Regulating Conglomerates: Evidence from an Energy Conservation Program in China*

Qiaoyi Chen
Fudan University

Zhao Chen
Fudan University

Zhikuo Liu
Fudan University

Juan Carlos Suárez Serrato
Duke University
& NBER

Daniel Yi Xu
Duke University
& NBER

PRELIMINARY AND INCOMPLETE

January 23, 2021

Abstract

How does energy regulation affect production and energy use within conglomerates? We study the effects of a prominent program aimed at reducing the energy use of large Chinese companies. Difference-in-differences analyses show that regulated firms significantly reduce their energy consumption and output but do not increase their energy efficiency. Using detailed business registration data, we link regulated firms to non-regulated firms that are part of the same conglomerate. We estimate large spillovers on cross-owned non-regulated firms, which increase both output and energy use. We then specify and calibrate a model of conglomerate production that fits our setting and the estimated effects of the regulation. The model quantifies the importance of conglomerate reallocation for aggregate outcomes, the shadow cost of the regulation, and the efficiency gains from using public information on business networks to improve the design of energy regulation.

*We are very grateful for comments from Hunt Allcott, Ashley Langer, Chris Timmins, and seminar participants at UCSC and the University of Oxford. All errors remain our own.

1 Introduction

Balancing economic growth with the negative side-effects of industrialization—such as pollution—is a central problem of governments in emerging economies. Nowhere is this problem more important or consequential than in China. As Figure 1 shows, energy regulation is of national and global importance given that the industrial energy use of China overshadowed that of other leading economics in the early years of the 21^{st} century.

This paper studies the effects of a large program aimed at curbing the energy use of Chinese industrial firms. As these firms are often part of much larger business networks, this paper asks three questions related to the impact of the program on regulated firms and on unregulated firms that are part of the same conglomerate. First, how do regulations impact the within-conglomerate allocation of production and energy use? Second, how does the ability of conglomerates to (partially) escape regulation impact the shadow cost of regulatory programs? Finally, can governments use information on conglomerate networks to improve energy regulation?

The regulation we study—the "Top 1000" program—targeted the largest energy-consuming firms in the most energy-intensive industries. The regulation was designed following examples of "voluntary agreement" programs in developed countries that were aimed at improving energy efficiency. The implementation of the program was adjusted to local institutions and constraints, with the result that, in practice, lowering energy consumption became the main regulatory objective. Since firms regulated by this program accounted for 47% of total industrial energy use in China in 2004, understanding the effects of this regulation is central to broader questions of energy regulation in China.

We combine difference-in-differences research designs with a quantitative model to study the effects of the Top 1000 program. First, using a difference-in-differences strategy, we estimate that regulated firms reduced their energy use by 13–16%. Regulated firms achieve these reductions by lowering output; we find no impact on their energy efficiency. Second, we study how conglomerates reallocate production across related firms. Using a second difference-in-differences design, we find that unregulated firms in the same conglomerate as regulated firms increase both output and energy use. This allows business groups to shift 40% of the output decline in regulated firms to unregulated affiliates. Third, we specify and calibrate a model of conglomerate production that matches our setting and the impacts of the policy on regulated firms. We quantify that the ability of conglomerates to shift production lowered the shadow cost of the regulation from 7.4% of energy costs to 3.4%. Finally, we show that the government can use public information on conglomerate networks to design a conglomerate-level regulation that would achieve the same energy-use reduction at a shadow cost of 2.25% of energy costs.

Overall, our findings show that, while the government was able to reduce the energy consumption of large firms, ignoring the reallocation of activity across business networks signifi-

cantly lowered the impact on energy reduction. While the ability to shift production lowered the shadow cost for regulated conglomerates, the Top 1000 program forced firms to shift production to less productive affiliates. Our results suggest that, by using publicly-available data on business networks, the government can significantly improve the design of energy regulation.

We develop these results in three steps. First, we implement a difference-in-differences strategy using firms in similar industries that were regulated in later years as controls. We use an event-study specification to show that Top 1000 firms had similar trends to unregulated firms prior to the regulation. We estimate that regulated firms reduce their energy use by about 13–16%. These estimates are robust to controlling for industry-by-year and province-by-year fixed effects, as well as to including controls for firm characteristics. Since the regulated firms consumed 670 million tons of coal equivalent in 2004, the direct reduction in energy use amounted to close to 100 million tons of coal equivalent. However, we also document that these firms see a decline in output of between 11–23%, and we do not find meaningful or statistically significant changes in energy efficiency. The lack of gains in energy efficiency suggest two hypotheses. First, that firms had limited potential to increase energy efficiency from a technological perspective—i.e., there were no "low-hanging fruits" (e.g., Allcott and Greenstone, 2012). A second hypothesis is that firms were able to escape the regulatory burden by shifting production to related parties.

Our second set of analyses leverages detailed business registration data to map the conglomerate networks of regulated firms. If regulated firms were able to escape the regulation by shifting production to related parties, we would expect to see an increase in both the output and energy use of firms linked to regulated firms through ownership networks. We test this hypothesis using a difference-in-differences strategy that compares non-regulated but related firms to non-regulated and non-related firms. To ensure that these two groups of firms are similar, we use a matching procedure based on pre-regulation characteristics to find a suitable set of control firms. These analyses show that, after the reform, regulated conglomerates shift production to unregulated firms. Specifically, we find an increase in firm output of 13%, as well as similar increases in other measures, such as profits, sales, assets, and employment. Importantly, we only find increases in the economic activity of related firms when their line of business coincides with the narrowly-defined (4-digit) industry classification of the regulated firm. As a placebo test, we show that firms in other industries do not see an increase in economic activity. Because related firms are smaller than regulated firms, we calculate that conglomerates are able to shift 40% of the output decline in regulated firms to related parties. Thus, while firms were able to escape a significant fraction of the regulatory burden, the lack of total shifting of output and the null effects on the energy efficiency of regulated firms suggest that firms still faced significant costs of improving their energy efficiency. Finally, we also show that, following the regulation, related firms also increased their energy use.

Our last set of analyses use a model of conglomerate production to quantify the economic mechanisms underlying our reduced-form results. The model matches key facts of conglomerate production as well as the reduced-form effects of the regulation. Using the model, we show that size-based regulations such as the Top 1000 program distort the allocation of production across related firms. The model clarifies that conglomerate-level spillovers are a distinct force from market spillovers and that ignoring this mechanism would over-estimate the impact of the policy on regulated firms and on industry-level outcomes. Our calibrated model quantifies the shadow cost of the Top 1000 program at 3.4% of energy costs. The ability of conglomerates to shift activity across related firms decreased the shadow cost from over 7.4% to 3.4% of energy costs, relative to a hypothetical case where regulators successfully prevented conglomerates from shifting production to related firms. Supposing instead that the government regulated the total energy use of the conglomerate would further reduce the shadow cost of the regulation by 34%. Similarly, we also show that an energy tax of 2.25% would yield a similar reduction in energy and that such a policy would not distort the within-conglomerate allocation of production.

This paper is organized as follows. Section 2 describes the policy context and the data we use to measure firm responses to the regulation as well as the ownership networks of regulated firms. Section 3 estimates direct effects of the Top 1000 program on regulated firms and Section 4 estimates indirect effects on unregulated firms that belong to the business networks of regulated firms. Section 5 describes our model of conglomerate regulation and Section 6 calibrates the model. Section 7 uses the model to quantify the shadow cost of the policy and to analyze distributional and aggregate effects of the regulation. Section 8 concludes.

2 Policy Background and Data

This section describes the Top 1000 energy savings program. We also describe the different datasets we use to measure economic activity and energy use as well as our strategy to map the ownership networks of Chinese conglomerates.

