

Zero “Waste”? Competition and Pollution Abatement in the Waste Incineration Industry

Kaicheng Luo *

March 15, 2026

Abstract

This paper studies China’s waste incineration industry where substantial environmental improvements take place absent tighter regulation or stronger enforcement. Using a novel dataset that combines 28 million daily plant-level emissions records with 300 thousand procurement contracts, I show that most of the decline in emissions is driven by competition generated by rapid entry. New plants—often promoted due to inter-jurisdictional competition—create substantial excess capacity. It is also the excess capacity that allows local governments to exert their environmental preferences by reallocating waste toward cleaner producers, effectively rewarding plants with better environmental performance. Facing the threat of demand reallocation, incumbent plants invest in pollution-control technologies and retrofit existing facilities. Back-of-the-envelope calculations suggest that the benefits of resolving environmental externalities outweigh the wasted fixed costs of duplicate investment by an order of magnitude.

Keywords: Waste incineration, Pollution abatement, Competition and regulation, China, Excess capacity

JEL Classification: Q53, L13, Q58, L51

*. Luo: Massachusetts Institute of Technology; kluo0630@mit.edu. I thank Daron Acemoglu, Amy Finkelstein, Jacob Moscona, Benjamin Olken, Tobias Salz, and Shaoda Wang for stimulating comments and suggestions. All errors remain my own.

1 Introduction

A central challenge in environmental economics is understanding how firms internalize environmental preferences when pollution is not directly priced. In most empirical settings, improvements in environmental performance are attributed to regulatory interventions—such as tighter standards, expanded monitoring, or stronger enforcement (Greenstone 2002, Duflo et al. 2013, Greenstone et al. 2022). Yet laws tend to be sticky and enforcement campaigns episodic, and in many industries, environmental outcomes continue to improve even when formal regulations remain largely unchanged. This pattern suggests that forces outside the traditional regulatory framework may also play an important role in shaping firms’ environmental behavior.

Competition is one such force. By reshaping market shares and firms’ incentives to invest in quality, competitive pressure can potentially influence environmental performance even in the absence of explicit pollution pricing. However, the direction of this effect is theoretically ambiguous and empirically difficult to measure. Competition may encourage firms to invest in cleaner technologies to attract environmentally sensitive buyers, but it may also discourage such investments by reducing profit margins. As a result, despite a large literature on environmental regulation, relatively little is known about how competition affects environmental outcomes.

In this paper, I ask three questions: First, how does increased competition affect firms’ incentives to invest in pollution abatement? Second, with regulations not changing and pollution not directly priced, through what channels can environmental preferences be transmitted to firms? Finally, balancing duplicated fixed costs against unpriced environmental externalities, is there currently too much or too little competition?

Answering these questions empirically is challenging. First, measuring competition and environmental performance at the firm level is difficult, as detailed data on emissions, abatement investments, and local market structure are rarely observed jointly. Second, even when such data are available, isolating the causal impact of competition is complicated by the absence of clear exogenous shocks to market structure. As a result, much of the existing literature relies on variations generated by regulatory reforms or trade liberalization that plausibly alter competitive pressure (Greenstone 2002, Shapiro and Walker 2018, Bombardini and Li 2020). However, policy reforms and trade liberalization often simultaneously affect multiple margins beyond competition, such as compliance cost and technology availability, making it difficult to attribute environmental changes to competitive forces alone.

Meanwhile, the mechanisms through which environmental preferences are transmit-

ted to firms are rarely directly observed. In settings where laws change little and pollution is not directly priced in formal markets, it remains unclear how firms internalize environmental preferences. In summary, despite a growing body of empirical work, it remains difficult to directly quantify how competition shapes firms' incentives to invest in pollution abatement, through which channels environmental preferences are transmitted, and whether the current degree of competition is socially excessive or insufficient.

This paper seeks to address these issues in the context of China's waste incineration industry. There are two defining features of this industry: it has expanded rapidly—the number of plants increased eightfold over the past decade, and its pollution intensity has declined just as substantially. Today, an average incineration plant emits 40% less particulate matter, 10% less nitrogen oxides, and 30% less carbon monoxide than its counterpart five years ago. These improvements occur despite unchanged national standards since 2014, and same monitoring infrastructure since 2017. I observe no bunching around regulatory thresholds, extremely rare ($<0.01\%$) instances of breaches, and no indication that firms anticipate further tightening of standards, which are already comparable to the world's strictest rules (e.g., EU regulations).

I show that much of the pollution reduction in China's waste incineration sector is actually induced by competition – a seemingly excessive entry dynamics that shifted market share toward cleaner producers. Using a novel dataset that combines 28 million daily plant-level emissions records from Continuous Emission Monitoring Systems (CEMS) with 300 thousand procurement contracts, I document that competitive pressure from new entrants explains almost all observed incumbent retrofitting efforts, and these responses account for about 20% of the aggregate decline in emissions.

These dynamics are underpinned by two forces that, albeit operating outside of price mechanisms, somewhat mimic the outcomes of an externality market. First, local governments systematically reallocate garbage toward cleaner producers. In effect, they behave like consumers with explicit environmental preferences, rewarding lower-emission plants with greater demand and creating incentives for both entrants' environmental positioning and incumbents' investment in pollution-control technologies. Second, entry is largely determined outside the market and is frequently subsidized or promoted for political reasons. This policy-driven entry erodes incumbents' market power and amplifies competitive pressure along the environmental dimension.

Entry is not always celebrated. The scale and speed of such often triggered recurring concerns about excess capacity, which is considered a salient issue in the policy circle (e.g. 13th Five Year Plan, 14th Five Year Plan). Canonical model predicts an excess number of firms and excessively high quality levels, because new entrants do not internalize the

market stealing effect, nor incumbents' fixed costs during quality upgrades. This concern is typically framed in terms of inefficiency, waste, or misallocation. Yet the sign of welfare calculations can be completely flipped in the presence of unpriced externalities. In those contexts, excess capacity becomes precisely what enables environmental reallocation to take place. When capacity is abundant, dirty incumbents are no longer insulated from competition. Demand can be steered toward cleaner plants, and entry becomes a policy tool that disciplines incumbents. By lowering entry barriers and intensifying competition for waste allocation, capacity expansion creates the conditions under which cleaner producers gain market share and dirtier plants endogenously respond through technology investment. In this sense, what is often criticized as overcapacity functions as a selection mechanism that accelerates environmental upgrading.

My findings suggest a different interpretation of China's green transition and a more nuanced understanding of excess capacity. Rather than being driven primarily by top-down enforcement and incumbent compliance, pollution mitigation in this sector operates to a large extent through a seemingly wasteful tide of entry. Excess capacity, in this setting, is not just a story of waste—it is also a story of a market-oriented way incorporating environmental preference into the production process.

This paper contributes to three strands of literature. First, it empirically studies how competition affects product "quality" in rapidly expanding markets using high-frequency firm-level data. Classic theory predicts that both directions are plausible: absent fixed cost, quality choices will align more with the average consumer's preference for quality (Spence 1975). Note that greater competition can either increase quality (to attract consumers) or reduce it (partly because margins shrink and incentives weaken), and such preference does not have to be welfare-improving from the socially planner's perspective. On the optimistic side, Kessler and McClellan (2000) finds that competition substantially lower medical costs for patients and significantly lower rates of adverse health outcomes in the healthcare industry. Matsa (2011) finds that competition from Walmart induces supermarkets to significantly improve product availability. On the pessimistic side, Becker and Milbourn (2011) finds that competition reduces the quality of credit ratings, while Mazzeo (2003) finds an inverted-U relationship between market concentration and airline service quality. This paper leverages the availability of granular, high-frequency firm-by-investment data and the fact that rapid, politically driven entry expansion generates numerous plausibly exogenous shifts in market structure across the country. These features allow the paper to shed new light on this long-standing question.

A second thread of literature takes fixed costs involved in market entry and quality investments more seriously. The tradition starts from the observation that entrants typ-

ically neglect the market stealing externalities, and therefore the canonical prediction is that entry is often excessive (Spence 1976, Mankiw and Whinston 1986, Bresnahan and Reiss 1991, Berry and Waldfogel 1999). In China, excessive capacity is more often a sign of political economy distortions than a curse of free entry conditions (Hsieh and Klenow 2009, Hsieh and Song 2015). This literature emphasizes that the GDP tournament underpins the preference for *local* investment projects (Li and Zhou 2005), leading to duplicated industrial policies and therefore duplicated industrial capacity across localities (Luo, Wang, and Yang 2025, Fang, Li, and Lu 2025). This paper contributes to this debate by highlighting a potential flip side of the argument: entry may sometimes be socially insufficient when gains from competition-induced environmental improvements outweigh the welfare losses from duplicated investment. From this perspective, slack capacity may be a feature rather than a distortion, and concerns about “overcapacity”, themselves overstated, may deter the government from taking socially responsible actions.

Finally, this paper contributes to a growing literature on the determinants of pollution abatement and environmental performance. Much of this literature emphasizes the role of environmental regulation and enforcement in reducing emissions. For example, studies show that tighter environmental standards, expanded monitoring technologies, and stronger regulatory enforcement can substantially improve firms’ environmental performance (e.g., Duflo et al. (2013); Zou (2021)). A related, but much thinner line of work highlights the role of market forces in driving environmental improvements, such as consumer demand for cleaner products (Barwick et al. 2024), investor pressure (Christensen, Hail, and Leuz 2021), and reputational concerns (Afsah et al. 2013). This paper contributes to this literature by highlighting the possibility that such market can be “engineered”. In the presence of positive externalities, there’s plenty of room for the government to take action without directly altering the regulatory framework from time to time.

