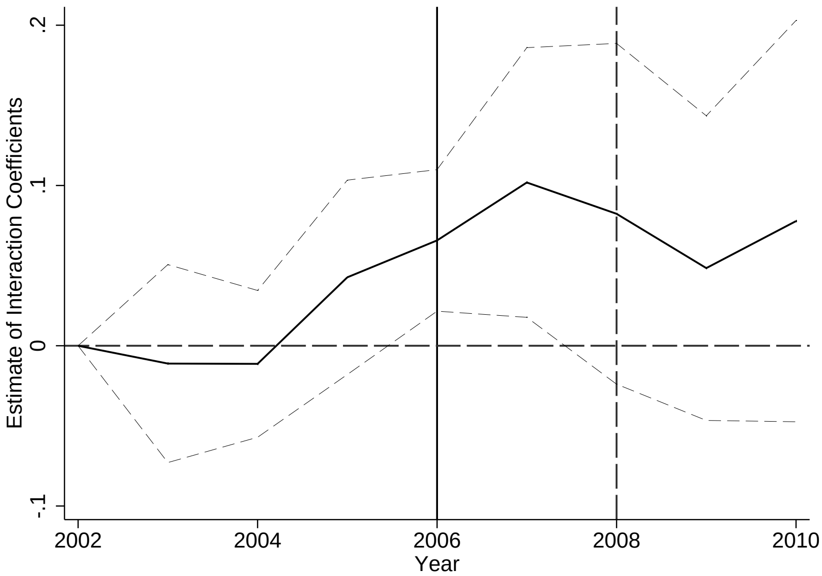

Notes: The solid line plots the point estimate for yearly coefficient estimates of the interaction coefficients $\sum_{t=1}^{T} \beta_{1t}$ in Equation (2) for the number of keywords related to air pollution in a province-year government work report. The dashed lines show the 95% confidence bands, based on heteroskedasticity-robust standard errors clustered at the province-level. The excluded time period $t = 1$ is the year 2002. The vertical line at 2006 marks the start of the policy and the vertical line at 2008 marks the start of SO_2 monitoring.

Figure 9: Dynamic treatment effects: rhetorical compliance (preceding period)

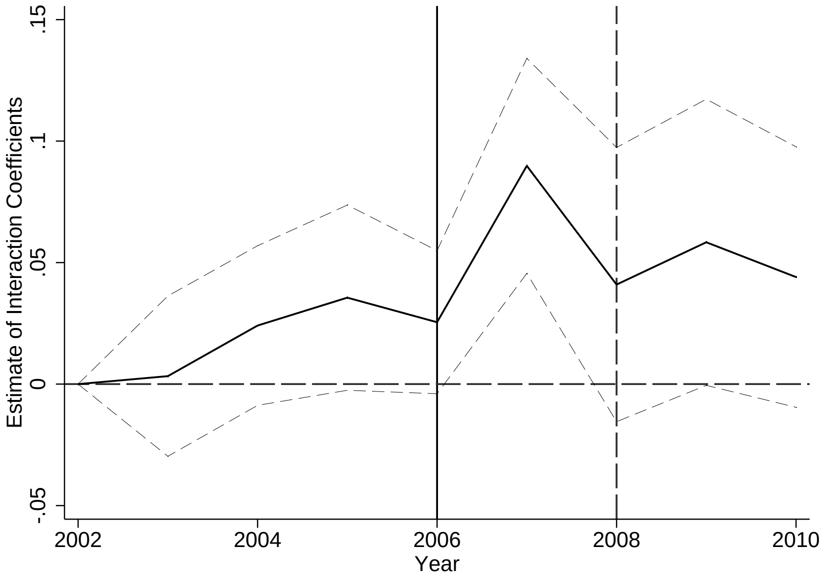

Notes: The solid line plots the point estimate for yearly coefficient estimates of the interaction coefficients $\sum_{t=1}^{T} \beta_{1t}$ in Equation (2) for the number of keywords related to air pollution in a province-year government work report *for work done in period preceding the report.* The dashed lines show the 95% confidence bands, based on heteroskedasticity-robust standard errors clustered at the province-level. The excluded time period $t = 1$ is the year 2002. The vertical line at 2006 marks the start of the policy and the vertical line at 2008 marks the start of SO_2 monitoring.

Figure 10: Dynamic treatment effects: rhetorical compliance (subsequent period)

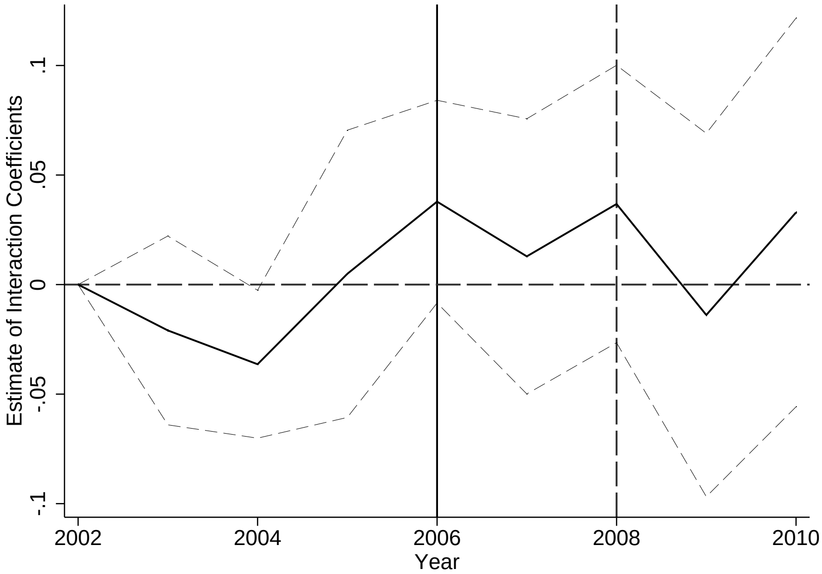

Notes: The solid line plots the point estimate for yearly coefficient estimates of the interaction coefficients $\sum_{t=1}^{T} \beta_{1t}$ in Equation (2) for the number of keywords related to air pollution in a province-year government work report *for work announced for the period following the report.* The dashed lines show the 95% confidence bands, based on heteroskedasticity-robust standard errors clustered at the province-level. The excluded time period $t = 1$ is the year 2002. The vertical line at 2006 marks the start of the policy and the vertical line at 2008 marks the start of SO_2 monitoring.

Figure 11: Real compliance: Installation of desulfurization devices

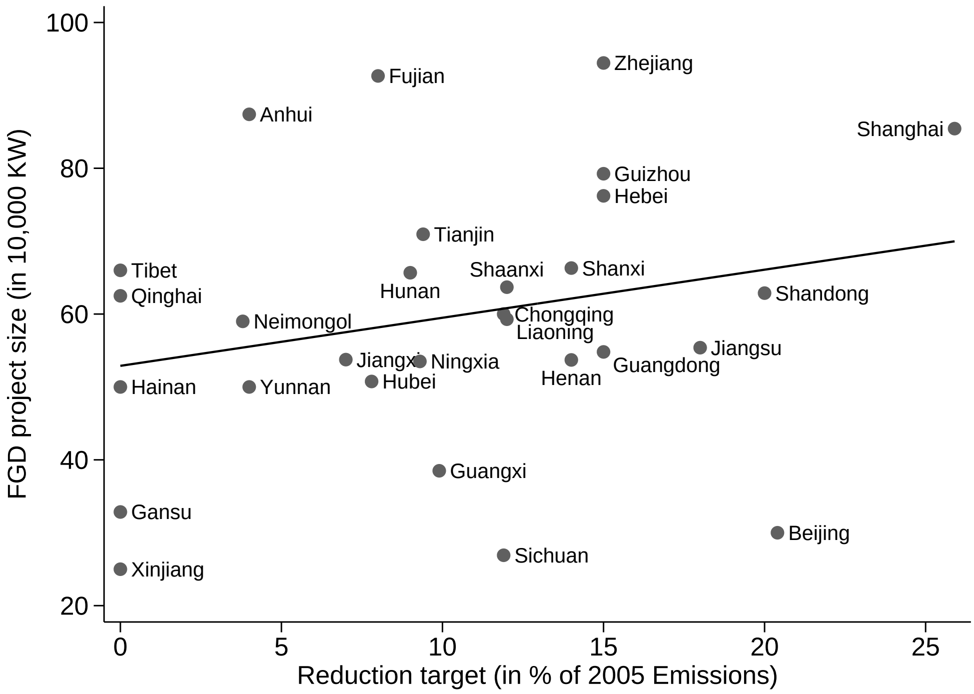

Notes: This graph plots the relationship between the SO_2 reduction target and the planned installation of desulfurization devices at the province level. The solid line fits a linear regression with slope parameter b=0.66 and p=022 computed from standard errors clustered at the province level.

Figure 12: Real compliance: Installation of desulfurization devices

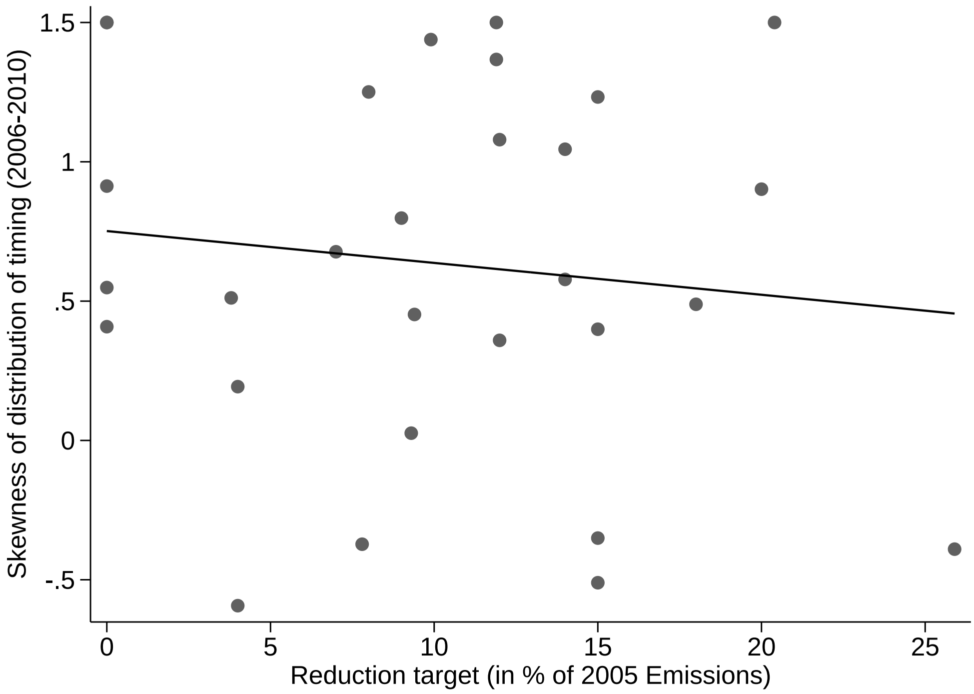

Notes: This graph tests whether provinces with a higher reduction target installed the planned desulfurization devices earlier. The vertical axis shows the skewness for each province of the 5 yearly observations from 2006-2010, where the weight is the capacity (in 10,000 KW) of planned desulfurization devices in each year. The solid line fits a linear regression with slope parameter b=-0.01 and p=0.56 computed from standard errors clustered at the province level.