2.1 The Top 1000 Program

To save energy and reduce related pollution, the Chinese government's "11th Five Year Plan" (11FYP) set an ambitious goal of reducing the country's energy intensity—defined as energy consumption per unit of GDP—by 20% between 2006 and 2011 (Price et al., 2010). Since the industrial sector accounts for 70% of total energy consumption, the government designed policies that focused on nine energy-intensive industries, which accounted for 80% of the country's industrial energy use. One of these key initiatives is the "Top 1000 Energy Saving Program," which targeted the highest energy-consuming firms in the most energy-intensive industries.

The Top 1000 program was first announced by the National Development and Reform Commission in April 2006, and the corresponding monitoring and assessment measures were released in 2007. The name Top 1000 refers to the 1008 industrial firms in the nine energy-intensive industries with energy consumption above 180 thousand tons of coal equivalent in 2004. The total energy consumption of these 1008 "super firms" was 670 million tons of coal equivalent in 2004, which accounts for 47% of China's industrial energy consumption and 33% of China's total energy consumption. Importantly, since the policy was announced in 2006 and selected firms based on their retrospective 2004 energy consumption, it was not possible to manipulate the list of participants in the program. Moreover, the list of firms regulated by the program did not change during the five year period. Table 1 reports the number of firms and their share of energy consumption in each of the regulated industries. Among Top 1000 firms, firms in the Iron and Steel, Chemical, and Electric Power industries account for around 63% of the firms and 68% of the regulated energy consumption in 2005.

The Top 1000 program was designed based on the belief that Chinese industries could significantly increase energy efficiency at a low cost (e.g., McKinsey & Co., 2009). The program was influenced by other "voluntary agreement" programs in developed countries and had two stated goals: to significantly increase the energy efficiency of these "super firms" and to reduce energy consumption by a total of 100 million tons of coal equivalent by 2011. Given the quick implementation of the program, many aspects of voluntary agreement programs (such as providing technological expertise or financing energy efficiency improvements) played a relatively minor role (Price et al., 2010). In practice, firms were only regulated based on energy use and not on energy efficiency.

To implement the policy, the central government assigned a target reduction in energy use to each provincial government. In turn, local officials negotiated individual quotas with each of the Top 1000 firms. Leaders of provincial governments and state-owned enterprises were then evaluated on whether these energy saving targets were met. As a result, local government officials monitored and enforced the energy saving targets of the Top 1000 firms very closely.[1] This strict supervision is evident in Table A.9, where we report the results of the government's annual assessment. This table shows that the compliance rate is very high. In fact, the total energy saving target was achieved in 2008, two years ahead of schedule. At the end of the 11FYP, the government estimated energy savings of 165.49 million tons of coal equivalent, far beyond the original target of 100 million tons. Among the regulated firms that successfully achieved their

[1]Under the "one vote veto" criteria, officials would not be considered for promotions or awards if the province or any of the local Top 1000 firms did not achieve its target. Similarly, the leaders of state-owned enterprise that did not meet the target did not receive annual bonuses. In this way, the Chinese setting contrasts with other developing country settings where the design of incentives for energy auditors plays a key role (e.g., Duflo et al., 2013).

reported energy saving targets, the average completion rate was over 250%.

Due to the successful experience during the 11FYP, the Top 1000 program was expanded into the "Top 10,000" Energy Savings Program during the "12th Five Year Plan" (12FYP) in 2012. In this case, the Top 10,000 refers to the 16,076 energy-intensive firms with energy consumption above 10 thousand tons of coal equivalent in 2010. These firms account for 60% of China's total energy consumption. As in the Top 1000 program, firms the Top 10,000 firms were required to improve their energy efficiency with a goal of saving a total of 250 million tons of coal equivalent during the 12FYP. Our primary analysis focuses on the Top 1000 firms between 2001 and 2011. Since the industrial firms in the Top 10,000 (but not in the Top 1000) were also energy intensive but were not regulated during the 11FYP, they serve as useful controls in our empirical analysis.

2.2 Firm Data

Our empirical analyses combine several rich data sources. The first dataset we use is the Annual Survey of Industrial Firms (ASIF) from the National Bureau of Statistics (2001-2009 and 2011).[2] This dataset provides detailed information on a firm's industry, address, ownership, output, and financial information, and covers all industrial firms with annual revenue above 5 million RMB (approximately $800,000 USD).

We merge these data with the lists of regulated firms using both firm name and a unique legal identifier. These lists are published by the National Development and Reform Commission. The lists of Top 1000 and Top 10,000 firms include information on the evaluations of each regulated firm, the firm-level energy saving target, and the annual compliance rate.

We complement these data with two additional datasets. First, we use China's Environmental Statistics Database (CESD) from China's Ministry of Environmental Protection from 2001 to 2010, which provides detailed information on firm energy consumption. Second, to fill missing data on firm investment and assets in the ASIF, we also merge data from the Annual Tax Survey (ATS) for years 2009 and 2010.[3]

Panel A of Table 2 shows the results of our data construction. Since Top 1000 and Top 10,000 firms are all large firms, the match rate with the ASIF is very high. We match over 99% of the Top 1000 firms and over 97% of the Top 10,000 firms. We also have a fairly good match rate with the CESD, where we match over 80% of Top 1000 and over 65% of Top 10,000 firms. Overall, our combined datasets capture the majority of the economic activity in the Top 1000 and Top 10,000 firms.

Panel A of Table 2 reports summary statistics for the Top 1000 and Top 10,000 firms in our

[2]As is well known in the literature, data for the 2010 ASIF display a number of irregularities and is often excluded from statistical analyses.

[3]To ensure a high-quality merge, we eliminate firms where differences in investment and assets between the ASIF and ATS exceed 25%.

sample. Our sample includes about 8,700 observations of Top 1000 firms and 87,000 observations of Top 10,000 firms over a period of 10 years. Because the CESD only reports energy consumption from primary sources (e.g., coal, oil, gas), our analyses of energy use and energy efficiency exclude firms in industries that mainly rely on electricity.[4] For this reason, the sample of firms with energy consumption data is smaller.

As we show in Panel A of Table 2, Top 1000 firms are larger, older, more likely to be state-owned, and more export-oriented than Top 10,000 firms. This table also shows that Top 1000 firms are slightly less energy efficient (defined as the ratio of energy use to output) than Top 10,000 firms. However, this difference is driven mostly by industry differences, since Top 1000 firms are more likely to be in energy-intensive heavy industries. Figure 2 plots the relation between firm output and energy efficiency and shows that we find strong scale economies in energy efficiency once we control for industry differences.[5] The fact that larger firms are more energy efficient is central to understanding the effects of regulating the energy consumption of large firms. As this graph shows, the firms targeted by the Top 1000 program are likely more energy efficient than most smaller firms. Output spillovers to smaller firms may then reduce aggregate energy efficiency.

2.3 Mapping Conglomerate Networks

To study how conglomerates responded to the regulation, we first need to identify firms' ownership networks. Our ownership data comes from China's Administrative Registration Database (CARD). These data are collected by the State Administration of Industry and Commerce and include registration information of all firms in China starting in 1980, including firm name, registration number, date of establishment, address, ownership, registered capital and related legal persons. Importantly, the data provides detailed shareholder information, which allows us to construct firm ownership networks at multiple levels.

We construct ownership networks using the four types of linkages displayed in Figure 4. First, we include wholly-owned affiliates of regulated firms as related parties. Second, we also include firms that are at least partially owned by regulated firms. Our data allow us to identify firms that are related through multiple levels of investment. We consider firms to be related if they are owned by a regulated firms by up to six levels of investment relations. Although in practice most related firms are fully owned, we require that the regulated firm owns at least 20% of the related firm at each level of investment. Third, we include shareholders of regulated firms and we allow up to three levels of shareholder links. Finally, we also include firms that are fully or partly owned by the shareholders of a regulated firm. In this case, we also allow for six levels of

[4] In practice, we exclude industries where electricity consumption accounts for more than 30% of the total energy consumption.