The remainder of the paper is organized as follows. Section 2 outlines the institutional background, and reviews the data sources. Section 3 documents the incumbents’ response to new incinerator entries. Section 4 examines its underlying causes and zooms in the regulated allocation of demand. Section 5 evaluates the benefits and costs of excess capacity through some back-of-the-envelope calculations. Section 6 concludes.

2 Institutional background and data

In this section, I describe some key features of the waste incineration industry, as well as the broader landscape of environmental governance in China. I then introduce two novel datasets I constructed, one from the continuous emission monitoring system (CEMS) and

the second from government procurement.

2.1 Waste Incineration in China: overview

There are three prominent methods of municipal solid waste management: incineration, landfilling, and biological decomposition (such as composting or anaerobic digestion). Incineration involves burning waste at high temperatures in specialized facilities to significantly reduce its volume while generating heat or electricity. Compared with landfills—where waste is buried and gradually decomposes over decades—incineration requires far less land and can recover energy, but it raises concerns about air pollutants and the need for advanced emission controls. Landfills are generally cheaper to operate and can handle mixed waste streams, yet they occupy large areas and may release methane, a potent greenhouse gas, as organic material decomposes. Biological decomposition methods, including composting and anaerobic digestion, focus on processing organic waste separately, converting it into compost or biogas; these approaches produce fewer direct emissions but require effective sorting systems, cannot treat most non-organic waste, and are not suited for large-scale waste management.

The waste incineration sector emerged in response to the severe urban waste crisis that accompanied China's rapid urbanization in the 2000s, when many cities faced overflowing landfills and limited land for new disposal sites. Today, the country operates 1,081 out of 2,800 waste incinerators around the globe (Havel 2024), handling more than 90% of trash collected in the country. Incineration offered a relatively fast solution: burning waste significantly reduces its volume while generating electricity through steam turbines.

The industry experienced rapid expansion in the past two decades. In 2005, roughly 85% of municipal waste was landfilled and less than 10% incinerated. Appendix Figure A.1 shows the number of new incinerators constructed each quarter. New construction peaked around 2021, when policymakers first expressed reservations about continuing their previously unwavering support for the industry, but still remained at high levels. Part of the reason behind shift in the policy landscape can be seen in Appendix Figure A.2: the expansion of incineration capacity outpaced the growth the trash volumes by more than five times. In 2024, the maximum designed capacity is already 150% higher than the total volume of collected municipal solid waste.

Concerns about overcapacity have long accompanied the development of waste incineration. Historically, many countries expanded incineration infrastructure in response to rising urban waste and tightening landfill constraints. In the 1980s and 1990s, parts of Eu-

rope and Japan built large fleets of incinerators designed to guarantee waste disposal and generate energy. However, as recycling and waste-reduction policies expanded, some regions began to worry that existing plants exceeded the amount of waste available, creating pressure to keep facilities running at high utilization. Critics argued that overcapacity could discourage recycling or waste reduction because operators need a steady waste stream to recover high fixed costs. Similar debates have re-emerged more recently in rapidly expanding markets such as China, where local governments have invested heavily in waste-to-energy facilities to address landfill shortages and urban environmental concerns.

More broadly, fears of overcapacity are not unique to waste management but recur across many infrastructure-intensive industries. Sectors such as electricity generation, steel, cement, and transportation networks frequently experience investment waves when governments or firms anticipate future demand growth. Because these projects require large upfront capital investments and long planning horizons, capacity expansions often occur simultaneously across regions, sometimes resulting in periods where supply exceeds demand. Economists have noted that such cycles can arise from policy incentives, competition among jurisdictions for investment, or strategic behavior by firms seeking to secure market share. When infrastructure is costly and durable, societies periodically grapple with the risk that ambitious investment programs may overshoot actual demand.

2.2 Entry

Entry into China's waste-incineration industry follows a two-step procedure. Local governments have the full authority to determine whether, when and where a new facility should be built. Municipal authorities then organize an auction where firms compete for the right to construct and operate. These projects are commonly implemented under public-private partnership arrangements, such as build-operate-transfer (BOT) contracts, in which a private developer finances and constructs the plant, operates it for a concession period—often 20 to 30 years—and then transfers ownership to the local government.

The process of entry typically unfolds over several years. The winning firm undertakes feasibility studies, environmental impact assessments, and construction of the facility, which generally takes two to three years before operations begin. The fixed costs are usually substantial. A new plant can cost somewhere between 10 to 100 million dollars. Once operational, plants function as long-lived infrastructure assets. Contracts with local governments are sometimes re-negotiated regarding the specificity of garbage processing

fee, minimal waste supply, and other configurations, but not often.

2.3 Product market

There are two major sources of revenue for the incineration plants. They are paid a flat fee per ton of garbage they process by the local government, and they sell the electricity they generated to the national grid, at a constant subsidized price. Based on listed-firm balance sheets, each of the two components contributes roughly half of their total revenue (e.g. Canvest Environmental Protection Group Co. Ltd., China Tianying Inc). The garbage process fee is usually fixed across firms within a same city but re-negotiated every couple years, and the (subsidized) electricity price is fixed across time and across the whole country until 2024.

There are a few notable features of on the product market (waste-burning) side of this industry. First, there's little horizontal heterogeneity across different plants, so excess capacity is unlikely to arise as an equilibrium outcome of monopolistic competition (Chamberlin 1951). Second, as shown in Appendix Figure A.2, demand is relatively stable and fully visible. The collected volume of municipal solid waste increases by 4 percent annually, and monthly data is publicly available on government portals. This rules out a canonical explanation of excessive capacity where demand fluctuation and cyclic expectations lead firms to build temporary buffer capacity or hold idle assets because they anticipate large short-term ups and downs in sales (Sun 2025). Third, there is effectively no international trade in municipal solid waste involving China. In 2017, the central government implemented a comprehensive ban on the import of foreign waste under the "National Sword" policy, which prohibited the entry of most solid waste streams that had previously been processed domestically (State Council, 2017). As a result, China's waste-incineration industry operates almost entirely as a closed domestic system: plants must rely on locally generated waste as their primary input, and excess capacity cannot be absorbed through international trade. Fourth, demand is highly localized. Garbage can be re-allocated within the city and nearby areas but rarely across cities and never across provinces.

These features combined enables us to define market boundaries and calculate the intensity of competition.

2.4 Regulation and emissions monitoring

National standards, issued by the Ministry of Ecology and Environment (MEE) in 2014, set strict limits on emissions of key pollutants, including particulate matter, sulfur diox-

ide, nitrogen oxides, dioxins, and heavy metals. These standards often approach or exceed U.S. and European benchmarks (Dugoua and Moscona 2025). Compliance is enforced through a combination of mandatory inspections, reporting requirements, and the installation of continuous emissions monitoring systems (CEMS). Starting from 2017, all modern incinerators are required to install CEMS, which provide real-time data on emissions that are automatically transmitted to local environmental authorities. Plants that fail to meet standards are subject to fines, forced operational suspension, or license revocation, creating strong regulatory incentives for investment in flue-gas treatment technologies and operational adjustments. The government starts making those data available to public in January 1, 2020.

Appendix Figure A.3, plots the evolution of emission level for each major pollutant, including PM, NO_x and CO. An average incineration plant in 2025 emits 40% less particulate matter, 10% less nitrogen oxides, and 30% less carbon monoxide than its counterpart merely five years ago. Appendix Figure A.4 plots the distribution of daily emission levels for these pollutants from 2020-2025, alongside the corresponding regulatory benchmarks indicated by dashed gray lines. I observe no bunching around the thresholds and rare occasions where breaches occur. Throughout the sample period, only 100 breaches of regulatory standards. That's less than 0.01% of all observations.

2.5 Data

This study draws on several novel datasets. Our two primary datasets include daily chimney-level pollutant monitoring data covering 1,081 firms and 2,349 plants from 2020 to 2025 (MEE online database) and a complete record of all contracts, both government and non-government, involving these firms from 2015 to 2025 (scraped from procurement databases and local bidding platforms). These datasets allow us to analyze firm-level emissions, compliance, and operational behavior with unprecedented granularity. In Appendix Table A.1, I present some summary statistics, including the designed capacity, power generation capacity, first year of operation, exit status, as well as emission readings for different pollutants.

In addition, I make use of a set of auxiliary datasets that are standard in the literature: geolocated firm-level information on registered capital, construction timing, and designed capacity (1990–2024, firm registration data which is released by the Chinese State Administration for Market Regulation); prefecture-year level garbage collection volumes and composition proxies (2000–2023, city yearbooks); and the universe of 3.3 million central and local policy documents (1980–2025, scraped from the Pkulaw database).

Together, these sources provide a detailed and multi-dimensional view of China's waste-incineration sector, combining market activity, regulatory oversight, and environmental outcomes.

3 Competition and emission reduction

Does competition encourage firms to invest in pollution-control retrofits or instead increase emissions? In this section, I answer this question by examining how incumbents respond to the entry of new incinerators.

3.1 Defining competition

Identifying the causal impact of competitor entry is empirically challenging because entry decisions are typically endogenous, reflecting latent profitability, incumbent capacity utilization, and strategic investments. To address this concern, researchers often model entry as a strategic game and estimate it structurally (Berry 1992; Jia 2008; Ryan 2012), using equilibrium conditions to recover underlying profit functions and fixed costs. In our setting, however, the politically driven expansion of the industry—unprecedented in both scale and speed—occurs largely independently of local market configurations. This institutional feature generates plausibly exogenous variation in competitor entry, allowing us to estimate reduced-form treatment effects.