Figure 13: Real compliance: Shutdown of small thermal units

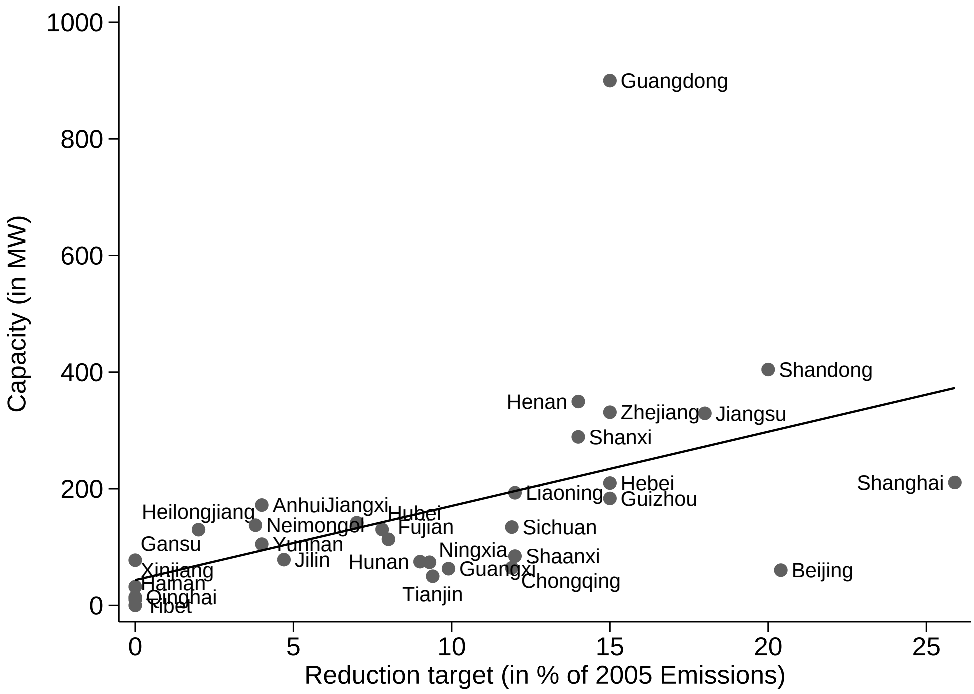

Notes: This graph plots the relationship between the SO_2 reduction target and the decommissioning of small thermal units at the province level. The solid line fits a linear regression with slope parameter b=12.72*** and p=0.004 computed from standard errors clustered at the province level.

Figure 14: Marginal abatement cost curve examples

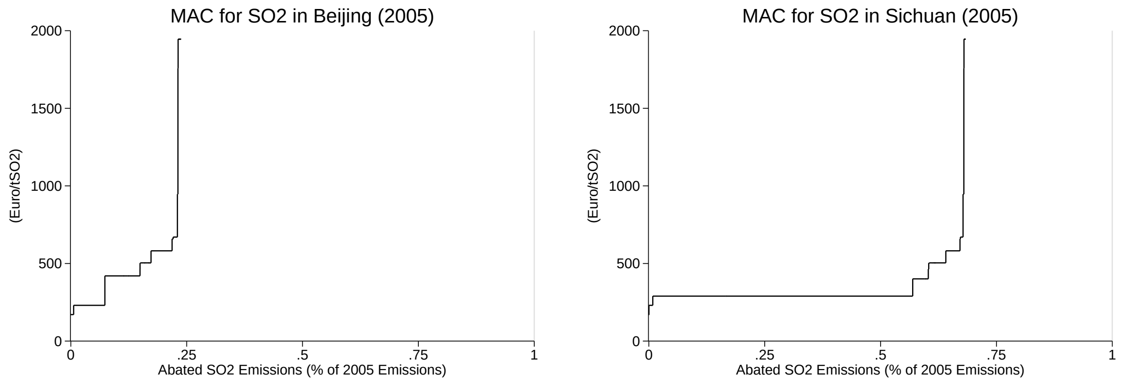

Notes: This graph shows the marginal SO_2 abatement cost curves for Beijing and Sichuan in 2005. The horizontal axis lists the abatement intensity relative to the 2005 SO_2 emissions level.

Figure 15: CONSISTENCY OF SO_2 MARGINAL ABATEMENT COST DATA AND OFFICIAL DATA

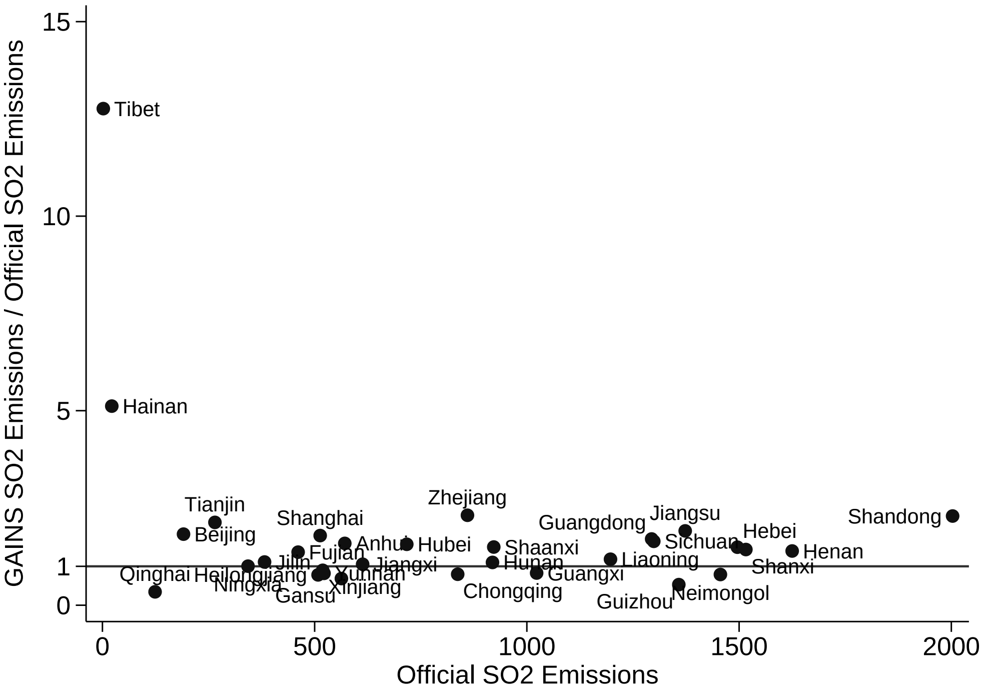

Notes: This graph shows the ratio of the SO_2 emissions data underlying the construction of the marginal abatement cost curve (from IIASA's GAINS model) and the official SO_2 emissions data from the Ministry of Environmental Protection for the year 2005 on the vertical axis. The datapoints are ordered according to the official SO_2 emissions level on the horizontal axis. The raw correlation between both data sources is 85.9%.

Figure 16: SO_2 MARGINAL ABATEMENT COST CURVE FOR CHINA

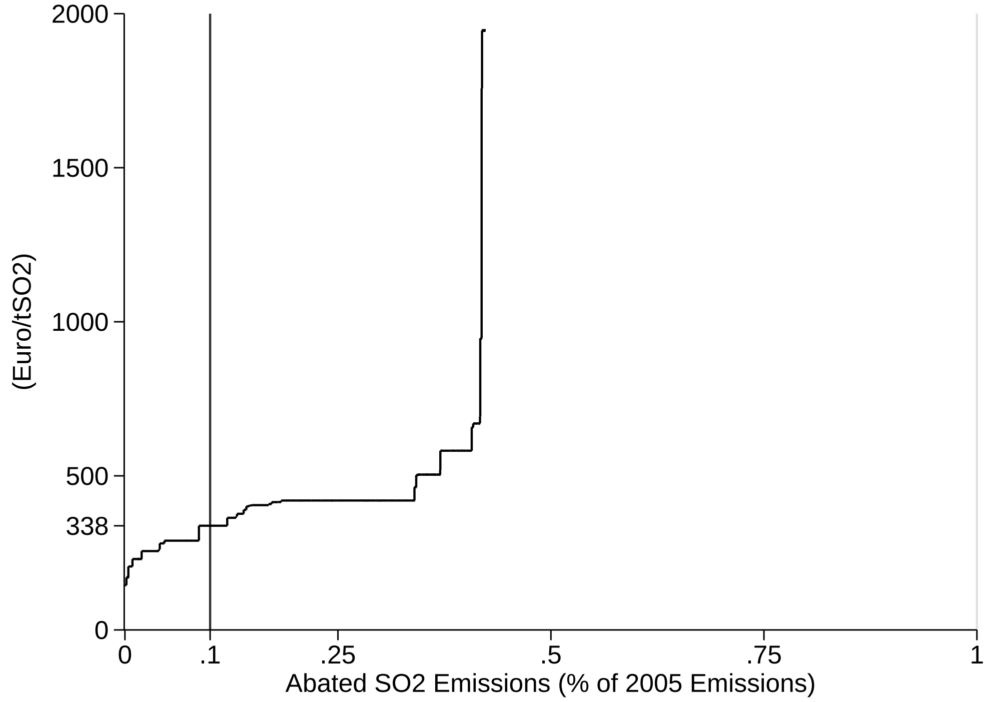

Notes: This graph shows the marginal SO_2 abatement cost curve for China in 2005. The horizontal axis plots the abatement intensity relative to the 2005 SO_2 emissions level. The vertical line at an abatement level of 10% illustrates the 10% national SO_2 emissions control target for China as a whole. Its intersection with the marginal abatement cost curves shows that the counterfactual marginal abatement cost for this abatement level is 338€/tSO_2.