[5] Figure 2 plots binned data after controlling for differences in energy efficiency across 4-digit industries.

investment and we restrict ownership to be at least 20% at each level.[6] Since the government has no incentives to help firms shift production, we exclude all related firms that are only connected to regulated firms through government entities.

Panel B of Table 2 reports the results of matching our registration data with our data on firm outcomes. Under our baseline definition, we can identify 80,341 related parties of Top 1000 firms in the CARD. Since most of these firms are small, we match 10,944 firms in the ASIF. In our baseline regressions, we require related firms to be in the same 4-digit industry as a related Top 1000 firm.[7] Since a large number of related parties are service firms or small firms that are not recorded in the ASIF, our main sample of related firms includes 2,776 industrial firms.[8] A potential concern with CARD data is that some of the related firms may not be engaged in production and may, in fact, be holding companies. One advantage of merging the CARD data with the ASIF and the CESD is that this ensures that our results are driven by real economic activity in industrial firms. Since it is likely very hard to shift production to firms in other narrowly-defined industries, we analyze firms in the same 2-digit industry but outside of 4-digit industries in a placebo test.

Panel B of Table 2 also examines the robustness of our network definitions to alternative assumptions. Restricting all types of relation to 2 levels does not have a large effect on our sample of related firms in the same 4-digit industry. Increasing the ownership requirements to 25% has a small effect on the number of related firms and the number of related parties is even somewhat stable when we increase the ownership ratio to 51%. These results suggests that, within narrowly-defined industries, firm ownership networks are very compact.

The merged CARD and ASIF data reveal some interesting patterns of firms that operate in the same 4-digit industry. First, we find that Top 1000 firms have an average of 2.48 related parties in narrowly-defined industries. Second, since Top 1000 firms are, in most cases, the largest firms in each industry, their related parties are smaller. On average, the output of related firms is 18.9% of the output of regulated firms. These facts imply that conglomerates may have significant scope to substitute production with similar technology across related firms. However, it is also unlikely that related parties can fully make up production declines in Top 1000 firms. Third, firms within conglomerates have an interesting relative size distribution. To produce Panel A of Figure 3, we compute each firm's size relative to the largest firm in the group; we then plot the average relative size by firm rank. A striking fact of this graph is that the average relative size within a conglomerate declines sharply in firm rank: the second-largest firm in a

[6]Figure A.11 depicts all the possible links we consider.

[7]A related party of multiple Top 1000 firms will be classified as a same-industry related party as long as it is in the same industry with one of its related Top 1000 firms.

[8]Omitting firms in unrelated industries is unlikely to affect our results since super firms like Top 1000 firms would not be able to shift production to service firms or to very small firms.

conglomerate is only 28% as large as the largest firm, on average. Interestingly, the decline in relative firm size follows an almost geometric decline, a fact that we use in our structural model. Finally, Panel B of Figure 3 shows the relation between the output of the largest firm and the number of firms in a conglomerate. The fact that conglomerates with more firms also have larger leading firms suggests that the number of firms in a conglomerate might depend on technological efficiencies that are shared by all firms in a conglomerate.

3 Effects of the Policy on Regulated Firms

As detailed in Section 2, the Top 1000 program sets energy saving targets for each regulated firm. To study the effects of the policy, we compare the effect on these firms relative other large firms that also operate in energy-intensive industries. Specifically, we use firms that are regulated after 2011 as part of the Top 10,000 program. Because firms in the same conglomerate as a regulated Top 1000 firm may be indirectly affected by the policy, we remove these firms from the set of control firms.

The identifying assumption of this difference-in-differences analysis is that, absent the Top 1000 regulation, the energy use and output of Top 10,000 firms would have trended similarly to that of Top 1000 firms. To provide evidence that these firms had similar trends prior to the implementation of this regulation, we use firm data from CESD to estimate an event-study analysis of the form:

$$Y_{ijkt} = \sum_{\tau=2002}^{2010} \beta_\tau \times Treat_i \times Year_\tau + \alpha_i + \eta_{jt} + \delta_{kt} + \varepsilon_{ijkt}, \tag{1}$$

where Y_{ijkt} is a dependent variable for firm i in industry j province k and year t. $Treat_i$ is an indicator for treatment group, which equals 1 for Top 1000 firms and 0 for Top 10,000 firms. The coefficients β_τ from this specification represent differences in the dependent variable between Top 1000 and Top 10,000 firms in each year. Given that the policy evaluation began in 2007, we identify the effects of the policy relative to 2006. We include firm-level fixed effects α_i in all regressions and we also allow for (4-digit) industry-by-year fixed-effects η_{jt} as well as province-by-year fixed effects δ_{kt}. We cluster standard errors at the firm level.

Figure 5 presents a visual implementation of our difference-in-differences estimation strategy. Panel A in Figure 5 displays the β_τ coefficients when the outcome variable is firm-level energy use (total coal consumption equivalent). This figure shows that, prior to the regulation, our treatment and control firms had similar trends. Additionally, this figure makes clear that, relative to non-regulated firms, the policy indeed succeeded at lowering the energy use of regulated firms.[9] Panel

[9]One potential concern is that our results may be contaminated by mean-reversion. Because firms were

B of this figure compares these year-by-year effects to the overall trend in energy consumption.[10] As this figure shows, the program successfully arrested the explosive growth in the energy use of regulated firms.

To quantify the overall effects of the policy, we estimate difference-in-differences specifications of the form:

$$Y_{ijkt} = \beta Treat_i \times Post_t + X'_{it}\gamma + \alpha_i + \eta_{jt} + \delta_{kt} + \varepsilon_{ijkt}, \tag{2}$$

where $Post_t$ is an indicator that equals one after 2006. In addition to different fixed effects, some specifications also control for firm characteristics X_{it}, which include indicators for state-owned firms and exporting firms, measures of profitability (e.g., return on assets), and firm age. Panel A of Table 3 shows that, on average, the total energy consumption of regulated firms decreased by 13-16%. These estimates are stable across specifications that control for different levels of fixed effects and for firm controls. To understand the magnitude of this effect, recall that regulated firms consumed 670 million tons of coal equivalent in 2004. The coefficients in Table 3 therefore imply annual reductions in energy use of close to 100 million tons of coal equivalent.

To discern whether this reduction in energy use is driven by changes in economic activity or in energy efficiency, we now estimate the effects of the program on firm output (i.e., revenue). Panels C–D of Figure 5 show that, after the reform, firm output in regulated firms also decreased significantly. Indeed, Panel B of Table 3 reports declines in output between 11% and 23%, depending on the specification. Accounting for the declines in output implies that the policy had limited impacts on energy efficiency. Panels E–F of Figure 5 show that we cannot reject the null hypothesis that the policy had no impact on energy efficiency. Based on the specification with both industry- and province-by-year fixed effects of Panel C of Table 3, the 95% confidence interval rules out that the policy increased energy efficiency by more than 4%.

The effects of the policy on regulated firms paint a mixed picture of the effects of the regulation. One one hand, the regulation succeeded at achieving a meaningful reduction in the energy use of energy-intensive firms. However, this reduction did not come about through a significant increase in energy efficiency which—while not directly targeted—was one of the underlying intents of the policy.