Even with exogenous entry, defining competitors is still non-trivial, especially when spatial spillovers are present. That entry decisions are independent across markets has been widely criticized as an assumption (Jia 2008), and many studies therefore allow for flexible competitive interactions among firms across locations. In our setting, however, institutional features sharply limit the geographic scope of competition. I define competing pairs as plants located within a 50 km radius and within the same prefecture. This definition reflects the institutional organization of the industry: prefectural governments are responsible for garbage collection and waste transportation, and waste allocation rarely crosses administrative boundaries. As a result, competition among incinerators primarily occurs within prefectures rather than across them.

I start by validating that I am capturing meaningful competitive relationships. I test a canonical prediction of entry in competitive markets: market stealing. The outcome I look at is the percentage of production downtime – days in which firms shut down their furnaces due to a lack of trash – within each six-month period. Regulations require plants to preheat furnaces over 850°C to maintain proper combustion. No trash can be burned

below this temperature and the firms typically have to spend 6–12 hours, consuming several thousand USD worth of fuel, to reheat the furnace. Therefore, shutting down and restarting an industrial incinerator furnace can be non-trivially costly. Although I do not observe actual daily garbage supply, I believe that the productive downtime provides a meaningful proxy for firm’s profit loss.

The main regression model at the level of plant i and semi-year t is

$$Y_{it} = \sum_{j=-L}^K \beta_j D_{i,t-j} + \alpha_i + \chi_t + \gamma' X_{it} + \varepsilon_{it} \quad (1)$$

where Y_{it} is the percentage of downtime of plant i within the six-month period. $D_{i,t-j}$ is an indicator for whether a nearby plant within 50 km radius is constructed at time period t . (α, χ) are fixed effects.

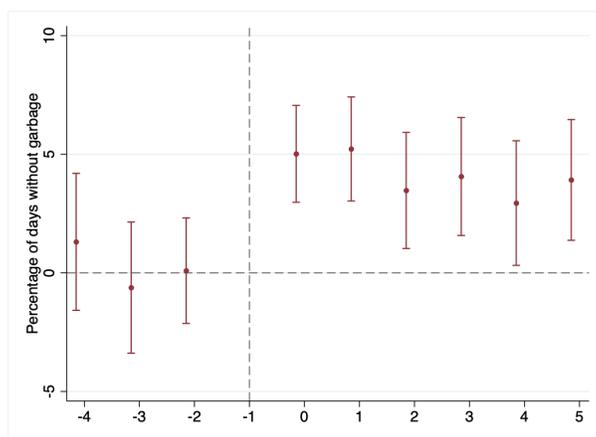


Figure 1: Market stealing effect. This figure reports event study estimates (β_j) following the specification of equation (1). The outcome variable is the total number of days within the 6-month period where the furnace is not operating. Treatment is defined as the construction of a nearby incinerator within 50 km radius in the same city. I focus on the sample where excess capacity exists. Plant fixed effects and time fixed effects are parsed out. Standard errors are clustered at plant level.

Figure 1 presents the results. I find that the entry of a nearby competitor immediately decreased the capacity utilization rate by 5 percentage points. Such effect is persistent at least for the next three years. Crucially, I find no pre-entry difference in trend between incumbents who later welcome a competitor and those who do not, indicating that firm entry is not selected according to capacity shortage.

The 50 km radius is chosen for illustrative purposes only. In Appendix Figure A.5, I estimate the market stealing effect using alternative radii (25 km, 75 km and 100 km). Reassuringly, I see the size of market stealing effects fading over distance, but all positive and qualitatively exhibiting the shape of a step-wise function.

I present two placebo exercises to gain confidence on our results. First, notice that market stealing should not happen when the market is not yet saturated. Since both the total volume of municipal solid waste collected upstream and the designed capacity of each plant are public information, I can identify those cities that are *under*-capacitated – cities where I expect market expansion to take place instead of market stealing. Conceptually, if I run the same regression on the subsample without excess capacity, I should see null effects on market stealing. Appendix Figure A.6, Panel A confirms this conjecture. Second, I explore placebo events using (i) entry of an additional competitor that is 300 km away; (ii) entry of an additional competitor that is within 50 km radius but across provincial boundaries. The left two panels of Appendix Figure A.7 shows that market stealing effect is completely muted around the entry of the placebo competitor, in both cases. Panel B additionally confirms our qualitative knowledge that markets are localized, and firms seldom contract with governments across the provincial border even if their distance are close.

3.2 Competitor entry and emission abatement

I now return to our motivating question, and ask how incumbents respond to the entry of new competitors. The first step towards answering this question is to be able to measure how polluting a plant is at a given time. From CEMS database, which offers daily average readings at plant \times pollutant level, I construct a proxy for "dirtiness" by taking the average of the standardized value of emission readings across three major pollutants (PM, NO_x and CO) and across time.

Our baseline estimation strategy follows the same empirical specification as (1), where the outcome of interest is now the level of emission concentration.

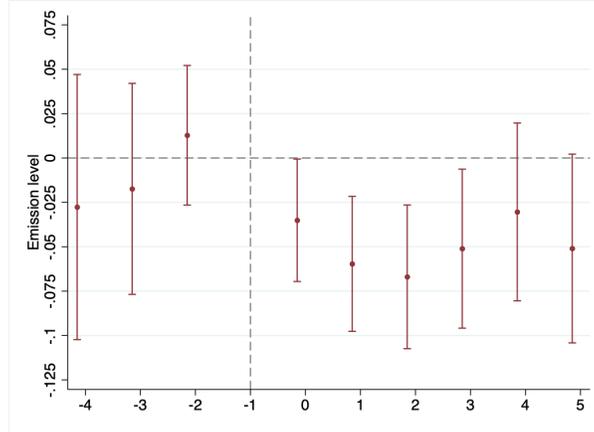


Figure 2: Emission abatement in response to entry. This figure reports event study estimates (β_j) following the specification of equation (1). The outcome variable is the average emission level of all pollutants within a six-month period. Treatment is defined as the construction of a nearby incinerator within 50 km radius in the same city. I focus on the sample where excess capacity exists. Plant fixed effects and time fixed effects are parsed out. Standard errors are clustered at plant level.

Figure 2 presents the results. I find that the entry of a nearby competitor induces incumbents to reduce their emission level by an average of 0.06, which is around 10% of the standard deviation. Throughout our sample period, I observe 1,242 pairs of entry events. A back-of-the-envelope calculation suggests that those entry alone explains more than 25% of all incumbent pollution abatement efforts throughout the sample period, which I detailed in Section 5.

Interestingly, I observe no pre-treatment emission reduction. This pattern suggests that the observed abatement is unlikely to be driven by unobserved local trends—such as tightening regulatory enforcement following political turnover—but instead reflects incumbents’ responses to entry. I also see that such reduction increases over time, reflecting the possibility that pollution-control investments take time to implement, which could mechanically delay observed emission reductions. To further probe the competitive mechanism, I conduct two placebo exercises. In markets without excess capacity, where neighboring entry does not lead to market stealing, entry is better interpreted as market expansion rather than intensified competition. Given our conjecture that retrofitting investments are primarily driven by competitive pressure, I should therefore observe little response from incumbents in these markets. Appendix Figure A.6, Panel B confirms this prediction. Second, I exploit the two placebo events of "non-competitor" entry. The right two panels of Appendix Figure A.7 show that no abatement effort will be exerted around the entry of the placebo competitor.

3.3 Emission abatement and large-scale technology adoption

Although CEMS data provide us with granular readings on final pollution outcomes. There are at least two drawbacks that are worth noting. First, combining evidences from Figure 1 and Figure 2, one might be worried that the drop in emission level comes from firms simply smoothing out the garbage dose across the day when they are endowed with less garbage supply. Second, I do not know how costly it is to reduce the pollution emission. Can firms easily achieve the goal by installing cheap end-of-pipe filters, or do they have to invest heavily in retrofitting the furnace? The former possibility contradicts with our preferred interpretation that firms are actively investing to choose the emission level.

To answer these questions, I turn to procurement contract data. According to Government Procurement Law, (2014 Rev.), all procurement contracts involving public facilities have to be disclosed online via one of the national or provincial platforms. These disclosure requirements generate detailed administrative records on procurement events, including project location, plant capacity, contract value, and the identity of the winning operator. I scrape the universe of those contracts, and compile a dataset of incineration procurement contracts from these public platforms and link them to facility-level information on plant characteristics and operation.

Our first step is to identify, from the universe of contracts, contracts that explicitly mention large-scale technology upgrades. I accomplish this in an AI-assisted way.¹ Appendix Table A.2 shows ten examples of large-scale technology upgrades. I am able to link them to the exact plant and the exact pollutant it targets, know the exact timing, and for some of them know the monetary value and configuration of technology.