Figure 17: THE COST-EFFICIENT ALLOCATION OF SO_2 REDUCTION TARGETS

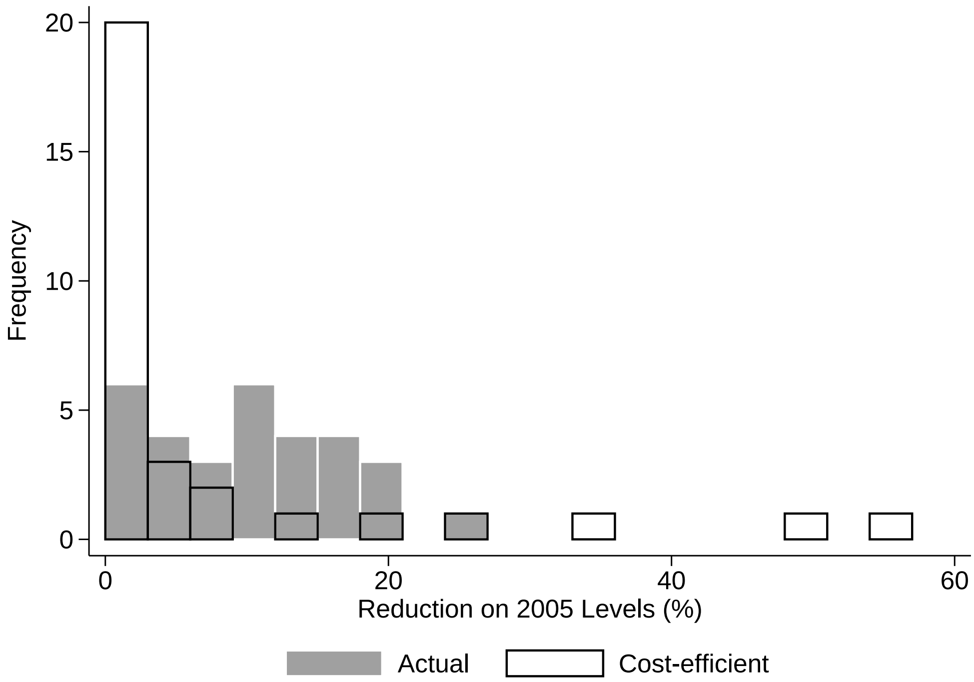

Notes: This graph shows the counterfactual cost-efficient allocation of SO_2 reduction targets under the 10% SO_2 total control target of the 11th Five-Year Plan (2006-2010), shown in white boxes. The three provinces with the highest targets under the cost-efficient allocation are Sichuan, Shandong and Zhejiang. Solid grey boxes show the actual allocation of reduction targets in comparison.

Figure 18: Difference between cost-efficient and actual allocation

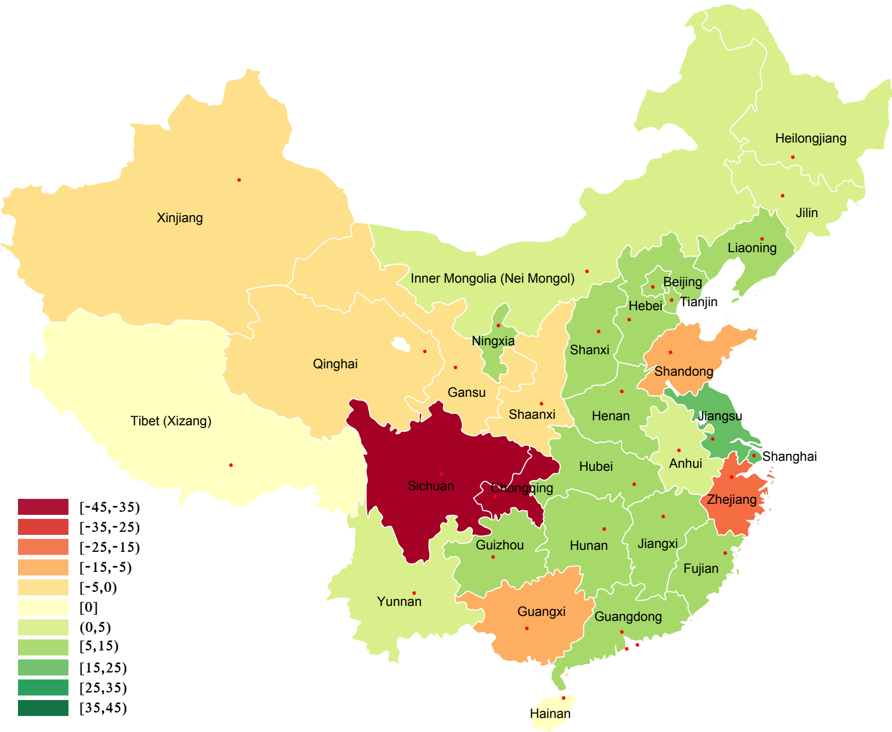

Notes: The graph shows the difference in percentage points between the SO_2 reduction target of each province under the actual minus the cost-efficient allocation (data from Table 9).

Figure 19: RELATION BETWEEN MAC AND SO_2 REDUCTION TARGETS

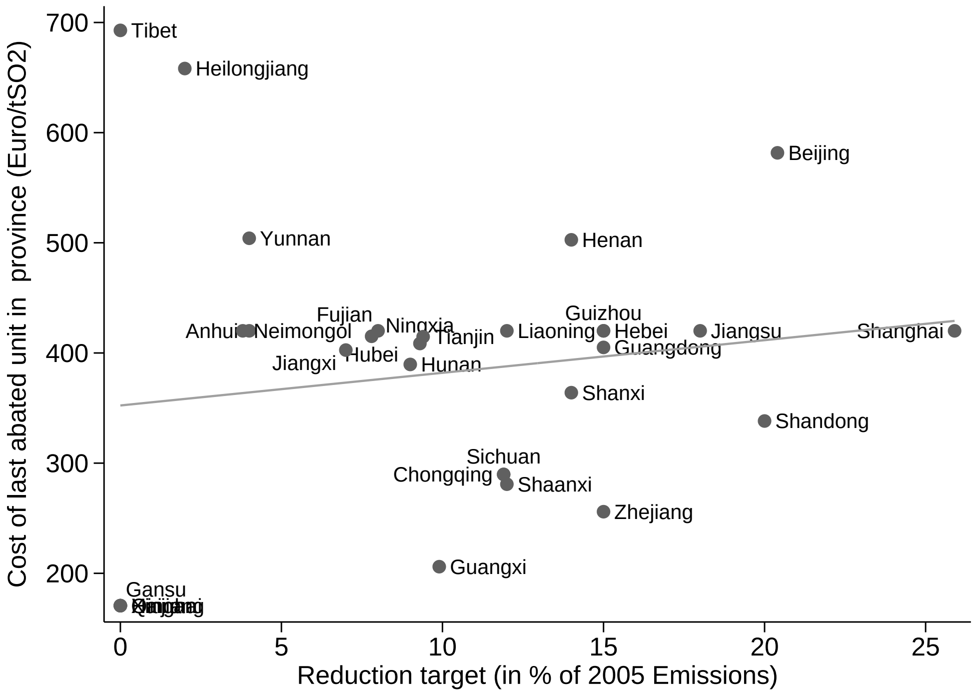

Notes: This graph shows the lack of correlation between a province's actual SO_2 reduction target in the 11th Five-Year Plan (2006-2010) and its marginal abatement cost at the level of the target. The solid line fits a linear regression with slope parameter b=2.96 and p=0.54 computed from standard errors clustered at the province level.

Figure 20: Marginal welfare impact and SO_2 reduction targets

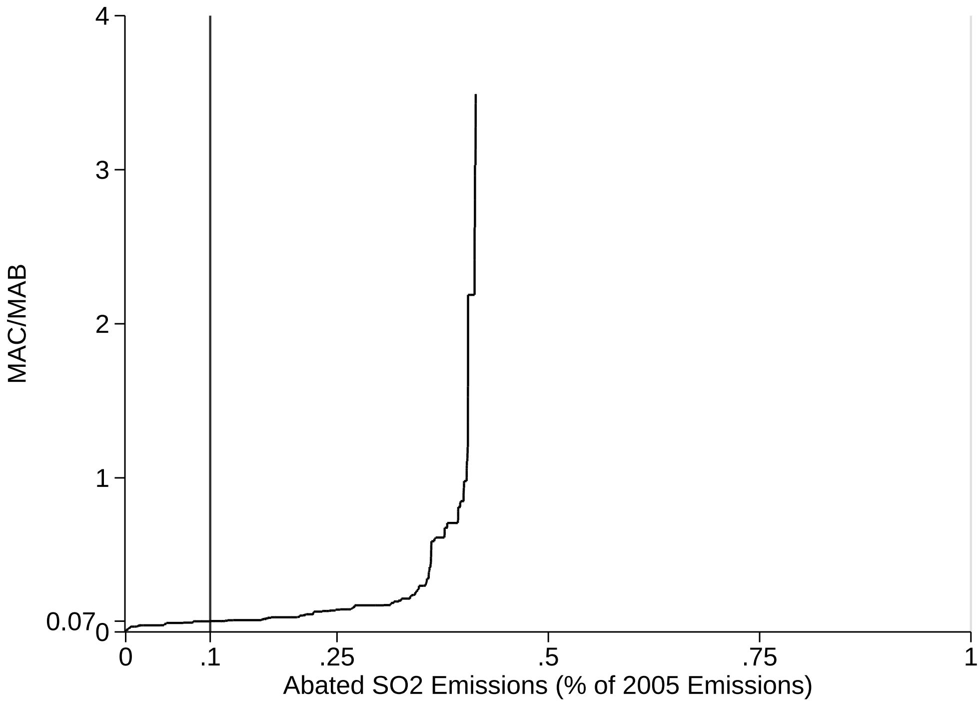

Notes: This graph shows the ratio of marginal abatement cost to marginal abatement benefits for SO_2 for China in 2005. The horizontal axis plots the abatement intensity relative to the 2005 SO_2 emissions level. The horizontal, longdashed line at 1 marks the welfare-optimal level of SO_2 abatement (40.3%). The vertical, shortdashed line at 0.1 marks the ratio of marginal abatement cost to marginal abatement benefit for the 10% SO_2 reduction target (0.07). Prohibitive cost ranges beyond a MAC/MAB ratio of 3.5 not shown.