One concern with these results is that our difference-in-differences specifications may be contaminated by spillover effects. Indeed, as we show below, non-regulated firms that operate in more tightly regulated industries also see increases in output following the implementation of the

regulated based on their 2004 energy use, one possibility is that regulated firms had idiosyncratically large levels of energy use in 2004 that revert to lower levels in later years. As this and other similar graphs show, outcomes for 2004 are not significantly different than those for 2001-2003 and we also do not see large differences with outcomes in years 2005–2006.

[10]For visual clarity, Panels B, D, and F of Figure 5 follow Ohrn (2018) by plotting trends for the control group that have the same average in the pre-period as the treated group.

regulation. Our results above temper the importance of this concern in two ways. First, Figure 5 shows that control firms do not grow at a faster rates after the regulation is in place. Second, our results are robust to controlling for industry-by-year fixed effects, which may partial out some of the market-level spillover effects. In Section 5, we use our model to clarify the importance of market-level spillovers. While our model is consistent with the existence of spillovers, the estimated effects of the policy are mostly due to declines in the energy use and output of regulated firms.

The next section studies whether conglomerates avoided the burden of the regulation by shifting economic activity to related parties. While the existing literature recognizes the importance of accounting for market-level spillovers, we show that spillovers that are transmitted through ownership networks constitute a distinct margin through which conglomerates adjust to energy regulations.

4 Spillover Effects of the Policy Through Ownership Networks

Regulated firms have strong incentives to shift production to related parties. By shifting production, conglomerates can partially offset the declines in economic activity in regulated firms and prevent competitors from increasing their market share. Moreover, shifting profits to related parties allows conglomerates to comply with the letter of the regulation—if not with its intent—without having to invest in potentially costly improvements in energy efficiency.

To measure the empirical importance of conglomerate spillovers, we use CARD data on the ownership networks of regulated firms to identify firms that may indirectly expand as a consequence of the Top 1000 regulation. We then use matching methods to identify controls firms that are (1) not part of the Top 1000 program, (2) not related to a regulated firm, and (3) that are in the same industry and have a similar size (measured in output) in years prior to the regulation. Using these firms as controls, we then estimate event-study and difference-in-differences analyses using specifications similar to those in Equations (1) and (2). In this setting, however, the $Treat_i$ variable is now an indicator of whether a firm is related to a Top 1000 firm. As we discuss in Section 2, we focus our study of spillovers on firms in the same 4-digit industry as the regulated firm. This restriction follows from the logic that only firms selling similar products as the regulated firms may be able to make up the decline in production of Top 1000 firms.

Figure 6 plots the results of these event-study analyses using ASIF data. Panel A shows that, prior to the regulation, related firms have similar trends in output as unrelated firms. After the regulation, however, firms that are related to Top 1000 firms see significant increases in output that persist for several years. Panel A of Table 4 shows that, after the regulation, related firms

expand by 13%, on average. This table also shows that we obtain very similar results across specifications that include different levels of fixed effects as well as firm-level controls.

To gauge the magnitude of these spillover effects, it is important to account for the number of related parties of each regulated firm as well as for their relative size. On average, Top 1000 firms have 2.48 related parties. However, since the average related firm is only 18.9% as large as its regulated counterpart, we calculate that conglomerates shift close to 41% ($\approx 2.48 \times 18.9\% \times 13.3\%/15.2\%$) of the output decline in regulated firms. This result is informative for a couple of reasons. First, this result shows that conglomerates are not able to fully avoid the regulation. Second, combined with the null effect of the program on the energy efficiency of regulated firms, this result shows that firms were unable or unwilling to increase the energy efficiency of their production processes even if it meant losing profits to competitors.

While the average spillover effect allows us to gauge the magnitude of the overall spillovers, we also document interesting differences in spillovers across related firms. Panel B of Table 4 shows that related firms that are in larger terciles of the size distribution see larger increases in output. This result suggests that larger firms were more able to expand, or, alternatively, that larger firms had excess production capacity. We also find that related firms that are more likely to be controlled by regulated firms see larger increases in output. Specifically, affiliates and firms that regulated firms have an investment stake on see larger output increases than firms related through shareholders or shareholders' investments (see Appendix Figure A.13).

The result that related firms see an increase in economic activity is robust to a number of checks. First, we show that we obtain similar results when we use one-to-one matching or three-to-one matching to find controls for related firms (see Panel A of Appendix Figure A.14). Second, Panel C of Table 4 shows that we find positive spillovers in other measures of economic activity including sales, profits, assets, fixed assets, and employment (see Appendix Figure A.14 for corresponding event studies). Third, we show that only those related firms that operate in the same narrowly-defined industries as the regulated firms—and that could possibly produce substitute output—see an increase in economic activity. Indeed, Panel B of Figure 6 and Panel A of Table 5 show that we find no impact on the output of related firms that operate outside the 4-digit industry of the regulated firm (but that are still in the same 2-digit industry). This placebo test is important, as it rules out the possibility that firms related to large conglomerates saw increases in economic activity after 2007, say, in response to the financial crisis or other shocks or trends.

Having established that conglomerates shift output across related parties, we now explore whether these firms also see changes in energy use and energy efficiency. Panels C and D of Figure 6 report these results using data from CESD. Panel C shows that related firms saw an increase in energy use after the regulation. Panel B of Table 5 shows that energy use in related

firms increased by 34–38% after the regulation. However, since related firms are smaller than regulated firms, they are less likely to be included in the CESD database. While the available data include firms across all affected industries, the number of observations in this panel is smaller than in Panel A of Table 4. For this reason, caution is warranted in ascribing these increases in energy use to all related firms. Panel D of Figure 6 and Panel C of Table 5 shows that these firms do not experience statistically significant changes in energy efficiency.

Overall, we find robust evidence that conglomerates shift production across related parties. On average, this shifting behavior allows conglomerates to recover 41% of the output reduction in regulated firms. Our structural model in the next section uses these results to quantify how the ability to shift production between related firms lowered the shadow costs of the regulation. We find larger spillover effects for larger related parties as well as firms that are more likely to be controlled by regulated firms. Finally, we also find suggestive evidence that related parties also increase their energy use. While we do not find significant changes in the energy efficiency of related firms, the fact that these firms are smaller suggests that they are—on average—less energy efficient than the regulated firms. Therefore, it is likely that the overall spillover effect on energy use was larger than the overall output spillover, leading to compositional losses in energy efficiency.

Market-Level Spillovers

Since related parties can not make up the entire output loss of Top 1000 firms, other firms in regulated industries may be indirectly affected by the energy saving program due to reduced competition. To examine this indirect effect of the regulation, we estimate the following difference-in-differences specification:

$$Y_{ijt} = \beta spillover_j \times Post_t + X'_{it}\gamma + \alpha_i + \tau_t + \varepsilon_{ijt}, \tag{3}$$

where $spillover_j$ is the proportion of the total energy saving targets of Top 1000 firms for industry j relative to the total energy consumption of industry j in 2006. To interpret the coefficient β as the average spillover effect, we normalize the $spillover_j$ variable by the average exposure across all industries. Since the variation in the independent variable is at the industry-year level, we do not include industry-by-year fixed effects in this regression, and we only control for firm fixed effects and year fixed effects.[11] Finally, to ensure that market-level spillovers are not contaminated by ownership-network spillovers, we exclude firms that are related to Top 1000 firms from this specification.

[11] Note also that the variation in $spillover_j$ is absorbed in our previous specifications that include industry-by-year fixed effects.

As shown in Table 6, unregulated firms in industries with stricter regulation increase their output significantly after the policy is implemented. Across all industries, we find that the average market-level spillover led to a 7–10% increase in the output of non-regulated firms. Note that the regressions in the first two columns of this table include both regulated and unregulated industries. We find larger increases (8–14%) when we only include firms in regulated industries. Note that, in this case, the identifying variation is driven solely by differences in regulation intensity across industries.