I start with a validation exercise showcasing that those technology upgrades, once implemented, lead to significant pollution abatement. Specifically, I estimate the following model:

$$Y_{ipt} = \sum_{j=-L}^K \beta_j D_{ip,t-j} + \alpha_{ip} + \chi_{pt} + \zeta_{it} + \gamma' X_{ipt} + \varepsilon_{it} \quad (2)$$

where Y_{ipt} is the average emission concentration of pollutant p at plant i within the six-month period t . $D_{ip,t-j}$ is an indicator for whether a technology upgrade has been adopted by i targeting pollutant p at time period t . (α, χ, ζ) denote fixed effects that ab-

1. Specifically, I use the following prompt: You are an expert in environmental technologies. I'll pass you some contracts involving trash incinerating firms, one at a time. Some contracts belong to the large-scale technology upgrade category, some not. There are 3 main emittants that the government monitors: PM, NOx, CO. I want you to only focus on upgrades that will matter, not generic contracts relevant to those upgrades, like post-upgrade examination. I want you to classify each into one of the categories. Please answer by one word: PM, NOx, CO, all, none. Here's the contract: {text}

sorb confounders invariant to plant–pollutant, pollutant–time, and plant–time pairs. For example, plant–pollutant fixed effects capture persistent differences in baseline emissions due to plant-specific technology or furnace design. Pollutant–time fixed effects absorb shocks affecting all plants for a given pollutant in a particular period, such as changes in monitoring technology or measurement protocols. Plant–time fixed effects control for plant-level operational conditions that vary over time but affect all pollutants simultaneously, such as fluctuations in waste composition, operating intensity, or maintenance cycles.

Appendix Figure A.8, Panel A, shows the results. Much as I expected, technology adoption strongly predicts subsequent emission abatement progress. To unpack the black box and to make magnitudes easier to interpret, I estimate the dynamic treatment effect pollutant by pollutant, with plant and time fixed effects. Results are shown in Panel B. An average technology upgrade targeting PM caused average emission level to decrease by 5 mg/m³, SO₂ level to decrease by 3.5 mg/m³, and NO_x level to decrease by 15 mg/m³. Effects are persistent.

How much of pollution abatement can be explained by those large-scale technology upgrades? I conduct a back-of-the-envelope calculation. (quantify)

Linking those environmental investments to firms, I ask if neighboring firm entry predicts technology upgrades. Empirically, I use the exact same strategy as outlined in Equation 1.

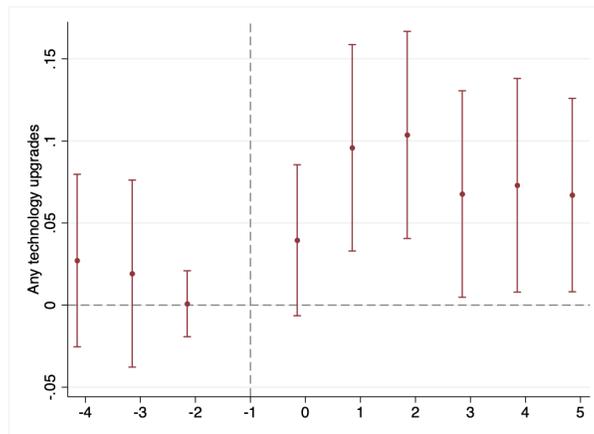


Figure 3: Emission abatement in response to entry. This figure reports event study estimates (β_j) following the specification of equation (1). The outcome variable is an indicator of whether firms adopt a large-scale technology upgrade. Treatment is defined as the construction of a nearby incinerator within 50 km radius in the same city. I focus on the sample where excess capacity exists. Plant fixed effects and time fixed effects are parsed out. Standard errors are clustered at plant level.

Figure 3 shows the results. A competitor entry increased the probability of large-scale

technology upgrades by 10 percentage points. (compare it with baseline rates). This effect is only present in cities where excess capacity exists and hence competition imminent, and not elsewhere (Appendix Figure A.6, Panel C).

4 Why competition works

We've established the fact that incumbents react to intensified competition by investing in pollution-reducing technologies. This section explores why. In Section 4.1 I argue that the retrofitting efforts do reflect the responsiveness to the government's environmental preference, not alternative concerns such as efficiency. Section 4.2 explores how such preference is exerted in the market via an demand-allocation mechanism.

4.1 Competing directly on environment

Competition always sounds like an umbrella term. What are the firms competing on? Is it price, fuel efficiency, or quality? What are they competing for? Is it a larger market share, a more stable relationship with the government, or a higher present-discounted profit flow?

I start by answering the first question. In particular, I ask how much of the incumbents' response is driven by environmental concerns. Our first empirical strategy is to add a triple-difference term to the event study regressions. The term captures how polluting an incumbent is as compared to his competitor-to-be. More specifically, I define z_{it} to be the difference between the incumbent's pre-period pollution level and the entrant's pollution level in the first six months, and estimate the following regression:

$$Y_{it} = \sum_{j=-L}^K (\beta_j D_{i,t-j} + \beta_{j,z} (D_{i,t-j} \times z_{it})) + \alpha_i + \chi_t + \gamma' X_{it} + \varepsilon_{it} \quad (3)$$

Using this structure. I can examine two outcomes of interest. First, does the scale of the market stealing effect vary with respect to the relative environmental positioning? Second, do dirtier firms differentially invest in retrofitting efforts? In a world where the consumer (i.e. the government) prefers plants with lower pollutant emission, I would expect that z_{it} predicts the magnitude of both the market stealing effect and the pollution abatement efforts.

I find both conjectures to be true in Figure 4 and Figure 5. Reassuringly, I find no pre-treatment predictive power of the relative environmental positioning of incumbents on either market stealing effects or pollution abatement efforts, but once a competitor

enters the arena, the relatively dirtier incumbent both gets more market share stolen, and is more likely to invest in retrofitting efforts.

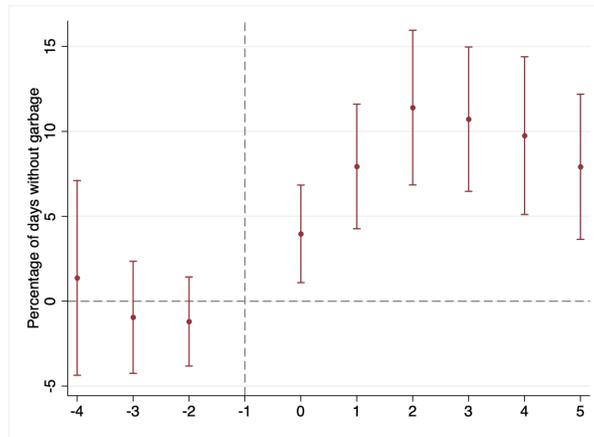


Figure 4: Heterogeneous market stealing effect by difference in emission level. This figure reports the estimates ($\beta_{j,z}$) following a triple-difference specification of equation (3). The outcome variable is the total number of days within the 6-month period where the furnace is not operating. Treatment is defined as the construction of a nearby incinerator within 50 km radius in the same city. I interact the treatment with the difference in emission levels between the incumbent and the new entrant in the first six months. I focus on the sample where excess capacity exists. Plant fixed effects and time fixed effects are parsed out. Standard errors are clustered at plant level.

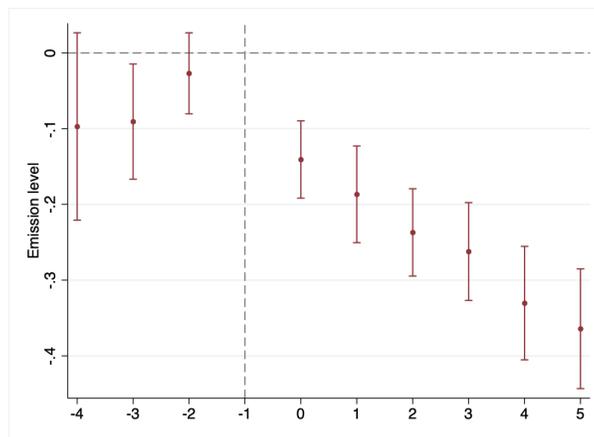


Figure 5: Heterogeneous emission abatement by difference in emission level. This figure reports the estimates ($\beta_{j,z}$) following a triple-difference specification of equation (3). The outcome variable is the average emission level of all pollutants within a six-month period. Treatment is defined as the construction of a nearby incinerator within 50 km radius in the same city. I interact the treatment with the difference in emission levels between the incumbent and the new entrant in the first six months. I focus on the sample where excess capacity exists. Plant fixed effects and time fixed effects are parsed out. Standard errors are clustered at plant level.

One might be concerned that environmental configurations are correlated with other features, such as age of furnace, fuel efficiency, and unit power of electricity generation.

If that is true, then Figures 4 and 5 merely captures co-movement with unobserved confounder, which contradicts with our preferred interpretation. I show two pieces of evidence that such concern is unlikely.

First, in Appendix Figure A.9, I run the same regression on placebo mediators that proxies for efficiency ($\tilde{z}_1 = \Delta$ designed capacity, $\tilde{z}_2 = \Delta$ power generation capacity per hour). I find no evidence that pollution abatement is differentially responsive to those features. Second, in Appendix Figure A.10, I again utilize the entire procurement dataset but now focus on investments that are *not* environment related. I find no evidence that incumbents are increasing other types of investments in response to firm entry.

4.2 Demand allocation

Section 4.1 argues that environmental pressure is the key element in firm-level competition, but the key question – how it is accomplished – is left ajar. In this section, I discuss the institutional backgrounds in greater detail, and then examine one particular mechanism through which environmental preferences can be exerted onto firms. Some alternative stories, which I find implausible, are then qualitatively discussed.

In 2024, a company in Guangzhou called Grandtop Yongxing Group Co Ltd., which operates the cleanest incinerator in the city, signed a deal with the prefectural government, where the city promises to dig up already-buried trash in the landfills and source them to one of their incinerating furnaces. It turns out that those explicit or implicit deals are prevalent. Since the decision to transport and allocate waste is commanded solely in the hands of the local administration, the government can strategically allocate demand to cleaner firms once capacity constraint is slack. Although this does not happen via any explicit price mechanisms, the consequence of demand allocation is isomorphic to the scenario where a set of environmentally-leaning consumers are purchasing more good from cleaner firms.