Figure 21: Economic growth and air pollution in China

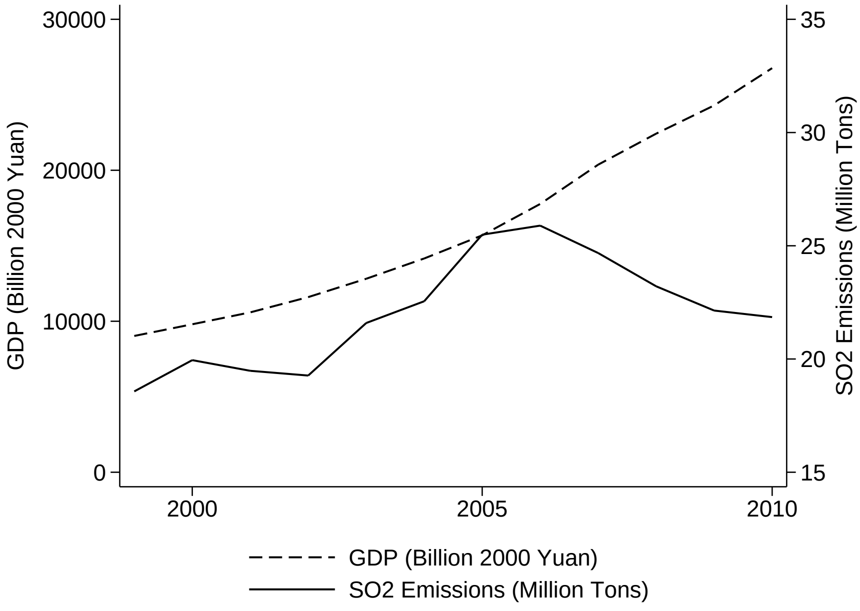

Notes: This figure shows the recent timeseries in GDP (in billion constant 2000 yuan) and in SO_2 emissions (in million tons) for China from 1999 to 2010. Data source: Fridley, Romankiewicz, and Fino-Chen (2013).

Tables

Table 1: Descriptive statistics

	Mean	Std. Dev.	Min.	Max.	Time Period
SO2 Emissions (kt)	739	471	1	2003	2002-2010
SO2 Satellite (DU)	0.37	0.39	-0.02	1.93	2005-2010
Selected SO2 Sat. (DU)	0.72	0.53	-0.03	2.57	2005-2010
SO2 Reduction Target (% on 2005 Baseline)	9.65	6.71	0	25.9	-

Note: Satellite data is measured in Dobson Units (DU). Selected SO_2 Sat. is the sample of polluted cities. Satellite SO_2 measurements below 0.2 Dobson Units are generally considered as clean air and negative values are likely noise from measurement error. Replacing negative values as either 0s or missing does not change the subsequent results.

Table 2: The effect of the policy for the whole period (2006-2010)

	SO2 Emissions (Kt)	Satellite SO2 (Dobson Units)	Selected Sat. SO2 (Dobson Units)
Reductiontarget × D(Post)	-6.3***	-0.0057	-0.0115*
	(1.6)	(0.0040)	(0.0050)
	[0.01]	[0.27]	[0.06]
Reductiontarget	0.3	0.0214***	0.0251***
	(1.34)	(0.0067)	(0.0076)
	[0.74]	[0.00]	[0.00]
D(Post)	107.87***	0.0482	0.0838**
	(17.7)	(0.0286)	(0.0386)
	[0.00]	[0.12]	[0.04]
Province FE	✓	✓	✓
Effect size (% of mean/σ)	*5.8%*	*10.5%*	*10.9%*
Observations	279	186	186
Provinces	31	31	31
R^2	0.96	0.90	0.84

Note: Heteroskedasticity-robust standard errors clustered at the province level are shown in parentheses. Wild bootstrap p-values from 1000 repetitions clustered at the province level are shown in square brackets. The effect size gives the estimated coefficient of the interaction term β_1 in Equation (1) as percentage of the mean of the dependent variable for a one standard deviation(σ)-increase in the SO_2 reduction target. $***p < 0.01$, $**p < 0.05$, $*p < 0.1$.

Table 3: THE EFFECT OF THE POLICY FOR EACH YEAR OF THE 11TH FIVE-YEAR PLAN (2006-2010)

	SO2 Emissions (Kt)	Satellite SO2 (Dobson Units)	Selected Sat. SO2 (Dobson Units)
2002× Reductiontarget	Excluded	-	-
2003× Reductiontarget	-0.32	-	-
	(3.93)		
	[0.93]		
2004× Reductiontarget	0.30	-	-
	(3.64)		
	[0.93]		
2005× Reductiontarget	2.17	Excluded	Excluded
	(4.04)		
	[0.66]		
2006× Reductiontarget	0.94	-0.00	0.00
	(4.08)	(0.00)	(0.01)
	[0.86]	[0.93]	[0.94]
2007× Reductiontarget	-2.27	0.01	0.01
	(3.70)	(0.00)	(0.01)
	[0.62]	[0.21]	[0.29]
2008× Reductiontarget	-6.40	-0.01	-0.01***
	(3.69)	(0.00)	(0.01)
	[0.14]	[0.11]	[0.00]
2009× Reductiontarget	-9.67**	-0.01**	-0.02**
	(3.79)	(0.00)	(0.00)
	[0.02]	[0.02]	[0.01]
2010× Reductiontarget	-11.41***	-0.02***	-0.03***
	(3.89)	(0.00)	(0.00)
	[0.00]	[0.00]	[0.00]
Province FE	✓	✓	✓
Mean dep. var.	739	0.37	0.72
Observations	279	186	186
Provinces	31	31	31
R^2	0.98	0.96	0.93

Note: Heteroskedasticity-robust standard errors clustered at the province level are shown in parentheses. Wild bootstrap p-values from 1000 repetitions clustered at the province level are shown in square brackets. The yearly interaction coefficients are estimates for $\sum_{t=1}^{T} \beta_{1t}$ in Equation (2), while 'Excluded' is the omitted time period. Note that the SO_2 reduction started in 2006 while government monitoring of SO_2 pollution became effective in 2008. $***p < 0.01$, $**p < 0.05$, $*p < 0.1$.

Table 4: Heterogeneous treatment effects depending on initial pollution levels

	SO2 Satellite Data			
	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Reductiontarget × D(Post)	0.0004	-0.0032	-0.0082	-0.0119*
	(0.0028)	(0.0030)	(0.0055)	(0.0053)
	[0.93]	[0.41]	[0.26]	[0.07]
Reductiontarget	0.0311***	0.0248***	0.0224***	0.0210***
	(0.0023)	(0.0025)	(0.0046)	(0.0044)
	[0.00]	[0.00]	[0.00]	[0.00]
D(Post)	0.0356*	0.0328	0.0382	0.0293
	(0.0178)	(0.0173)	(0.0274)	(0.0365)
	[0.06]	[0.10]	[0.26]	[0.47]
Province FE	✓	✓	✓	✓
Effect size (% of mean/σ)	-1.5%	5.8%	9.9%	9.4%
Observations	47,376	47,592	47,016	48,096
Provinces	31	31	31	31
R-squared	0.54	0.60	0.59	0.61

Note: The table reports estimates for Equation (1) at the county-yearmonth level for different subsamples. Quartile marks the quartile of initial based on its 2005 SO_2 pollution, calculated from SO_2 satellite data, relative to the mean SO_2 pollution within the same province. Heteroskedasticity-robust standard errors clustered at the province level are shown in parentheses. Wild bootstrap p-values from 1000 repetitions clustered at the province level are shown in square brackets. $***p < 0.01$, $**p < 0.05$, $*p < 0.1$.

Table 5: Controlling for marginal abatement cost and marginal welfare impact

	SO2 Emissions (Kt)		SO2 Satellite (Dobson Units)		Sel. Satellite SO2 (Dobson Units)	
Reductiontarget × D(Post)	-6.26**	-6.15***	-0.0057	-0.0058	-0.0116*	-0.0117**
	(1.62)	(1.63)	(0.0040)	(0.0041)	(0.0050)	(0.0051)
	[0.01]	[0.00]	[0.23]	[0.23]	[0.06]	[0.04]]
Reductiontarget	0.33	1.33	0.0214***	0.0364***	0.0252***	0.0461***
	(1.36)	(1.44)	(0.0067)	(0.0055)	(0.0077)	(0.0062)
	[0.81]	[0.22]	[0.00]	[0.00]	[0.00]	[0.00]
MAC × D(Post)	0.01		-0.0000		-0.0000	
	(0.02)		(0.0000)		(0.0000)	
	[0.55]		[0.90]		[0.19]	
MAC	-0.12***		-0.0003***		-0.0003***	
	(0.02)		(0.0000)		(0.0001)	
	[0.00]		[0.00]		[0.00]	
MAC/MAB × D(Post)		-0.59		0.0001		0.0009
		(0.61)		(0.0015)		(0.0016)
		[0.45]		[0.90]		[0.39]
MAC/MAB		0.71*		0.0053***		0.0068***
		(0.60)		(0.0011)		(0.0012)
		[0.08]		[0.00]		[0.00]
D(Post)	99.97**	113.20***	0.0492	0.0470*	0.0992**	0.0753*
	(21.02)	(20.20)	(0.0316)	(0.0271)	(0.0465)	(0.0395)
	[0.02]	[0.00]	[0.11]	[0.08]	[0.03]	[0.07]
Province FE	✓	✓	✓	✓	✓	✓
Observations	279	279	186	186	186	186
Provinces	31	31	31	31	31	31
R^2	0.96	0.96	0.90	0.90	0.84	0.84

Note: Heteroskedasticity-robust standard errors clustered at the province level are shown in parentheses. Wild bootstrap p-values from 1000 repetitions clustered at the province level are shown in square brackets. The table reports the results from estimating Equation (3). MAC refers to the marginal abatement cost given the actual SO_2 reduction target and MAC/MAB is the ratio of the marginal abatement benefits to the marginal abatement cost. $***p < 0.01$, $**p < 0.05$, $*p < 0.1$.

Table 6: The effect of monitoring: Overview

y_{pt}	Misreporting Unlikely	Misreporting Likely	Hypothesis
SO2 emissions	$\hat{\beta}_1$	$\hat{\beta}_1 + \hat{\beta}_7$	$\hat{\beta}_7 < 0$
SO2 satellite data	$\hat{\beta}_1$	$\hat{\beta}_1 + \hat{\beta}_7$	$\hat{\beta}_7 > 0$
Normalized difference	$\hat{\beta}_1$	$\hat{\beta}_1 + \hat{\beta}_7$	$\hat{\beta}_7 > 0$

Note: This table provides an overview on how to interpret the estimated effect of the SO_2 reduction target for each province from separate estimations of Equation (4) for both the official SO_2 emissions indicator and the SO_2 satellite data. The sample runs through the first policy period until 2008 when the government could not monitor SO_2 emissions. The column labelled 'hypothesis' captures the ex ante expected coefficient for the triple interaction term under the hypothesis of misreporting. 'Normalized difference' captures the difference between the SO_2 satellite data and the SO_2 emissions data, both normalized to 2005 levels.