These results yield a couple of insights. First, these findings further confirm our previous estimates that related parties were not able to make up for all the output loss of the Top 1000 firms. Second, a full accounting of the spillover effects of the regulation needs to include both within-conglomerate spillover as well as market-level spillovers. To study the distribution and aggregate impacts of the Top 1000 program, the next section develops a model of conglomerate production where regulations impact regulated firms, non-regulated firms that are part of the conglomerate, as well as unrelated competitors.

5 A Model of Conglomerates with Regulation

This section presents an industry equilibrium model of conglomerates that is consistent with cross-sectional data patterns and reduced-form responses to the policy of energy regulation. We use the model to interpret our difference-in-differences estimates, to compute the shadow cost of the regulation at the conglomerate level, and to quantitatively evaluate the aggregate and distributional impacts of the regulation.

5.1 Demand and Technology

Our industry equilibrium model draws the structure of product differentiation and monopolistic competition from Melitz (2003). We consider an individual sector with an exogenous aggregate expenditure R. The representative consumer has CES preference over a continuum of varieties $\omega \in \Omega$:

$$U = \left[\int_{\omega \in \Omega} q(\omega)^{\rho} d\omega \right]^{1/\rho},$$

where $q(\Omega)$ represents the consumption level of variety ω and $\sigma = 1/(1-\rho) > 1$ denotes the elasticity of substitution between varieties. Utility maximization by the representative consumer yields the following residual demand curve for each variety ω:

$$q(\omega) = RP^{\sigma-1} p(\omega)^{-\sigma},$$

where $P = [\int_{\omega \in \Omega} p(\omega)^{1-\sigma} d\omega]^{\frac{1}{1-\sigma}}$ is the aggregate price index.

We define a "conglomerate" in our model as a variety ω, which can be manufactured by its multiple affiliates. Each conglomerate starts with a central producer—the model counterpart of a Top 1000 firm. Conglomerates have heterogeneous efficiencies ϕ, which are drawn from the distribution $G(\phi)$ with density $g(\phi)$. Production at each firm requires total capital k subject to decreasing returns to scale α.[12] The assumption that $\alpha < 1$ matches the observation that conglomerates often expand capacity by setting-up multiple affiliates. Intuitively, conglomerates may operate more firms as a way to escape decreasing returns to scale, and as a way to share production knowledge ϕ across firms. However, as we show in Panel A of Figure 3, conglomerates are not able to replicate the same scale across related firms. To match this fact, we assume that the productivity of the i^{th} affiliated firm is $\delta^{i-1}\phi$. This assumption can be interpreted as either a limit on the span of managerial control or as a measure of imperfect knowledge sharing across firms. In addition to the regular production capital, each manufacturing establishment incurs a fixed outlay of capital denoted by f. This assumption is motivated by two facts. First, that conglomerates have a finite number of affiliates. Second, as shown in Figure 2, that firms face economies of scale in energy efficiency. These assumption imply that, for a conglomerate with efficiency ϕ, the production technology of the i^{th} affiliate is then:

$$q_i = \delta^{i-1}\phi(k_i - f)^{\alpha}.$$

Similarly, the total output of a conglomerate with efficiency ϕ that controls n affiliates is:

$$q(\phi, n) = \sum_{i=1}^{n} \delta^{i-1}\phi(k_i - f)^{\alpha}.$$

5.2 Utilized Capital and Conglomerate Size

The conglomerate faces two decisions: how many firms to operate and how to utilize capital across affiliates. To match the empirical setting, we interpret utilized capital as the general energy use of a firm.[13] Formally, the conglomerate solves the following problem:

$$\pi(\phi) = \max_{n} \{\pi(\phi, n)\}$$

$$\pi(\phi, n) = \max_{k_1, k_2 \ldots k_n} \left\{ R^{1-\rho} P^{\rho} \left[\sum_{i=1}^{n} \delta^{i-1}\phi(k_i - f)^{\alpha} \right]^{\rho} - r \sum_{i=1}^{n} k_i \right\}, \tag{4}$$

[12] To focus on the production plans of firms in a conglomerate, we abstract from other variable inputs, such as labor or intermediate goods. This abstraction is justified under the assumption that these other inputs can be flexibly adjusted and are already "optimized out."

[13] This interpretation is consistent with research studying the macroeconomic and productivity roles of energy use (e.g., Burnside et al., 1995).

where $\pi(\phi, n)$ is the total profit of a conglomerate with efficiency ϕ given that it chooses to control n affiliates. The conglomerate takes the industry aggregates R and P and the price of utilized capital r as given.

Assuming that a conglomerate has n affiliates, the capital utilization across affiliates is determined by the n first order conditions of Equation 4:

$$\frac{\partial \pi}{\partial k_i} = \underbrace{R^{1-\rho}P^{\rho}}_{\text{Market Demand}} \underbrace{\rho\left[\sum_{i=1}^{n} \delta^{i-1}\phi(k_i - f)^{\alpha}\right]^{\rho-1}}_{\text{Residual Revenue}} \underbrace{\delta^{i-1}\phi\alpha(k_i - f)^{\alpha-1}}_{\text{Marginal Product}} - r = 0 \quad (1 \leq i \leq n). \tag{5}$$

An important insight of this expression is that conglomerates internalize the marginal product of inputs across firms through the residual revenue term, which is common to all firms in the conglomerate. As we discuss below, the impact of energy regulations on the residual revenue term is key to understanding the difference between within-conglomerate and market-level spillovers. The following proposition characterizes the choice of optimal inputs for each affiliated producer using the first order conditions described in Equation 5.

Proposition 1 (Within Conglomerate Distribution)**.** *Absent regulation, the variable input and the output of affiliated producers in a conglomerate follows a decreasing geometric sequence given by:*

$$\frac{q_i}{q_1} = \frac{k_i - f}{k_1 - f} = \delta^{\frac{i-1}{1-\alpha}}. \tag{6}$$

The within-conglomerate distribution described in Proposition 1 is broadly consistent with the empirical pattern in Panel A of Figure 3, where the average output of the second largest affiliate in a conglomerate is less than 30 percent of that of the largest one, and where the output of other affiliated producers in the conglomerate decreases exponentially with their rank i. Equation 6 links this distribution to two model parameters. First, the size gap among affiliates will be larger if the within-group knowledge depreciation is more severe (lower δ). Second, if firms are closer to having constant-returns-to-scale production (α closer to one), the conglomerate will concentrate more activity in its top producer, which increases the dispersion of the within-group size distribution.

We now characterize the optimal output of a conglomerate of size n with efficiency ϕ:

$$q^*(\phi, n) = t_n(\phi)^{\frac{1-\alpha}{1-\alpha\rho}} \phi^{\frac{1}{1-\alpha\rho}} \left(\frac{R^{1-\rho}P^{\rho}\alpha\rho}{r}\right)^{\frac{\alpha}{1-\alpha\rho}}. \tag{7}$$

The term $t_n(\phi) = \frac{1-\delta^{\frac{n(\phi)}{1-\alpha}}}{1-\delta^{\frac{1}{1-\alpha}}}$ sums-up the conglomerate-level knowledge across its affiliates. $t_n \geq 1$ and is increasing in n, but at decreasing rate. This expression shows that the optimal conglom-

erate output $q^*(\phi, n)$ is increasing in ϕ. The output of the i^{th} affiliated producer is then:

$$q_i^*(\phi, n) = \left[\frac{\delta^{\frac{i-1}{1-\alpha}}}{t_n(\phi)}\right] q^*(\phi, n).$$

We now consider the optimal number of affiliates. A conglomerate of size n and efficiency ϕ will add a new affiliate as long as:

$$\pi(\phi, n+1) - \pi(\phi, n) = \left(t_{n+1}{}^{\frac{(1-\alpha)\rho}{1-\alpha\rho}} - t_n{}^{\frac{(1-\alpha)\rho}{1-\alpha\rho}}\right)\left(\frac{1}{\alpha\rho} - 1\right)\left(\frac{R^{1-\rho}P^\rho\alpha\rho}{r^{\alpha\rho}}\right)^{\frac{1}{1-\alpha\rho}} \phi^{\frac{\rho}{1-\alpha\rho}} \geq rf. \quad (8)$$

The first term in this expression shows that adding a new affiliate can improve the conglomerate's revenue and profit by lowering its overall marginal cost. On the other hand, the conglomerate incurs a fixed cost of rf when adding a new firm. While the marginal benefit of adding a new affiliate is increasing in ϕ, it is also decreasing in the number of existing affiliates n. Since the fixed cost is the same for all affiliates, Equation 8 guarantees the existence of a cutoff value ϕ_n, where conglomerates with efficiency $\phi > \phi_n$ will operate at least n affiliated producers.