I show this by estimating the following model:

$$Y_{it} = \beta \cdot \text{Emission}_{it} + \chi_{j(i) \times t} + \gamma' X_{it} + \varepsilon_{it}$$

where Y_{it} is the unutilized capacity of burner i during time period t , and Emission_{it} is the average pollutant emission level over the past 6 months. $\chi_{j(i) \times t}$ is the city \times semi-year fixed effects. This creates scope for reallocation with environmental preference without price competition.

Table 1: Consumer (government) preference for environment and garbage allocation.

	Production downtime			
	(1)	(2)	(3)	(4)
Panel A: Full sample				
Avg pollutant concentration	0.0271** (0.0113)	0.0264** (0.0112)	0.0388*** (0.0113)	0.0299** (0.0141)
Panel B: Placebo: sample with no excess capacity				
Avg pollutant concentration	0.00651 (0.00481)	0.00683 (0.00483)	0.00352 (0.00812)	-0.00344 (0.00578)
Mean of DV	0.167	0.167	0.167	0.167
Time FE	No	Yes	Yes	Yes
Prefecture FE	No	No	Yes	Yes
Firm FE	No	No	No	Yes

Notes: The outcome variable is the proportion of days when no waste is being burned, as documented by CEMS monitors. The independent variable is the standardized average pollutant concentration of three major pollutants of the specific plant over the past six months. Semi-year fixed effects, prefecture fixed effects and firm fixed effects are sequentially parsed out. In Panel A, I estimate the regression on the full sample. In Panel B, I focus on cities where excess capacity does not exist, comparing the total designed capacity of all incineration plants combined with the total municipal solid waste collected, as reported in city year books. Standard errors clustered at prefecture level are reported below the estimates.

Table 1 presents the estimates. I find that one unit increase of pollutant concentration increased production downtime by 2.71 percentage points, which is an 18.0% increase from the mean. Since garbage allocation is usually decided by the local government, I control for prefectural-by-year fixed effects, effectively focusing on within-local-market variation over different plants. In Panel B, I run the same regression focusing on the subsample where excess capacity does not exist. Aligned with our interpretation, I find the effect to be entirely muted in those cases.

Qualitatively, I discuss two alternative interpretations of Table 1. First, to the extent that environmental configurations are typically correlated with marginal costs, cleaner firms may be able to compete on price (i.e., the garbage processing fee) to attract customers. In practice, however, this mechanism is unlikely to operate because institutional constraints require governments to set a uniform price for garbage processing across all plants within the same administrative jurisdiction. Second, if governments and firms can contract *ex ante* on both quantity and environmental configurations, such contract design

could itself induce firms to adopt cleaner technologies. Leveraging the fact that all contracts are observed in the data, I manually reviewed these agreements. I find that most contracts do not specify any quantity commitment, and among the few that do, the minimum disposal amounts are set at extremely low levels. There is therefore little reason to believe that these quantities are ever binding.

5 Quantification

5.1 Cost and benefit of excess capacity

Leveraging the estimates, I conduct a simple back of the envelop calculation.

On the benefit side, observe that we can decompose total emission into three objects that I can measure:

$$\underbrace{y_{it}}_{\text{total emission}} = \gamma \times \underbrace{\bar{y}_{it}}_{\text{Avg. CEMS emission readings}} \times \underbrace{C_i}_{\text{designed capacity}} \times \underbrace{u_{it}}_{\text{capacity utilization rate}} \quad (4)$$

where γ is a conversion factor, \bar{y}_{it} is measured in mg/m^3 , C_i in tons of trash per day, and u_{it} as the percentage of days the furnace is underutilized.

It is important to note that y_{it} is observed for only a subset of firms, using data reported in their Execution Report of Pollutant Discharge Permit. To facilitate downstream calculations linking emissions to health outcomes, I express y_{it} in metric tons. The key parameter to estimate is γ , which converts the product of \bar{y}_{it} , C_i , and u_{it} into total emissions. I calibrate $\gamma \simeq 0.20$ by estimating Equation (4) via OLS in logarithmic form.

Taking reduced form estimates (0.07) from Figure 2, and multiply it by the average number of incumbents it affected (1.76) as well as the standard deviation of each emitant, I find that one additional entry induced total emission reduction of PM by 99.0 tons, NOx by 1959.4 tons and reduction of SO₂ by 672.13 tons.

We then draw on off-the-shelf estimates from the environmental economics literature that translate emission volumes into the economic costs of health damages. Within each pollutant category, I adopt conservative estimates: \$3,800 per ton of NOx, \$14,000 per ton of SO₂, and \$88,000 per ton of PM. Even using these lower-bound values, I estimate that the environmental benefits generated by one additional entrant amount to approximately \$25.6 million per year (Heo, Adams, and Gao 2016).

From a social planner's perspective, these benefits must be weighed against the costs associated with additional entry, including duplicated fixed costs from excess capacity and the cost of technology upgrades themselves. An average incineration plant in China

costs roughly \$50 million to build, while the average technology upgrade in our procurement data costs about \$1 million. Even under the extreme assumption that the entrant generates no additional output and merely engages in business stealing, the environmental benefits alone would offset these costs in less than two years. Given that the typical plant operates for 20–30 years, the aggregate social benefits of an additional entrant exceed its costs by roughly an order of magnitude.

6 Conclusion

This paper studies how competition shapes firms' incentives to invest in pollution abatement in the absence of explicit pollution pricing or changing environmental regulations. Using high-frequency emissions data from China's waste incineration sector, I show that competitive pressure induced by rapid entry plays a central role in driving environmental improvements. Entry shifts demand toward cleaner producers and induces incumbents to retrofit their plants, accounting for a substantial share of the observed decline in emissions. These dynamics are facilitated by the institutional structure of the industry, in which local governments allocate waste across plants and systematically favor cleaner facilities. In effect, environmental preferences are transmitted to firms through administrative allocation of demand rather than through formal price mechanisms.

How special is the waste incineration sector? On the one hand, several features make it particularly well suited for this type of mechanism. Waste disposal is a local public service, and the government typically acts as the sole buyer of processing capacity. This centralized allocation of demand makes it possible for environmental preferences to influence firms' market shares even when pollution is not priced. In addition, emissions are continuously monitored through standardized technologies, allowing regulators to observe environmental performance with relatively high precision. These institutional characteristics create unusually favorable conditions for environmental competition.

On the other hand, the underlying mechanism highlighted in this paper may extend beyond waste incineration. In many sectors where governments play an important role as purchasers or allocators of demand—such as electricity generation, public transportation, water treatment, and other environmental infrastructure—environmental performance is increasingly observable through monitoring technologies. In these contexts, regulators may be able to influence firms' incentives by reallocating demand toward cleaner producers rather than by continuously tightening formal regulations. More broadly, the results suggest that competition can serve as an indirect policy instrument for environmental improvement when governments can steer demand based on observable environmental

performance.

The findings also speak to a broader debate about overcapacity in China. Rapid capacity expansion across a range of industries has often been interpreted as evidence of inefficiency, driven by political incentives for local investment and resulting in duplicated fixed costs. While these concerns are valid in many settings, the analysis in this paper highlights a potential flip side. In the presence of unpriced environmental externalities, excess capacity can facilitate the reallocation of production toward cleaner firms. A policy that targets zero waste may in fact be extremely "wasteful".

Online Appendix

Appendix A Additional figures and tables

Table A.1: Summary statistics

	# obs.	Mean	Std.	Min	Max
Full sample					
Electric power (kilowatt)	1081	21849.76	24425.24	0.00	300000.00
Designed capacity (ton per day)	2349	511.36	175.38	34.00	1600.00
First year of operation	2349	2018.47	4.38	2000.00	2025.00
Exit status	2349	0.05	0.21	0.00	1.00
Concentration (PM)	5749472	3.14	2.57	0.00	60.49
Concentration (NOx)	5749265	131.50	50.83	0.00	349.24
Concentration (CO)	5742537	7.49	9.64	0.00	336.30

Table A.2: Ten examples of major technology upgrades

Firm	City	Year	Winning Bid (RMB)	Description	
Everbright Environmental (Siyang) Co., Ltd.	Environmental Energy	Suqian	2023	246,000	Technical retrofit installing valveless filters in the circulating water system.
Guizhou Environmental Technology Ltd.	Xinyuan Co.,	Liupanshui	2021	3,550,000	Retrofit adding soot blowers to the second and third flue passes of Boilers #1 and #2 (third tender).
Everbright Environmental (Hangzhou Fuyang) Co., Ltd.	Environmental Energy	Hangzhou	2022	66,000	Procurement of flue gas recirculation flowmeter for the Fuyang waste incineration power generation project.
Everbright Environmental (Hangzhou) Ltd.	Environmental Energy Co.,	Hangzhou	2020	140,000	Retrofit service for the lime storage tank top dust collector at the leachate treatment station.
Chengfa Environmental (Runan) Co., Ltd.	Environmental Energy	Zhengzhou	2020	439,000	Procurement of activated carbon for the municipal solid waste incineration power generation project.
Ningbo Shimao Energy Co., Ltd.	Environmental Energy	Ningbo	2023	2,600,000	Procurement of demister and packing for wet scrubber tower in the flue gas wet desulfurization system.
Everbright Environmental (Weifang) Co., Ltd.	Environmental Energy	Weifang	2021	162,000	Procurement of activated carbon.
Everbright Environmental (Suqian) Co., Ltd.	Environmental Energy	Suqian	2021	218,000	Wastewater system retrofit in the lime slurry preparation room and fly ash chelation room.
BEWG Environmental Renewable Energy (Zhangjiagang) Co., Ltd.	Environmental Renewable Energy	Suzhou	2023	2,864,100	Supply and installation of PNCR furnace denitrification system.
Wuxi Xidong Environmental Co., Ltd.	Environmental Energy	Wuxi	2025	986,100	Major overhaul services for baghouse dust collectors No. 1–No. 4.