Table 7: The effect of monitoring: Evidence

Misreporting likely due to	y_{pt}	$\hat{\beta}_7$	Significance $\hat{\beta}_7$
Career concerns	SO2 emissions	-0.17	p=0.99
	SO2 satellite	0.008	p=0.38
	Normalized difference	0.025	p=0.35
Marginal abatement cost	SO2 emissions	-3.08	p=0.52
	SO2 satellite	-0.009	p=0.29
	Normalized difference	0.008	p=0.34
Marginal welfare impacts	SO2 emissions	7.64	p=0.13
	SO2 satellite	0.002	p=0.80
	Normalized difference	0.002	p=0.86

Note: The table shows the estimated effect of the SO_2 reduction target for each province from separate estimations of Equation (4). The outcome variables are the official SO_2 emissions indicator, the SO_2 satellite data as well as the 'Normalized difference', which captures the difference between the SO_2 satellite data and the SO_2 emissions data, both normalized to 2005 levels. The sample is for the first policy period until 2008 when the government could not monitor SO_2 emissions. The significance of $\hat{\beta}_7$ is the p-value using wild bootstrap p-values from 1000 repetitions clustered at the province level.

Table 8: Reactions by the regulated agents in response to the targets

	Political attention to air pollution			Placebo outcomes	
	All statements	Past period	Future period	'Forest'	'Medical care'
Reductiontarget	0.07*	0.03**	0.04**	-0.12	0.02
× D(Post)	(0.02)	(0.01)	(0.02)	(0.07)	(0.07)
	[0.08]	[0.03]	[0.03]	[0.31]	[0.79]
Reductiontarget	-0.01	-0.00	-0.00	-0.07	0.35***
	(0.03)	(0.02)	(0.02)	(0.10)	(0.07)
	[0.74]	[0.66]	[0.78]	[0.40]	[0.00]
D(Post)	1.19***	0.48**	0.71**	3.59**	5.68***
	(0.24)	(0.13)	(0.16)	(0.96)	(0.79)
	[0.00]	[0.01]	[0.01]	[0.04]	[0.00]
Province FE	✓	✓	✓	✓	✓
Effect size	30.9%	29.6%	31.5%	-7.2%	1.2%
Observations	274	273	273	274	274
Provinces	31	31	31	31	31
R^2	0.51	0.42	0.43	0.69	0.52

Note: Heteroskedasticity-robust standard errors clustered at the province level are shown in parentheses. Wild bootstrap p-values from 1000 repetitions clustered at the province level are shown in square brackets. The effect size gives the estimated coefficient of the interaction term β_1 in Equation (1) as percentage of the mean of the dependent variable for a one standard deviation(σ)-increase in the SO_2 reduction target. $***p < 0.01$, $**p < 0.05$, $*p < 0.1$.

Table 9: Actual and counterfactual SO_2 target allocations

	SO2 reduction target under different allocations (in % of 2005 emissions)			Difference between target allocations (in percentage points)		
	Cost-efficient	Optimal	Actual	Actual - cost-efficient	Actual - optimal	Cost-efficient - optimal
Anhui	1.5	8.3	4	2.5	-4.3	-6.8
Beijing	7.3	22.9	20.4	13.1	-2.5	-15.6
Chongqing	50.5	0	11.9	-38.6	11.9	50.5
Fujian	0.7	0.7	8	7.3	7.3	0
Gansu	2.5	0.8	0	-2.5	-0.8	1.7
Guangdong	0.7	48.3	15	14.3	-33.3	-47.6
Guangxi	20	0	9.9	-10.1	9.9	20
Guizhou	5.8	0	15	9.2	15	5.8
Hainan	0	0	0	0	0	0
Hebei	1.7	0.1	15	13.3	14.9	1.6
Heilongjiang	1.7	0	2	0.3	2	1.7
Henan	0.6	0.1	14	13.4	13.9	0.5
Hubei	2.5	35.2	7.8	5.3	-27.4	-32.7
Hunan	1	1	9	8	8	0
Jiangsu	0.9	25.3	18	17.1	-7.3	-24.4
Jiangxi	1.6	1.6	7	5.4	5.4	0
Jilin	1.6	1.6	4.7	3.1	3.1	0
Liaoning	4.2	0	12	7.8	12	4.2
Neimongol	1.1	0	3.8	2.7	3.8	1.1
Ningxia	0.1	0	9.3	9.2	9.3	0.1
Qinghai	3.5	0	0	-3.5	0	3.5
Shaanxi	13.9	0	12	-1.9	12	13.9
Shandong	25.5	0.6	20	-5.5	19.4	24.9
Shanghai	6.2	6.2	25.9	19.7	19.7	0
Shanxi	0.6	0	14	13.4	14	0.6
Sichuan	56.2	0	11.9	-44.3	11.9	56.2
Tianjin	2.2	52.1	9.4	7.2	-42.7	-49.9
Tibet	0	0	0	0	0	0
Xinjiang	1.6	0	0	-1.6	0	1.6
Yunnan	1.3	1.3	4	2.7	2.7	0
Zhejiang	35.4	35.4	15	-20.4	-20.4	0

Note: This table shows the SO_2 emissions reduction targets under each policy regime. 'Cost-efficient' are the counterfactual targets for achieving the cost-efficient allocation given the national 10% SO_2 reduction target, 'Optimal' are the counterfactual targets for achieving the welfare-optimizing allocation given the national 10% SO_2 reduction target and 'Actual' are the provincial reduction targets used in the 11th Five-Year Plan (2006-2010).

Table 10: Efficiency gains from trade

Allocation	Efficiency measure: Abatement cost at the margin (in €/tSO2)	Welfare impact at the margin (MAC/MAB)
Actual	658	2.19
Cost-efficient	338	1.2

Note: The efficiency measures report the highest cost of all last abated units across provinces for the actual allocation from the 11th Five-Year Plan (2006-2010) and the counterfactual cost-efficient allocation reported in Table 9. Estimates for the cost of the actual allocation exclude outliers. The welfare impact at the margin is the ratio between marginal abatement cost and marginal abatement benefits.

Appendix

A.1 Additional Figures and Tables

Figure A.1: The Great Recession and SO_2 reduction targets

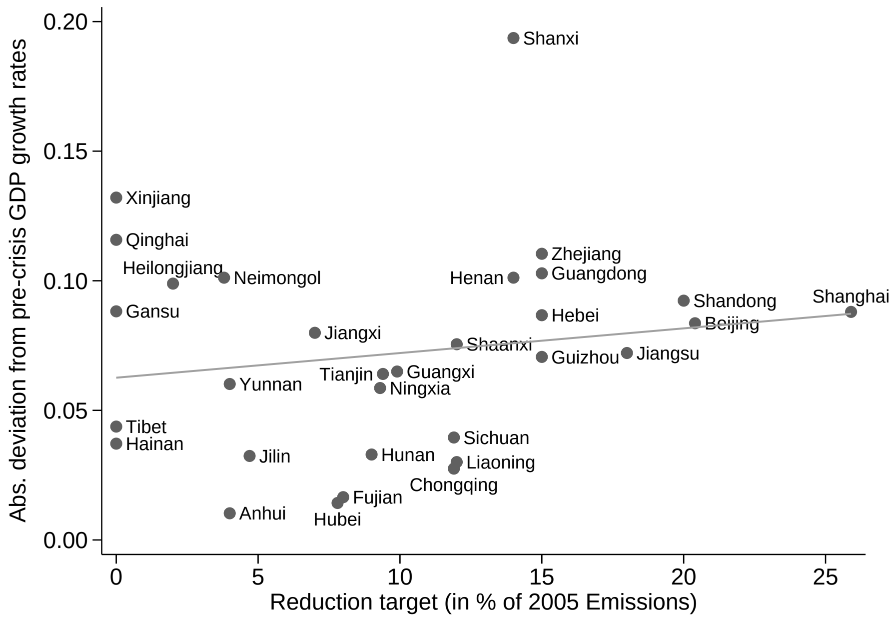

Notes: This graph plots the SO_2 reduction targets on the horizontal axis against the absolute deviation from the pre-crisis (2002-2007) GDP growth rate for each province in China in 2009, the year China was struck by the Great Recession. The solid line fits a linear regression with slope parameter b=0.001 and p=0.305 computed from standard errors clustered at the province level.

Figure A.2: Likely misreporting over time

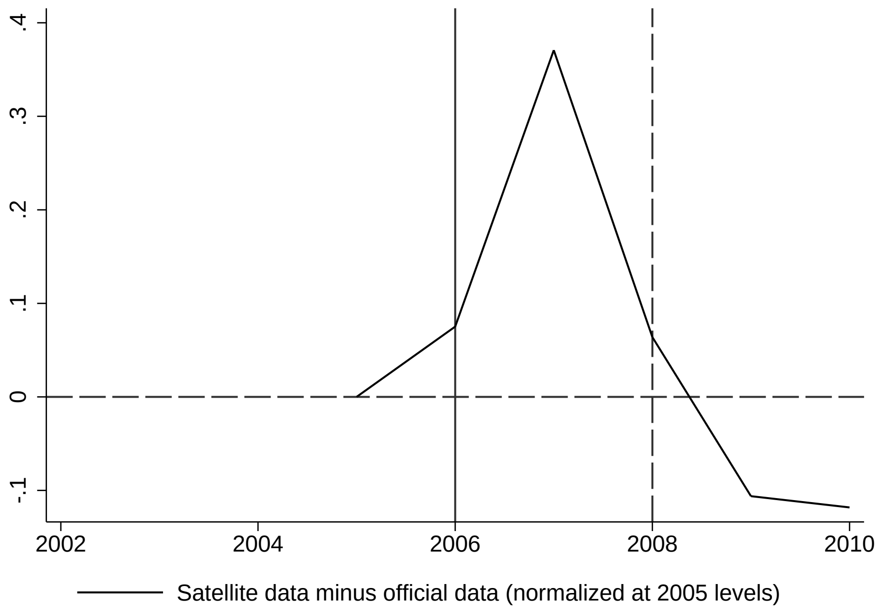

Notes: The graph shows the difference between the SO2 satellite data from NASA and the official SO2 data from the Chinese government. Both variables are normalized to the 2005 levels for China.