Proposition 2 (Optimal Conglomerate Size)**.** *Without regulation, the optimal number of firms in a conglomerate n is non-decreasing in its fundamental efficiency ϕ. For $n > 1$, a conglomerate chooses to have n affiliated producers when $\phi_n \leq \phi < \phi_{n+1}$, where*

$$\phi_{n+1} = \frac{r^{\frac{1}{\rho}} f^{\frac{1-\alpha\rho}{\rho}}}{\left(t_{n+1}{}^{\frac{(1-\alpha)\rho}{1-\alpha\rho}} - t_n{}^{\frac{(1-\alpha)\rho}{1-\alpha\rho}}\right)^{\frac{1-\alpha\rho}{\rho}} C_\rho^{\frac{1}{\rho}} \left(R^{\frac{1-\rho}{\rho}} P\right)}, \quad (9)$$

and where $C_\rho = (1-\alpha\rho)^{1-\alpha\rho}(\alpha\rho)^{\alpha\rho}$.

The prediction from Proposition 2 is consistent with our empirical observation in Panel B of Figure 3 that conglomerates with higher efficiency have, on average, a larger number of affiliated firms.

5.3 Equilibrium

The unique equilibrium of the model is characterized by the zero profit condition and the free entry condition. Conglomerates operate whenever

$$\pi(\phi) \geq 0 \;\Rightarrow\; \phi \geq \phi^* = \frac{r^{\frac{1}{\rho}} f^{\frac{1-\alpha\rho}{\rho}}}{C_\rho^{\frac{1}{\rho}} \left(R^{\frac{1-\rho}{\rho}} P\right)}. \quad (10)$$

Equation 10 shows that only firms with $\phi > \phi^*$ choose to participate in the market.[14] Letting M denote the mass of active firms, the aggregate price index is given by:

$$P = \left[\int_{\phi^*}^{\infty} p(\phi)^{1-\sigma} \frac{g(\phi)M}{1 - G(\phi^*)} d\phi\right]^{\frac{1}{1-\sigma}}.$$

[14]Comparing Equations 9 and 10 shows that $\phi^* < \phi_2$, which defines the range of efficiency of single-firm producers.

To enter the market, an entrepreneur pays an entry cost f_e. Upon entry, the efficiency of the conglomerate ϕ is realised. Since the conglomerate will only operate if $\phi > \phi^*$, the free entry condition is given by:

$$\int_{\phi^*}^{\infty} \pi(\phi) g(\phi) d\phi - f_e = 0. \tag{11}$$

An equilibrium is given by the exit threshold ϕ^* and the mass of active conglomerates M, such that (1) conglomerates make optimal allocation and size decisions, (2) the product market clears, and (3) the zero profit and free entry conditions (Equations 10–11) are satisfied.

5.4 Effects of the Top 1000 Program

Since the Top 1000 program targeted very large firms, we assume that only conglomerates with ϕ above an efficiency level $\tilde{\phi}$ are subject to the regulation. For each conglomerate, the regulation sets a proportional input quota for its largest firm, which is the model counterparts of a Top 1000 firm. Specifically, the input of regulated firms can not exceed $\bar{k}_1(\phi) = \xi k_1^*(\phi)$, where $\xi < 1$ and k_1^* is the unregulated optimal input use.

Letting λ be the Lagrange multiplier associated with the regulatory constraint, we can re-state the conglomerate's first order conditions for k_i $(1 \leq i \leq n)$ as:

$$\frac{\partial \pi}{\partial k_i} = \underbrace{R^{1-\rho} P^{\rho}}_{\text{Market Demand}} \underbrace{\rho \left[\sum_{i=1}^{n} \delta^{i-1} \phi (k_i - f)^{\alpha} \right]^{\rho-1}}_{\text{Residual Revenue}} \underbrace{\delta^{i-1} \phi \alpha (k_i - f)^{\alpha-1}}_{\text{Marginal Product}} = r + \underbrace{\lambda(\phi) \mathbb{I}[i=1]}_{\substack{\text{Shadow Cost} \\ \text{of Regulation}}}. \tag{12}$$

This equation shows that the regulation distorts the allocation of inputs within a conglomerate by adding a shadow cost $\lambda(\phi)$ to the input of the regulated firm. This distortion leads conglomerates to allocate more inputs to the unregulated firms, relative to the case without the regulation. While conglomerates substitute production across firms, the regulation leads to an overall reduction in the conglomerate's output.

The regulation also distorts input use across conglomerates. Since only conglomerates with $\phi > \tilde{\phi}$ are part of the Top 1000 program, the regulation increases the overall input cost of regulated conglomerates relative to unregulated conglomerates. Additionally, because conglomerates with more affiliates can shift more production to related parties, conditional on being regulated, more efficient conglomerates (higher ϕ) are subject to a smaller shadow cost $\lambda(\phi)$.

To see how the regulation in our model connects to our difference-in-differences analysis, note that we can write the conglomerate j's revenue from affiliate i as follows:

$$\ln \text{Revenue}_{ij} = \underbrace{\ln(\text{Production Share}_{ij})}_{\text{Allocation Effect}} + \underbrace{\rho \ln \left(\sum_{i \in j} q_{ij} \right)}_{\text{Residual Revenue}} + \underbrace{\ln(R^{1-\rho} P^{\rho})}_{\text{Market Demand}}, \tag{13}$$

where Production Share$_{ij} = q_{ij} / \sum_{i \in j} q_{ij}$.[15] Equation 13 clarifies the three ways in which the Top 1000 program impacts the revenue of regulated firms. First, as we discuss above, when firm i is regulated, the conglomerate is forced to reallocate inputs to other firms, which lowers the production share in regulated firms. Second, since the marginal cost goes up at the conglomerate level, the market share of its variety decreases, which lowers the group's residual revenue as well as the revenue of the regulated firm. Finally, since all regulated firms in the industry contract, the Top 1000 program has an equilibrium impact on the industry-level price index P, which, in turn, has a countervailing effect on the revenue of the regulated firm.

We can also use Equation 13 to characterize the impact of the regulation on the control firms in our difference-in-differences analyses in Section 3. Since these firms are not regulated or related to Top 1000 firms, the regulation does not impact the within-conglomerate allocation of production. Control firms see an increase in their residual and firm-level revenue as the market reallocates demand. As in the case of the regulated firms, non-regulated firms also benefit from the equilibrium impact on market demand.

This discussion clarifies the interpretation of our estimates of the impacts of the regulation on the output of Top 1000 firms. One benefit of our difference-in-differences approach is that the common market demand effect cancels out.[16] Our estimates then capture three effects: (1) a negative allocation effect on regulated affiliates, (2) a negative residual revenue effect on regulated conglomerates, and (3) a positive residual revenue effect on control firms.