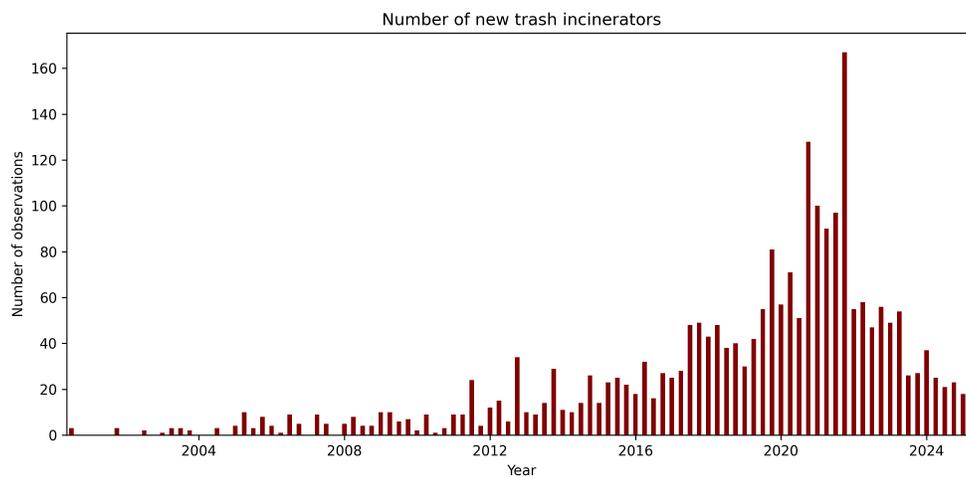


Figure A.1: Number of new incinerators by quarter. Incinerators are registered with the Ministry of Ecology and Environment once construction is completed and CEMS is installed. In our data, that's on average 3 months before they start to burn trash.

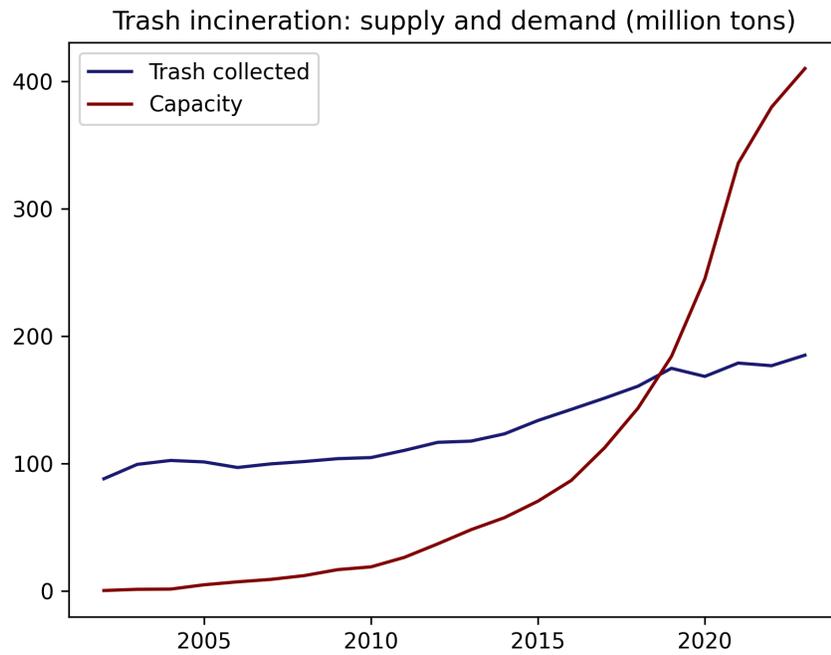


Figure A.2: Landscape of excess capacity of waste incinerators. This figure plots the total volume of municipal solid waste across all cities in China against the aggregate designed capacity of all trash incineration plants combined.

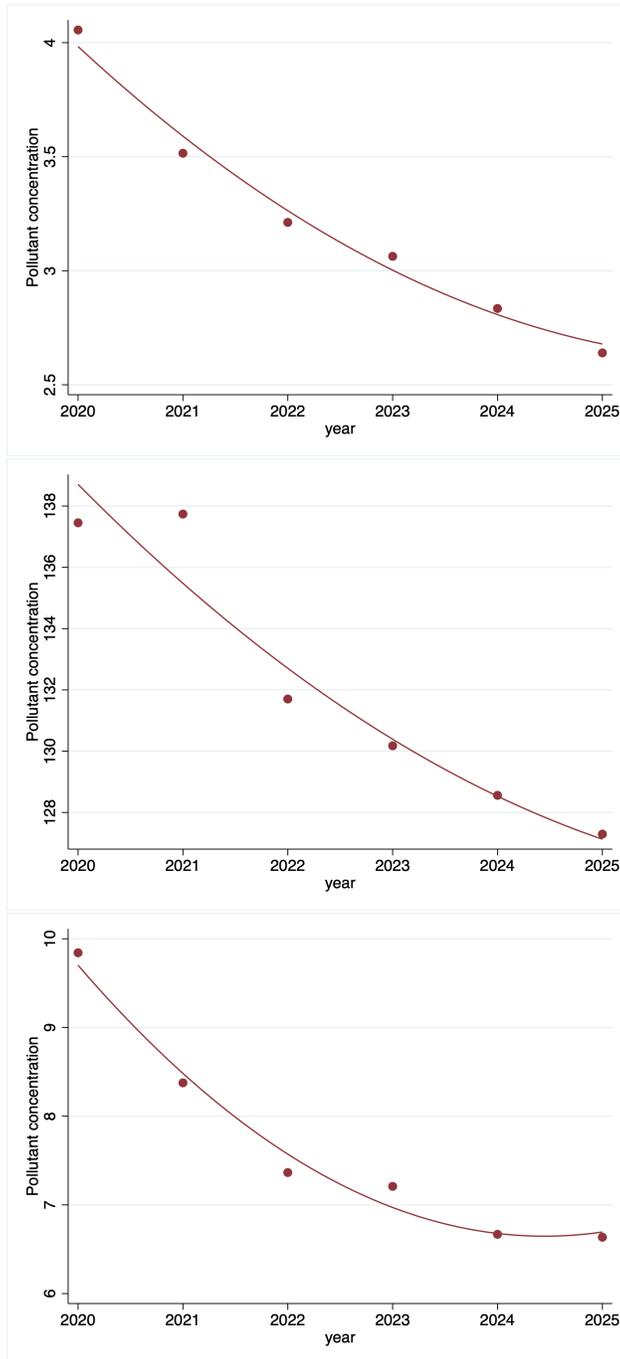


Figure A.3: Emission level of incinerators by year. This figure plots the average emission level for particulate matters (PM), nitrogen oxides (NOx) and carbon monoxide (CO), across all incineration plants.

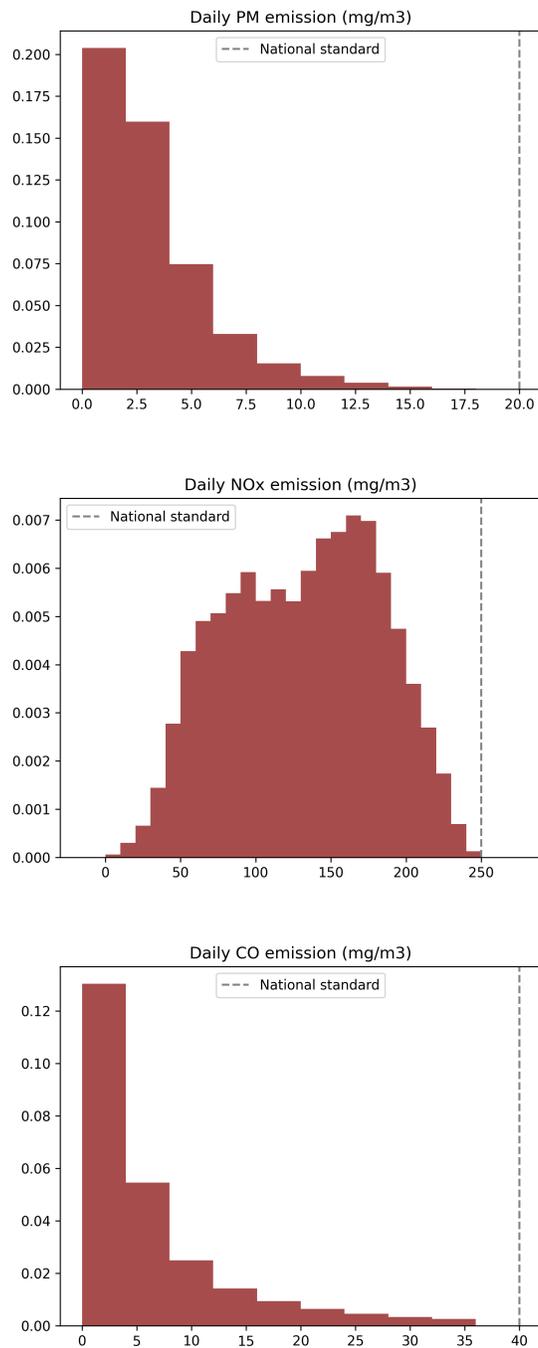


Figure A.4: Distribution of pollution emission levels. This figure plots the distribution of emission levels for particulate matters (PM), nitrogen oxides (NOx) and carbon monoxide (CO), across all incineration plants across the entire sample period (2020-2025).