Figure A.3: Dynamic treatment effects: Likely misreporting

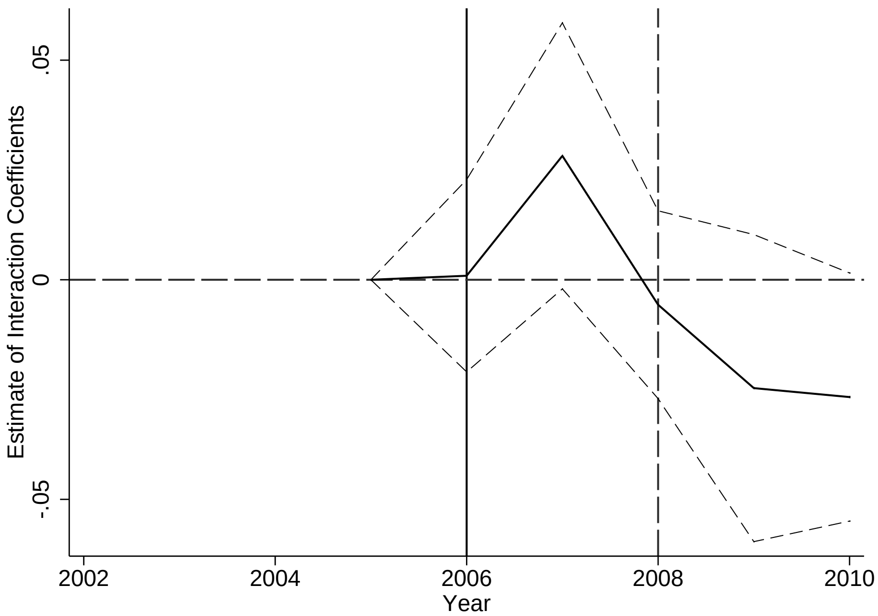

Notes: The solid line plots the point estimate for yearly coefficient estimates of the interaction coefficients $\sum_{t=1}^{T} \beta_{1t}$ in Equation (2) for the difference between the SO2 satellite data from NASA and the official SO2 data from the Chinese government. Both variables are normalized to the 2005 levels for China. The dashed lines show the 95% confidence bands, based on heteroskedasticity-robust standard errors clustered at the province-level. The excluded time period $t = 1$ is the year 2002. The vertical line at 2006 marks the start of the policy and the vertical line at 2008 marks the start of SO_2 monitoring. The p-value for the 2007 interaction is p=0.13 based on the wild bootstrap procedure with 1000 repetitions or p=0.08 based on the cluster-robust standard errors.

Figure A.4: Political attention to air pollution over time (normalized by report length)

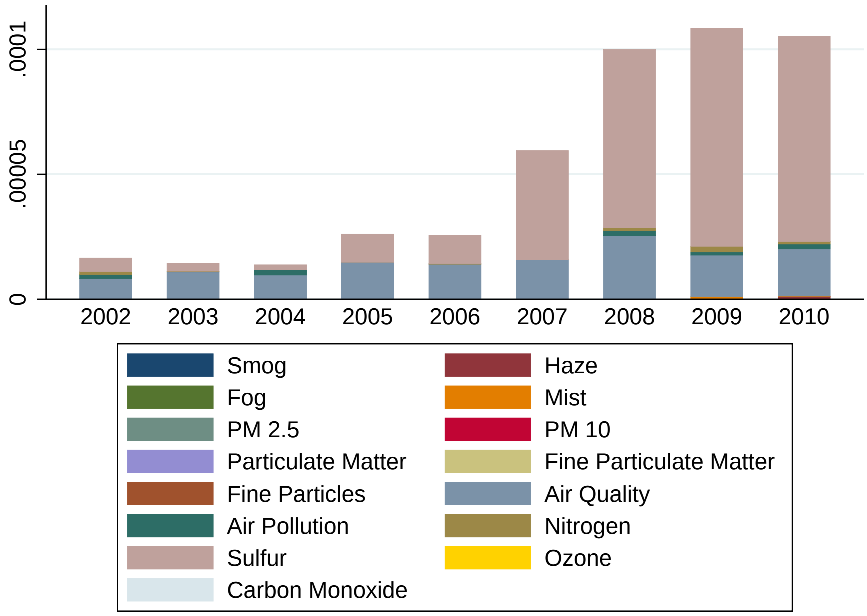

Notes: The graph shows the mean count of keywords related to air pollution for all provincial government work reports in a given year from 2002 to 2010. The count of keywords is normalized by the length of each government work report.

Figure A.5: Relation between welfare impact and SO_2 reduction targets

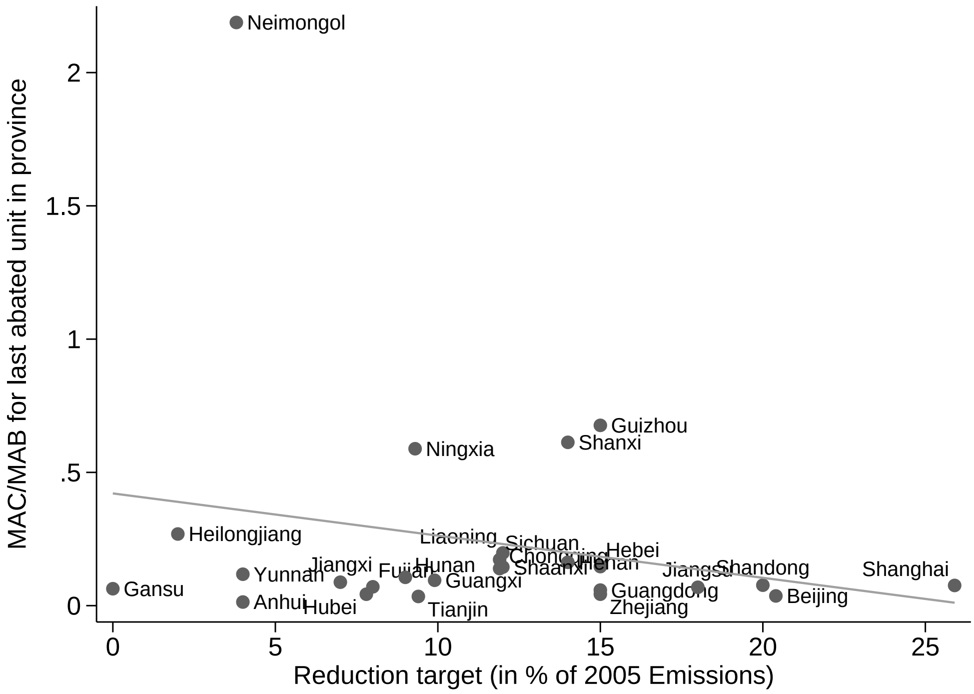

Notes: This graph shows the lack of correlation between a province's actual SO_2 reduction target in the 11th Five-Year Plan (2006-2010) and the marginal welfare impact at the level of the target. The solid line fits a linear regression with slope parameter b=-0.16 and p=0.36 computed from standard errors clustered at the province level.