A similar discussion allows us to interpret our estimates of the spillover effects of the regulation through ownership networks. We first note that the control firms have the same impact as in our estimates of the direct effects of the regulation. Firms related to Top 1000 firms share the residual revenue and market demand with the regulated firms. In contrast to regulated firms, the allocation effect increases the revenue of related parties. Since the market demand effect also cancels out in these estimations, our difference-in-differences estimates of spillover effects capture the following three mechanisms: (1) a positive allocation effect on related affiliates, (2) a negative residual revenue effect on regulated conglomerates, and (3) a positive residual revenue effect on control firms.

These insights highlight the importance of interpreting quasi-random estimates through the lens of a model that accounts for within- and across-conglomerate reallocation of production, as well as equilibrium impacts on industry-level prices. In Section 6, we calibrate the model parameters to fit both distributional features of the data as well as our estimates of the regulation on regulated firms. We the show that the predictions of the calibrated model for spillover effects on relates firms and at the market level are broadly consistent with our estimates from Section 4.

[15] Equation 13 follows by multiplying the conglomerate j's inverse residual demand by affiliate i's production.

[16] Even in the cases where the equilibrium response is industry- or location-specific, the market demand effect would be absorbed by industry-by-year and province-by-year fixed effects.

In Section 7, we use our calibrated model to decompose the importance of the different mechanism behind our difference-in-differences estimates, to measure the shadow cost of the regulation, and to quantify the aggregate and distributional impacts of the Top 1000 program.

6 Model Calibration

This section calibrates the key parameters of the model to quantitatively match the data patterns prior to the regulation. We then show that the model can match the observed firm-level effects of the regulation in a simulated counterpart of the Top 1000 program.

6.1 Parameters and Identification

We briefly describe the set of structural parameters of the model and how they are identified by the data. We parameterize the conglomerate efficiency distribution $G(\phi)$ with a Log-Normal distribution with mean zero and standard deviation σ_m. Our model has the following parameters $(\sigma, \alpha, f, \delta, f_e, \sigma_m)$, which include the elasticity of substitution σ, the returns to scale parameter α, the affiliate-level fixed cost f, the within-conglomerate size depreciation δ, the conglomerate-level entry cost f_e, and the dispersion of the efficiency distribution σ_m.

We start by setting the values of two parameters based on previous estimates. We follow the literature by calibrating $\sigma = 4$ (Melitz and Redding, 2015). We use the estimate of returns to scale of $\alpha = 0.9$ from Burnside et al. (1995), who also use energy data to proxy for utilized capital.[17]

We then use our data to calibrate the remaining model parameters in two steps. First, we use the within-conglomerate relative size distribution in Panel A of Figure 3 to obtain the parameter δ. Recall that Equation 6 shows that the output of affiliates depreciates in rank by the factor $\delta^{\frac{1}{1-\alpha}}$. We therefore set $\delta^{\frac{1}{1-\alpha}} = 0.28$ to match the average relative size of the three largest affiliates within each conglomerate. In our model, the fixed cost f generates economies of scale in energy efficiency. To calibrate f, we use the observed relation between energy use and firm sales and we set $f = 37$ thousand RMB (approximately 6,000 USD) to match the intercept in this relation (see Figure A.16).

In the second step of our calibration, we simulate the industry equilibrium to match the observed within-industry firm size distribution prior to the implementation of the Top 1000 program. This process identifies the last two parameters of the model: σ_m and f_e. Figure 8 shows that our model does a good job of replicating both the observed firm-size distribution and the concentration of output.

[17] Conventional estimates of returns to scale range from 0.8 to 0.9 depending on aggregation and time period.

6.2 Model Response to the Top 1000 Program

We calibrate two additional parameters to implement the Top 1000 program in our model. As discussed in Section 5, our version of the regulation targets conglomerates with efficiency level ϕ above $\tilde{\phi}$. We choose the threshold $\tilde{\phi}$ to match the share of total energy consumed by regulated firms within energy-intensive industries. We choose the value of $\tilde{\phi} = 8.84$, which reproduces the fact that regulated firms account for 58% of the total energy consumption in energy-intensive industries.

Finally, we choose the input quota, ξ, by matching the difference-in-differences estimate of the program on the output of regulated firms. To match our estimate of a 15.7% reduction in output from Panel A of Table 3, we set the value of $1 - \xi$ to 13%. This value is also consistent with the intent of the policy to save 100 million tones of coal equivalent and the fact that regulated firms consumed 800 tones of coal equivalent in 2006, i.e., $\frac{100}{800} = 12.5\%$, (Price et al., 2010). Panel A of Table 7 collects the calibrated parameters of the model.

Panel B of Table 7 confirms that this calibration replicates the difference-in-differences estimate on the output of regulated firms. The table also shows that our model does a good job of matching several un-targeted moments including the effects on the energy use of regulated firms and the spillover effects on the output of related firms. Our model also implies reasonable market-level effects that are broadly consistent with the market-level spillovers. Relative to our large spillover estimate on the energy use of related firms, the model predicts a smaller effect that is marginally outside of the 95% confidence interval of our estimated effect.[18]

7 Policy Simulations

This section uses our calibrated model to conduct three exercises. First, we use our model to decompose the roles of market-level and conglomerate-level spillovers. Second, we quantify the shadow cost of the regulation and measure the extent to which the ability of conglomerates to shift production across related firms lowered the shadow cost of the policy. Finally, we consider whether the government can use information on the business networks of regulated firms that achieve the same reduction in energy use, but at a potentially lower shadow cost.

7.1 Aggregate and Distributional Effects of the Policy

Table 8 reports the effects of simulating the Top 1000 program under different assumption. Panel A reports the impact of the policy on regulated conglomerates. The first column shows that the

[18] One potential reason why the estimated effect is larger than the model prediction, is that, as we discuss in Section 4, the data coverage on the energy use of related firms is relatively sparse in the CESD data, which implies that our spillover estimate may not be representative of all firms.

output of regulated conglomerates would decrease by 8.4% if the conglomerate was not able to reallocate production across related firms and absent equilibrium effects. The second column shows that allowing conglomerates to reallocate activity would result in a smaller decline of 5% for regulated conglomerates. Alternatively, column 3 shows that allowing the industry price to adjust would lead to a smaller decline in output. The fourth column shows the importance of accounting for both market- and conglomerate-level spillovers. Ignoring the conglomerate-level spillover would result in a predicted 6% reduction for regulated conglomerates, relative to a 3% reduction when we account for both margins.

We now provide intuition for the different mechanisms behind the conglomerate and the market spillovers. In the case of the conglomerate spillover, our model assumes that the regulation forces conglomerates shift production to less productive affiliates. While the output of these affiliates is perfectly substitutable, decreasing returns to scale and knowledge depreciation limit the degree to which it is profitable for conglomerates to make up the reduction in the output of regulated firms. In contrast, competing firms face an increase in residual demand, but their output is imperfectly substitutable with that of regulated conglomerates. Based on the parameterization of our model, we find that the within-conglomerate substitution accounts for 60% of the overall adjustment between the no spillover case and the case with both spillovers (with the market adjustment accounting for the remaining 40%).

Panel A of Table 8 also computes impacts on input use and energy efficiency. According to out model, we see slightly larger reductions in energy use and small gains in energy efficiency at regulated conglomerates.