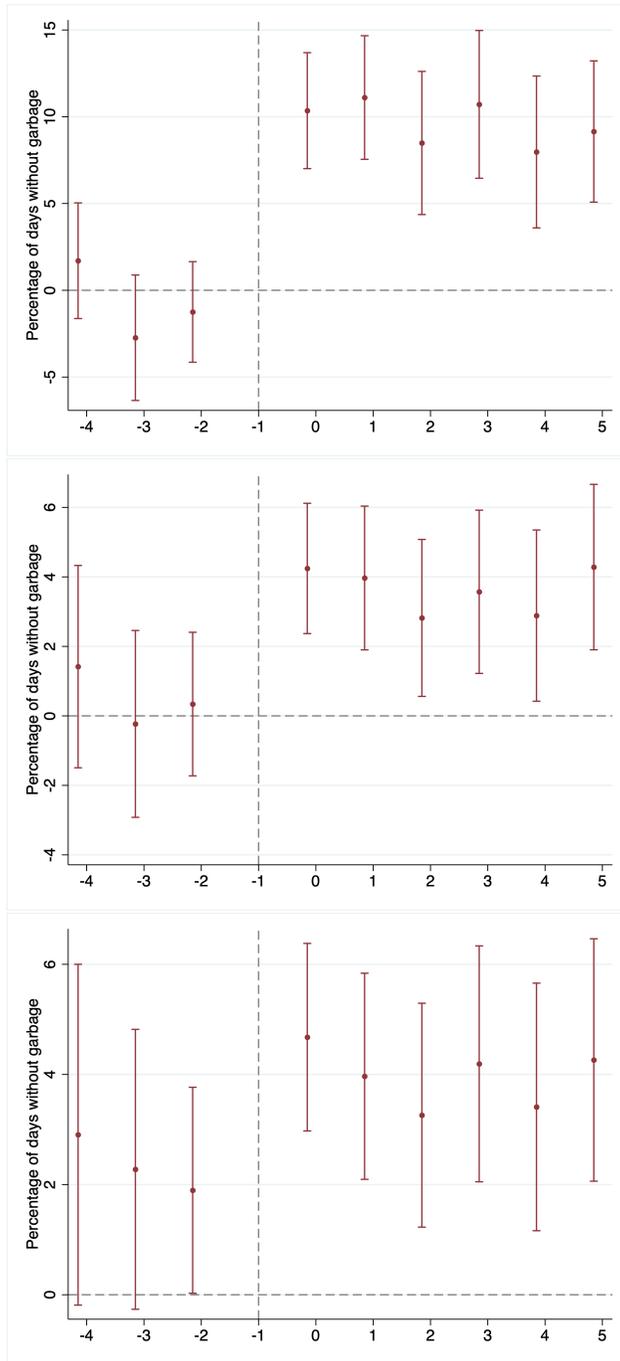


Figure A.5: Re-defining competition. This set of figures replicates our baseline Figure 1 with alternative choices of bandwidth defining competition. In Panel A, I set bandwidth to be 25 km. In Panel B, I set bandwidth to be 75 km. In Panel C 100 km. Plant fixed effects and time fixed effects are parsed out. Standard errors are clustered at plant level.

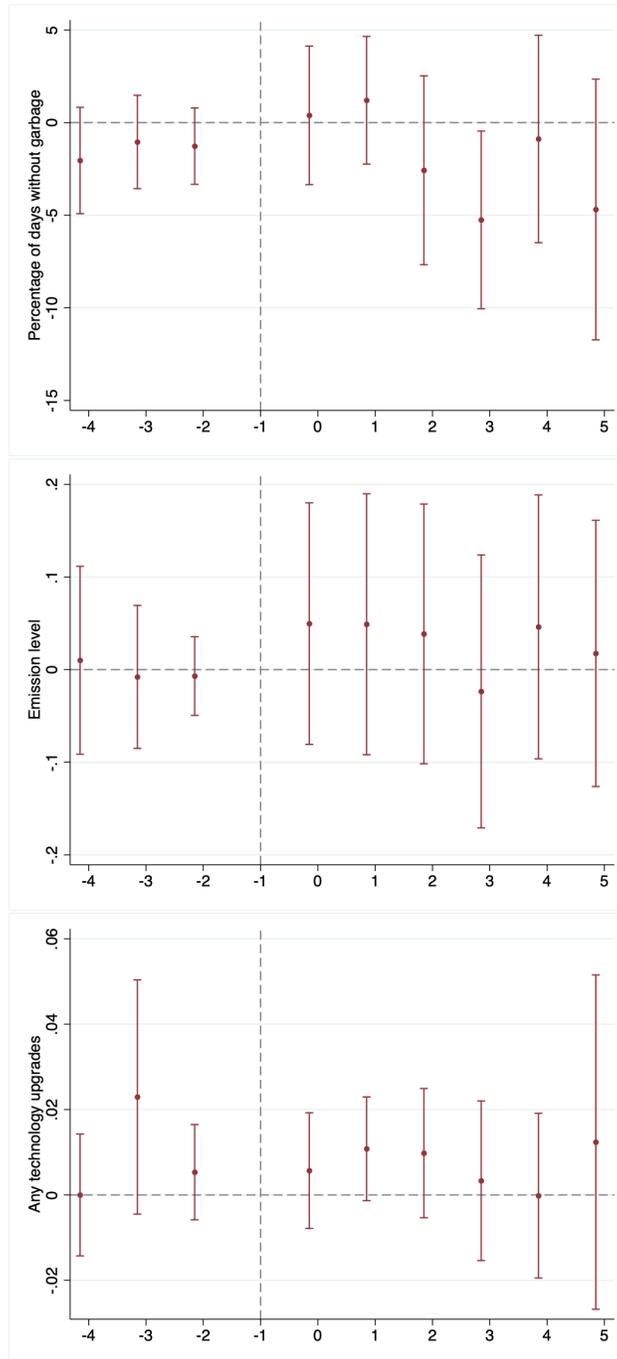
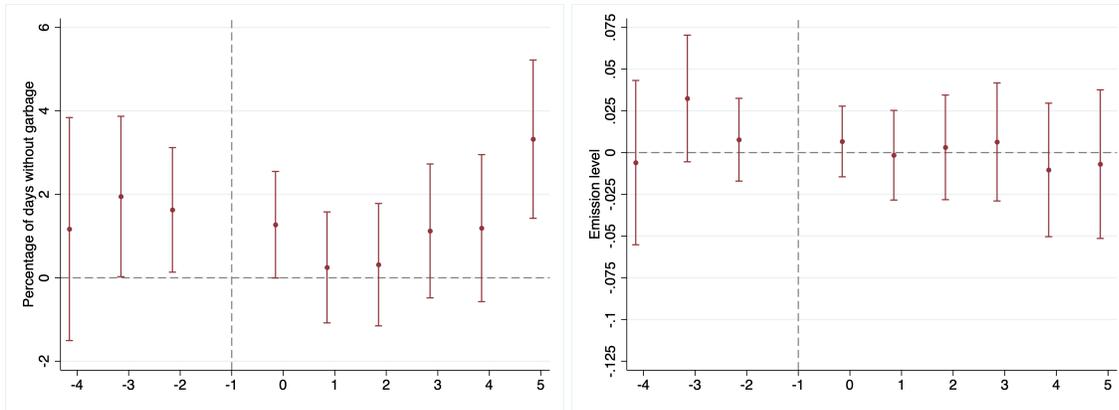
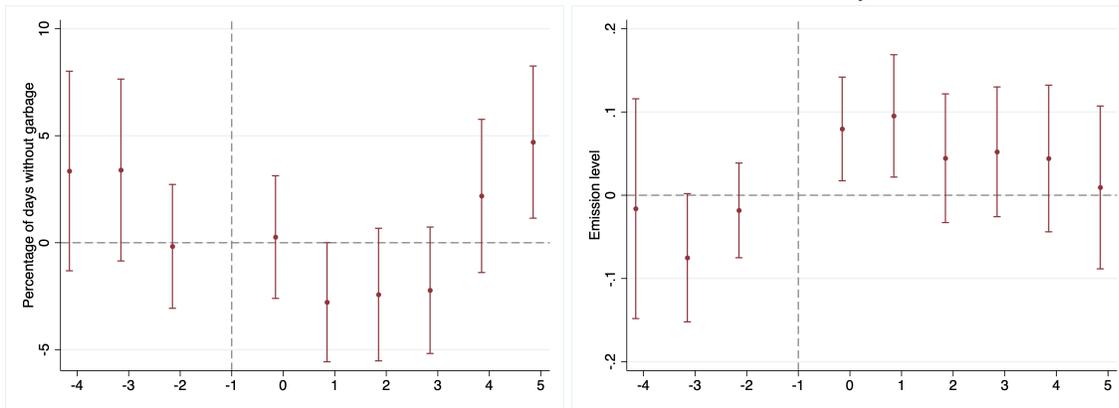


Figure A.6: Market share, emission abatement, technology upgrades not responsive to entry without excess capacity. This figure reports event study estimates (β_j) following the specification of equation (1). The outcome variable is percentage of no-garbage days, the average emission level of all pollutants within a six-month period, and instances of technology upgrades geared towards cleaner emission, respectively in three panels. Treatment is defined as the construction of a nearby incinerator within 50 km radius in the same city. I focus on the sample where excess capacity does not exist. Plant fixed effects and time fixed effects are parsed out. Standard errors are clustered at plant level.