A.2 Marginal SO_2 abatement cost curves for all provinces in China

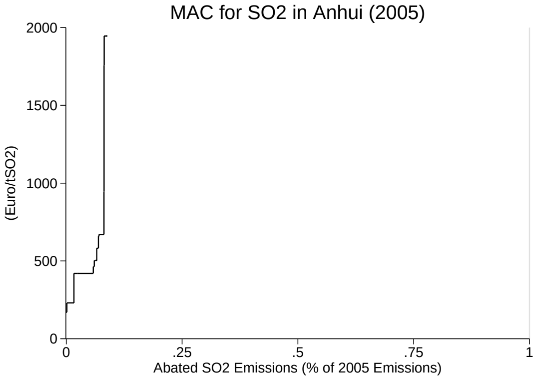

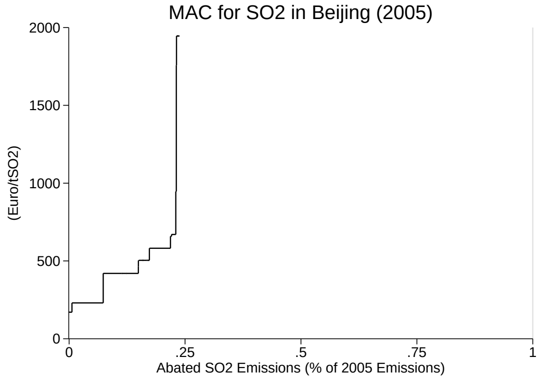

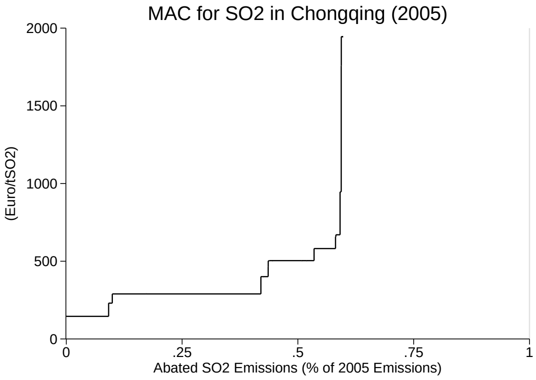

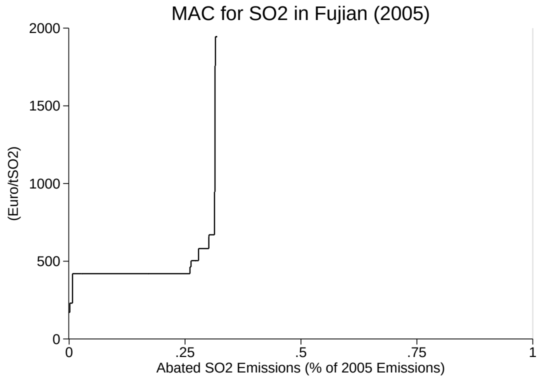

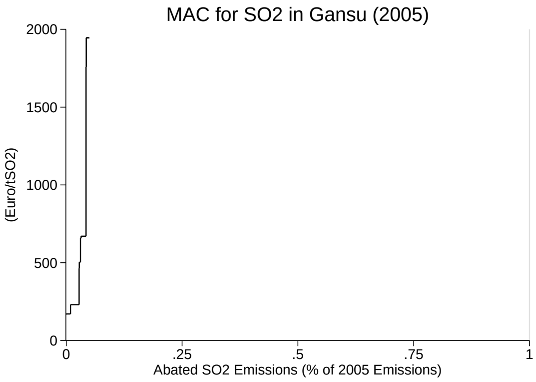

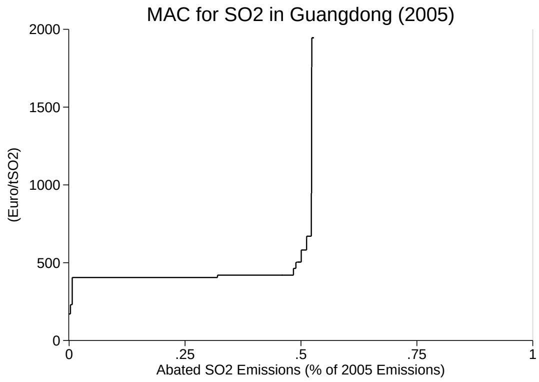

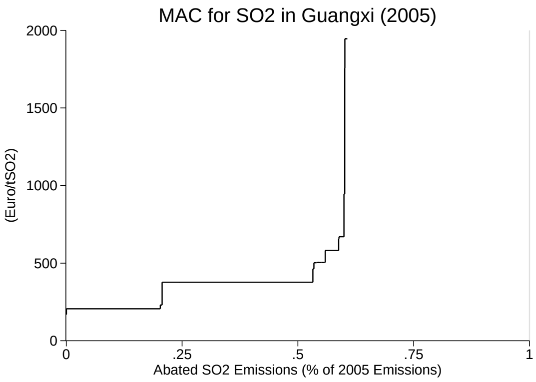

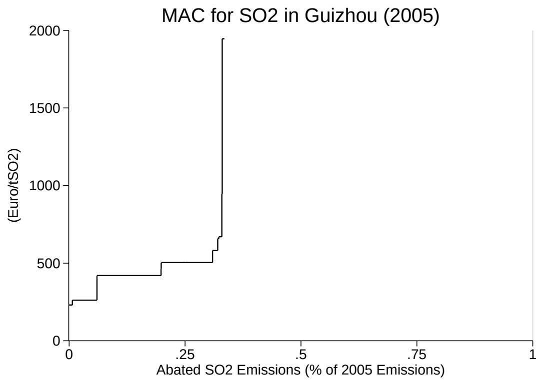

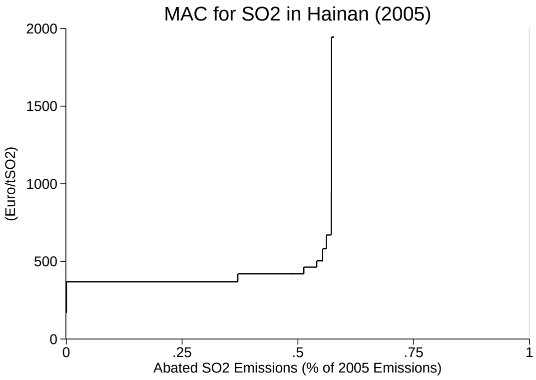
MAC for SO2 in Hainan (2005)
(Euro/tSO2)
0
500
1000
1500
2000
0
.25
.5
.75
1
Abated SO2 Emissions (% of 2005 Emissions)

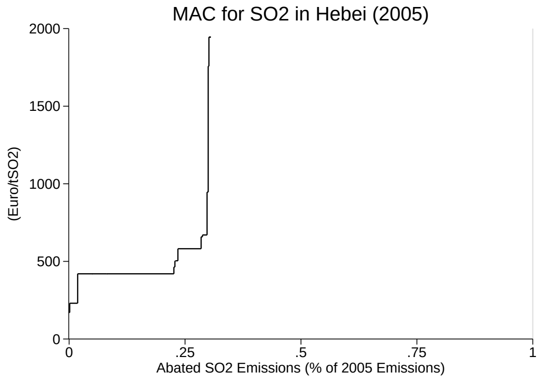
MAC for SO2 in Hebei (2005)
(Euro/tSO2)
0
500
1000
1500
2000
0
.25
.5
.75
1
Abated SO2 Emissions (% of 2005 Emissions)

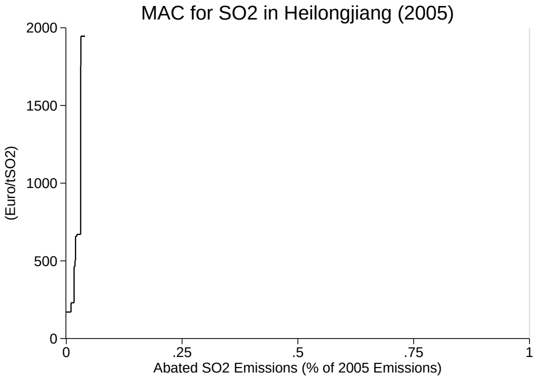
MAC for SO2 in Heilongjiang (2005)
(Euro/tSO2)
0
500
1000
1500
2000
0
.25
.5
.75
1
Abated SO2 Emissions (% of 2005 Emissions)

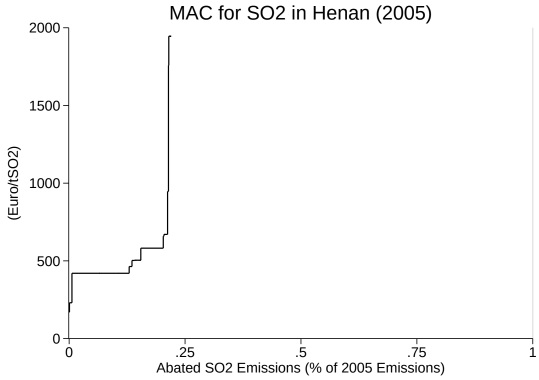
MAC for SO2 in Henan (2005)
(Euro/tSO2)
0
500
1000
1500
2000
0
.25
.5
.75
1
Abated SO2 Emissions (% of 2005 Emissions)

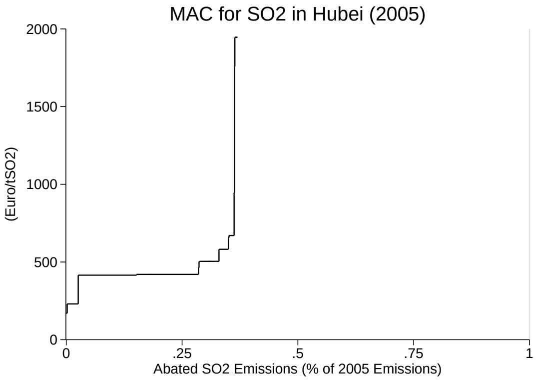
MAC for SO2 in Hubei (2005)
(Euro/tSO2)
0
500
1000
1500
2000
0
.25
.5
.75
1
Abated SO2 Emissions (% of 2005 Emissions)

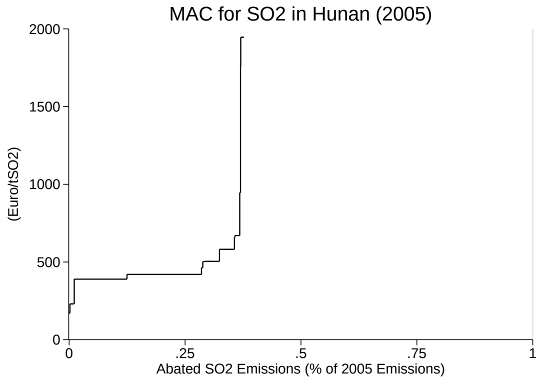
MAC for SO2 in Hunan (2005)
(Euro/tSO2)
0
500
1000
1500
2000
0
.25
.5
.75
1
Abated SO2 Emissions (% of 2005 Emissions)

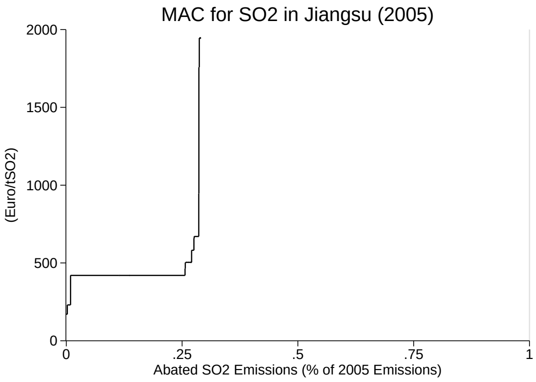
MAC for SO2 in Jiangsu (2005)
(Euro/tSO2)
0
500
1000
1500
2000
0
.25
.5
.75
1
Abated SO2 Emissions (% of 2005 Emissions)

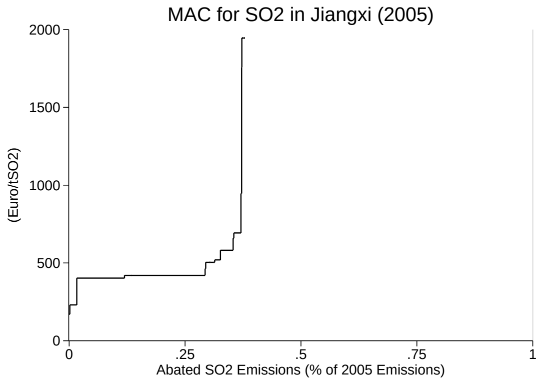
MAC for SO2 in Jiangxi (2005)
(Euro/tSO2)
0
500
1000
1500
2000
0
.25
.5
.75
1
Abated SO2 Emissions (% of 2005 Emissions)

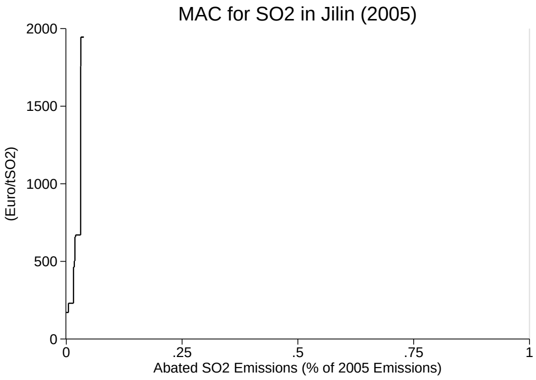
MAC for SO2 in Jilin (2005)
(Euro/tSO2)
Abated SO2 Emissions (% of 2005 Emissions)

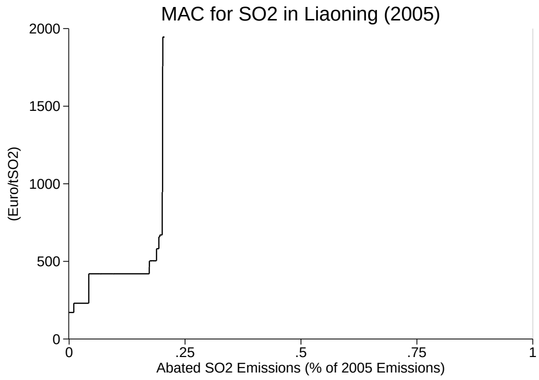
MAC for SO2 in Liaoning (2005)
(Euro/tSO2)
Abated SO2 Emissions (% of 2005 Emissions)

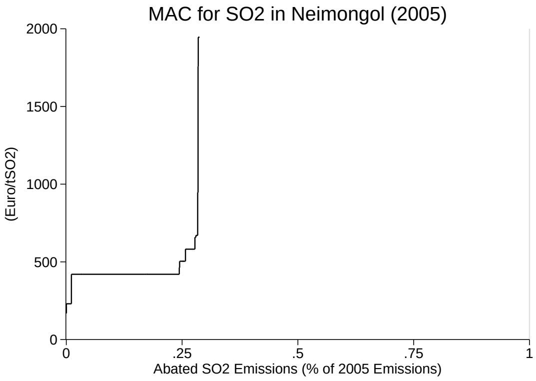
MAC for SO2 in Neimongol (2005)
(Euro/tSO2)
Abated SO2 Emissions (% of 2005 Emissions)

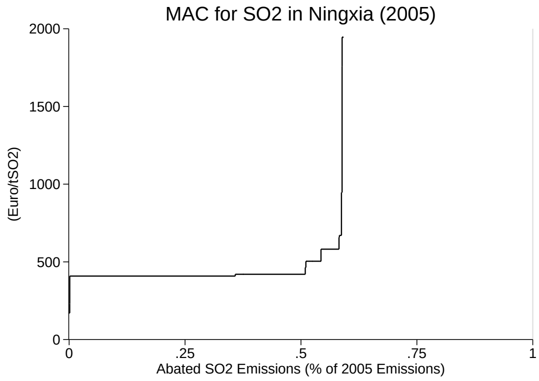
MAC for SO2 in Ningxia (2005)
(Euro/tSO2)
Abated SO2 Emissions (% of 2005 Emissions)

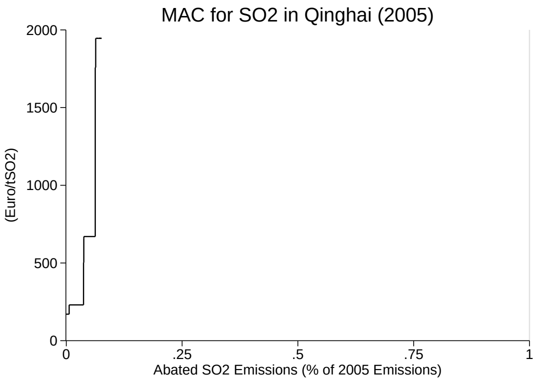
MAC for SO2 in Qinghai (2005)
(Euro/tSO2)
Abated SO2 Emissions (% of 2005 Emissions)

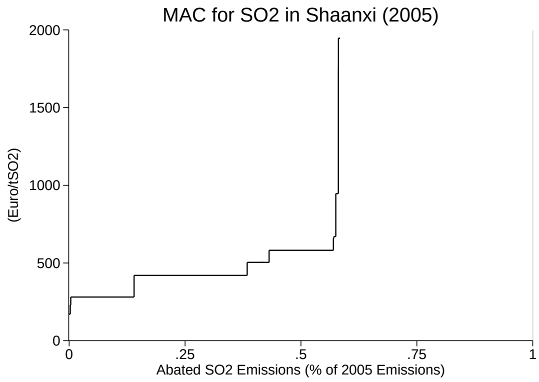
MAC for SO2 in Shaanxi (2005)
(Euro/tSO2)
Abated SO2 Emissions (% of 2005 Emissions)

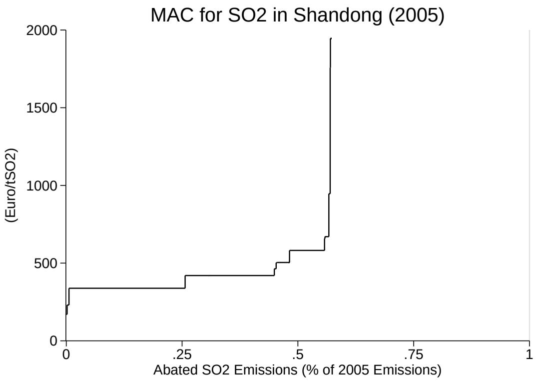
MAC for SO2 in Shandong (2005)
(Euro/tSO2)
Abated SO2 Emissions (% of 2005 Emissions)

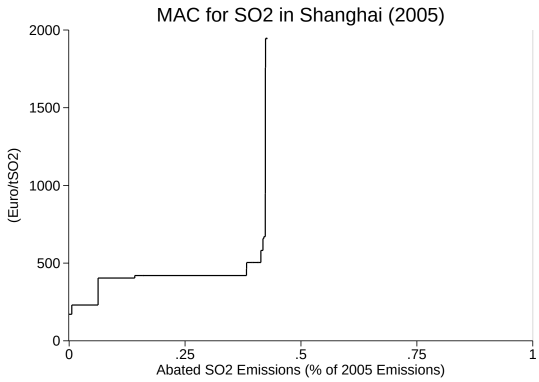
MAC for SO2 in Shanghai (2005)
(Euro/tSO2)
Abated SO2 Emissions (% of 2005 Emissions)

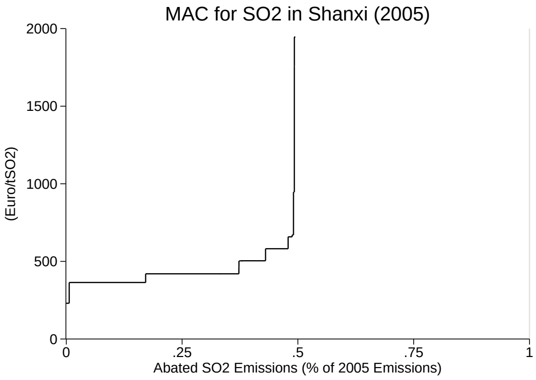
MAC for SO2 in Shanxi (2005)
(Euro/tSO2)
0
500
1000
1500
2000
0
.25
.5
.75
1
Abated SO2 Emissions (% of 2005 Emissions)

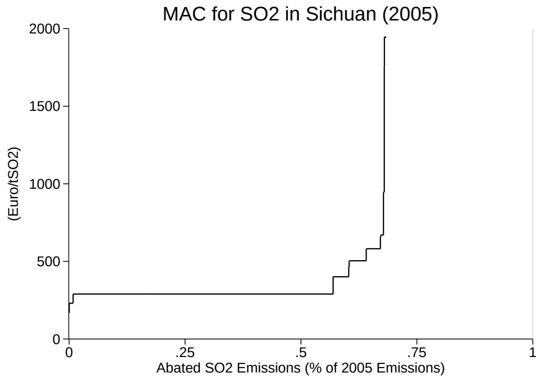
MAC for SO2 in Sichuan (2005)
(Euro/tSO2)
0
500
1000
1500
2000
0
.25
.5
.75
1
Abated SO2 Emissions (% of 2005 Emissions)

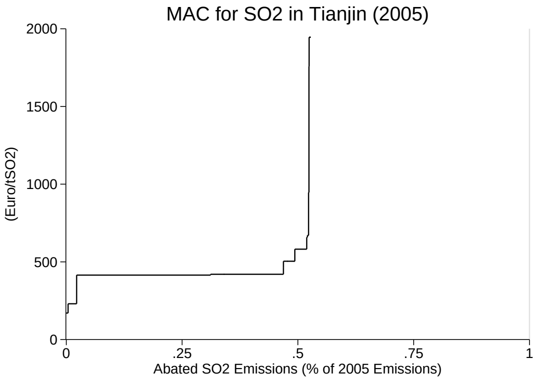
MAC for SO2 in Tianjin (2005)
(Euro/tSO2)
0
500
1000
1500
2000
0
.25
.5
.75
1
Abated SO2 Emissions (% of 2005 Emissions)

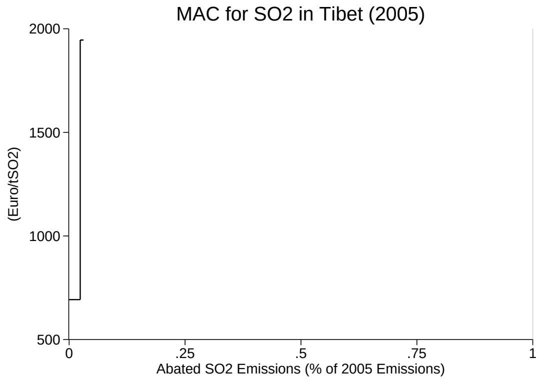
MAC for SO2 in Tibet (2005)
(Euro/tSO2)
500
1000
1500
2000
0
.25
.5
.75
1
Abated SO2 Emissions (% of 2005 Emissions)

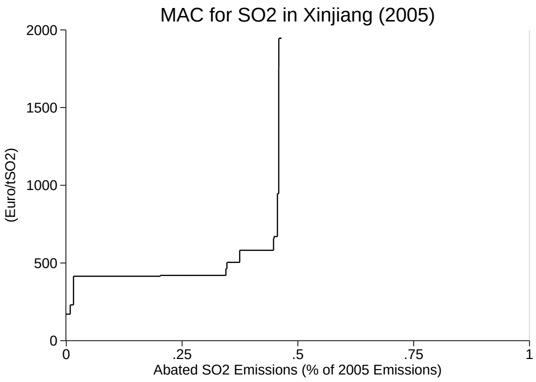
MAC for SO2 in Xinjiang (2005)
(Euro/tSO2)
0
500
1000
1500
2000
0
.25
.5
.75
1
Abated SO2 Emissions (% of 2005 Emissions)

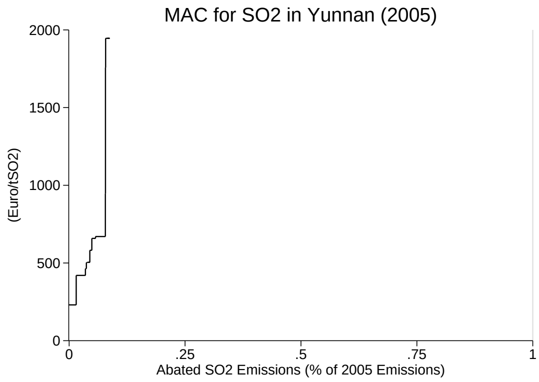
MAC for SO2 in Yunnan (2005)
(Euro/tSO2)
0
500
1000
1500
2000
0
.25
.5
.75
1
Abated SO2 Emissions (% of 2005 Emissions)

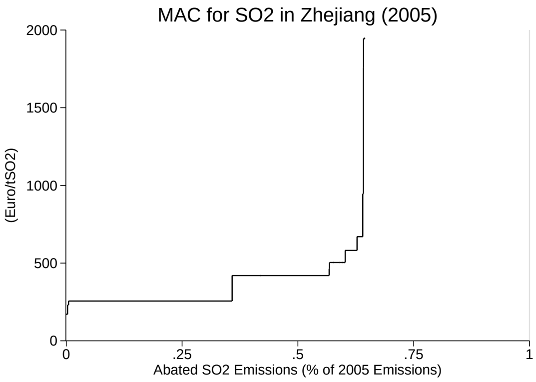
MAC for SO2 in Zhejiang (2005)
(Euro/tSO2)
0
500
1000
1500
2000
0
.25
.5
.75
1
Abated SO2 Emissions (% of 2005 Emissions)