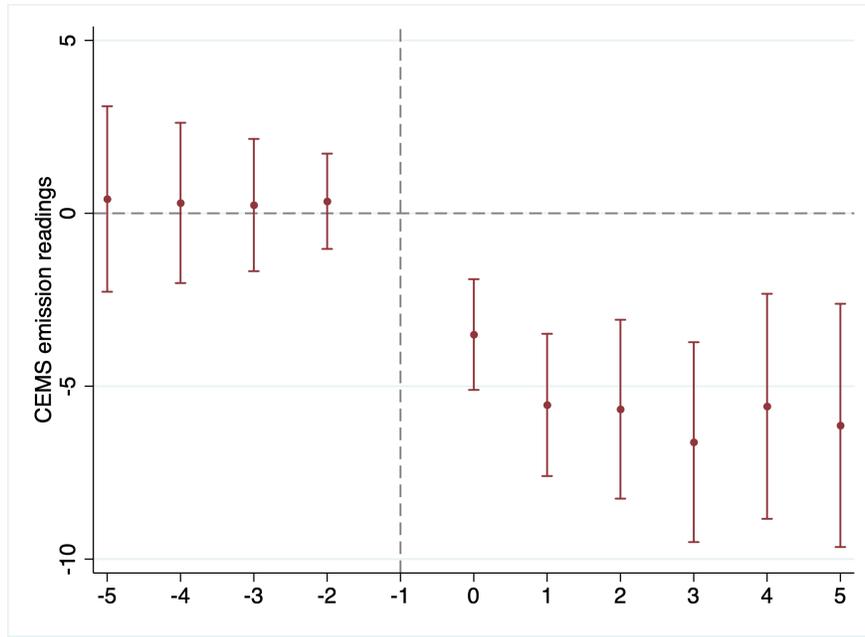


Panel A: Placebo, entrants > 300 km away

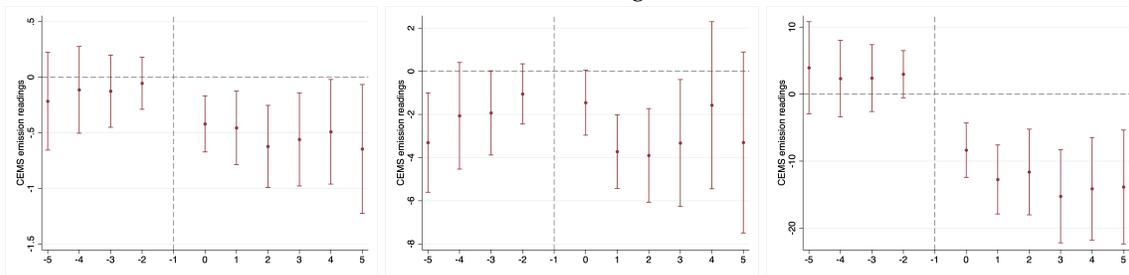


Panel B: Placebo, entrants < 50 km but across provincial boundaries

Figure A.7: Market share, emission abatement not responsive to placebo entry. This figure reports event study estimates (β_j) following the specification of equation (1). The outcome variable is percentage of no-garbage days, the average emission level of all pollutants within a six-month period, and instances of technology upgrades geared towards cleaner emission, respectively in three panels. Placebo treatment is defined as the construction of a nearby incinerator 300 km radius away, in a different city in Panel A, and the construction of a nearby incinerator within 50 km radius, but across provincial boundaries in Panel B. Plant fixed effects and time fixed effects are parsed out. Standard errors are clustered at plant level.



Panel A: Pooled regression



Panel B: Separate regression for PM, SO₂ and NO_x

Figure A.8: Technology upgrade and emission abatement. This set of figures report event study estimates, studying the effects of technology upgrades on emission levels. The dependent variable in the three panels are emission concentration of PM, SO₂, NO_x, respectively. Treatment is defined as the implementation of a large-scale technology upgrade and identified from the universe of procurement contracts. Plant fixed effects and time fixed effects are parsed out. Standard errors are clustered at plant level.

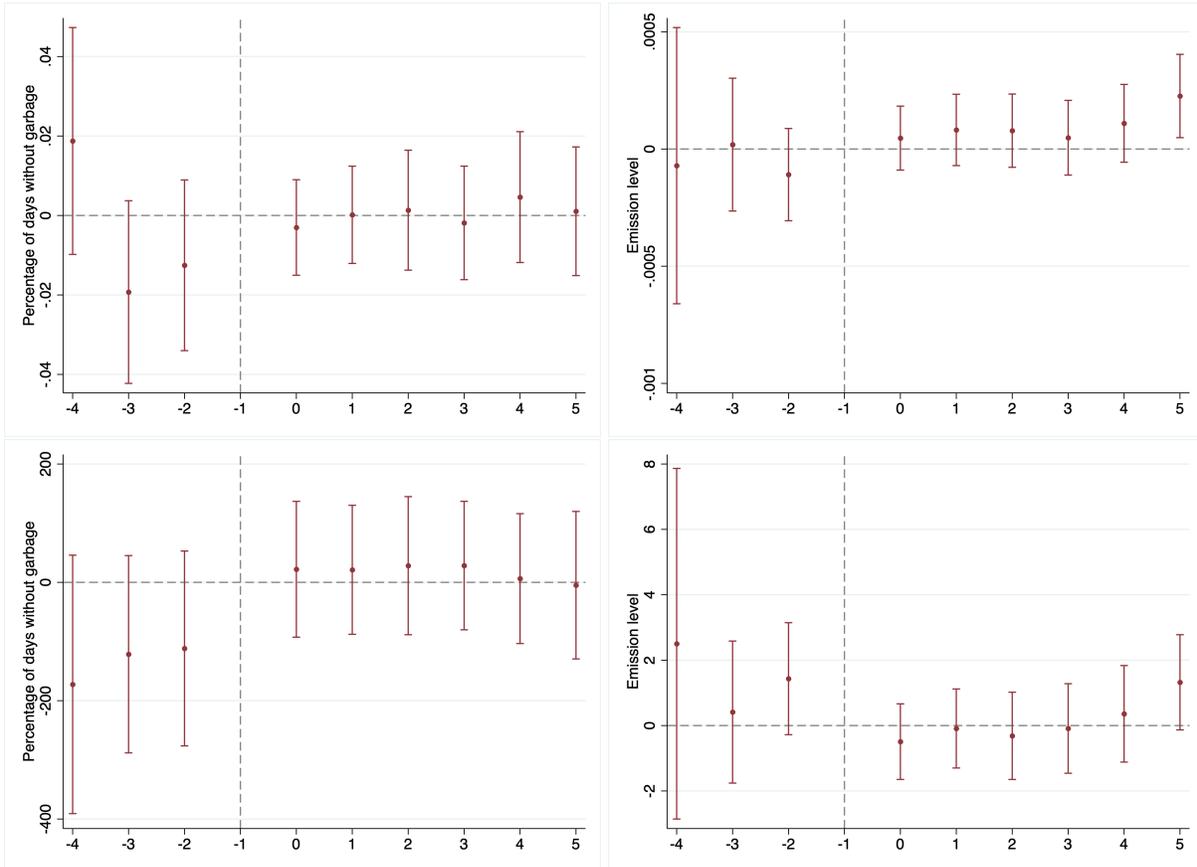


Figure A.9: Heterogeneous market stealing effect, and abatement response by difference in designed capacity and power generation capacity. This figure reports the estimates ($\beta_{j,z}$) following a triple-difference specification of equation (3). The outcome variable is the total number of days within the 6-month period where the furnace is not operating (left panels) and emission readings (right panels). Treatment is defined as the construction of a nearby incinerator within 50 km radius in the same city. I interact the treatment with the difference in designed capacity (top panels) and difference in power generation efficiency (bottom panels) between the incumbent and the new entrant in the first six months. I focus on the sample where excess capacity exists. Plant fixed effects and time fixed effects are parsed out. Standard errors are clustered at plant level.

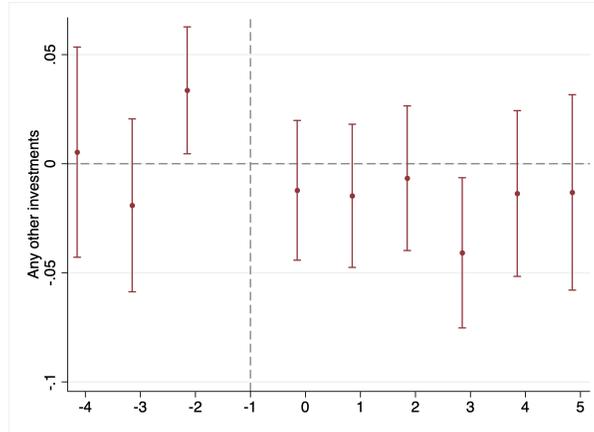


Figure A.10: Other investments in response to entry. This figure reports event study estimates (β_j) following the specification of equation (1). The outcome variable is an indicator of whether firms make a non-environment investment that year. Treatment is defined as the construction of a nearby incinerator within 50 km radius in the same city. I focus on the sample where excess capacity exists. Plant fixed effects and time fixed effects are parsed out. Standard errors are clustered at plant level.