Panel B of Table 8 computes the aggregate impacts of the policy. The model implies smaller reductions on aggregate, as non-regulated conglomerates expand in response to the regulation. This table continues to show the importance of separately accounting for conglomerate- and market-level spillovers. Note that since this table considers the output of unregulated firms, the market-level spillovers are relatively more important than in Panel A. At the aggregate level, we find slightly smaller declines in energy use as well as small declines in energy efficiency. The decline in energy efficiency is also consistent with the result that, as production is shifted away from large firms, the average size of firms in the industry declines.

So far, we focused our discussion on the short-run effects of the policy, which ignore the entry of new conglomerates. This focus matches the short-run nature of our empirical analysis. However, we can also use the model to simulate longer-run impacts once new firms enter the market. In the long run, Panel A shows that regulated conglomerates lose market share as new firms enter. In Panel B, however, the entry of new, smaller firms completely makes up the output loss and the reductions in energy use, which further lowers the average firm size in the industry.

7.2 Shadow Cost of the Policy

We now use the model to quantify the shadow cost of the policy. This exercise allows us to measure the extent of conglomerate-specific distortions and to benchmark the cost of the program relative to alternative regulations.

We compute the shadow cost of the policy using Equation 12. Recall that in this equation, the shadow cost only enters the first order condition for the regulated firm. However, the shadow cost impacts the choice of energy utilization at all related firms through the residual revenue term. In our model, firms face a lower shadow cost if they are able to substitute production away from the regulated firm. Firms with higher efficiency ϕ and that have a larger of affiliates will therefore face a lower shadow cost. Additionally, the size of the shadow cost depends on our model parameters, which capture both returns to scale α and the depreciation of knowledge δ.

Panel A of Figure 9 plots the implied shadow cost as a function of efficiency ϕ. The blue line plots the shadow cost of our simulated Top 1000 program. This shadow cost is zero for firms with $\phi < \tilde{\phi}$ and jumps to close to an average of 3.4% for regulated firms. Since the shadow cost has the same scale as the cost of inputs, we can interpret this value as an equivalent tax on energy use. While 3.4% might seem like a small number, recall that that inputs constitute a much larger tax base, especially relative to profits.[19]

Our model allows us to quantify how the ability of conglomerates to shift production to related firms lowered the shadow cost of the regulation. The orange line in Panel A of Figure 9 plots the shadow cost assuming that regulated firms were not able to shift production to related parties. In this case, the shadow cost of the regulation would have been 7.5% of input costs, which is 120% larger than in the baseline case.

Finally, the model allows us to consider alternative forms of energy regulation. To make the regulations comparable, we ensure that they result in a similar energy reduction as the simulated Top 1000 regulation. S uppose that instead of targeting the energy use of large firms, the government used information on the business networks of energy intensive firms—which is publicly available—to design the regulation at the conglomerate level. The green line in Panel A of Figure 9 shows that this conglomerate-level regulation would lower the shadow cost of the regulation by 34%, to 2.25%. Intuitively, the Top 1000 regulation limited output at the most productive firms. A conglomerate-level regulation is less distortionary since it allows conglomerates to equalize the marginal product of inputs across firms, while reducing their overall input use.

Panel B of Figure 9 provides more detail into how this shadow cost varies by ϕ. The green line shows that conglomerates with a higher ϕ face a lower shadow cost of the regulation, and

[19] Indeed, in models with constant marginal cost and with a similar value of σ, inputs are $(\sigma - 1) = 3-$times as large as profits. An equivalent profit tax would then be 10.2%.

that $\lambda(\phi)$ drops discontinuously at the threshold values ϕ_n where conglomerates add additional affiliates. However, the magnitude of these differences pales in comparison to the difference in the shadow cost between regulated and non-regulated firms. Finally, this graph also plots a size-dependent energy tax that would result in the same energy reduction as the Top 1000 program. The pink line shows that, instead of monitoring large firms, the government could have instead reduced energy use by the same amount with an energy tax of 2.25%.

8 Conclusion

This paper studies the effects of a prominent energy conservation program in China. We combine detailed data on energy use and business networks to study the effects of the regulation on both regulated firms and on firms that are not regulated but that are part of the same conglomerate. While the program leads regulated firms to decrease energy use, this decrease is driven by a decline in production output and not by an increase in energy efficiency. We show that the program leads to large increases in the output and energy use of unregulated firms that are part of the same conglomerate. By shifting production to related firms, regulated conglomerates escape close to 40% of the regulation-driven output reduction. The fact that regulated conglomerates are unable to fully shift lost output to related firms and that we find no impacts on the energy efficiency of regulated firms implies that regulated firms found it costly to increase their energy efficiency.

We calculate the shadow cost of the regulation using a model of conglomerate production that matches our setting and the reduced-form effects of the regulation. The model shows that—even with the ability to shift some production to related firms—the regulation increased the cost of conglomerate production by 3.4%. We show that this cost could be lowered by targeting the regulation at the conglomerate level or by imposing an energy use tax.

Overall, this paper shows that the economic effects and the efficacy of policies that target large firms are modulated by substitution along ownership networks. Since ownership networks are public information, the results of our paper reveal a potential improvement to existing energy regulations.

References

Allcott, Hunt and Michael Greenstone, "Is There an Energy Efficiency Gap?," *Journal of Economic Perspectives*, February 2012, *26* (1), 3–28.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo, "Capital utilization and returns to scale," *NBER macroeconomics annual*, 1995, *10*, 67–110.

Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan, "Truth-telling by Third-party Auditors and the Response of Polluting Firms: Experimental Evidence from India*," *The Quarterly Journal of Economics*, 09 2013, *128* (4), 1499–1545.

McKinsey & Co., "Unlocking Energy Efficiency in the U.S. Economy," Technical Report 2009.

Melitz, Marc J., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 2003, *71* (6), 1695–1725.

Melitz, Marc J and Stephen J Redding, "New trade models, new welfare implications," *American Economic Review*, 2015, *105* (3), 1105–46.

Ohrn, Eric, "The Effect of Corporate Taxation on Investment and Financial Policy: Evidence from the DPAD," *American Economic Journal: Economic Policy*, May 2018, *10* (2), 272–301.

Price, Lynn, Xuejun Wang, and Jiang Yun, "The challenge of reducing energy consumption of the Top-1000 largest industrial enterprises in China," *Energy Policy*, 2010, *38* (11), 6485–6498.

Figures

Figure 1: Cross-Country Differences in Industrial Energy Use

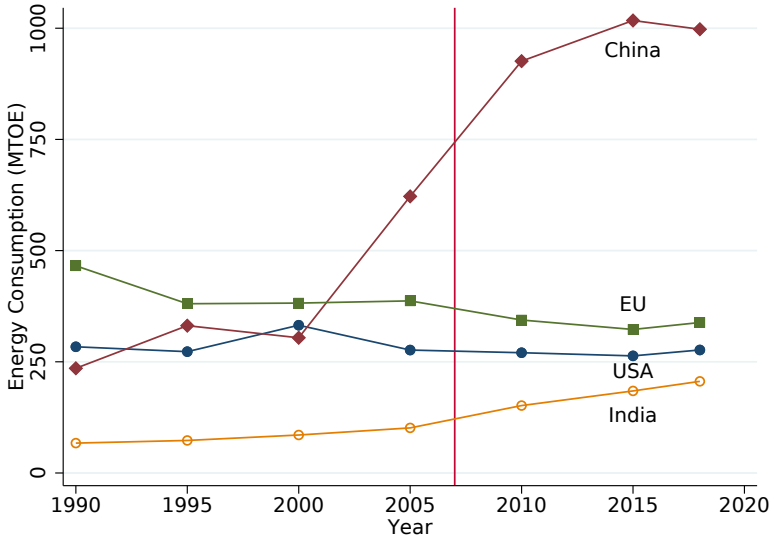


Figure 2: Economies of Scale in Energy Efficiency

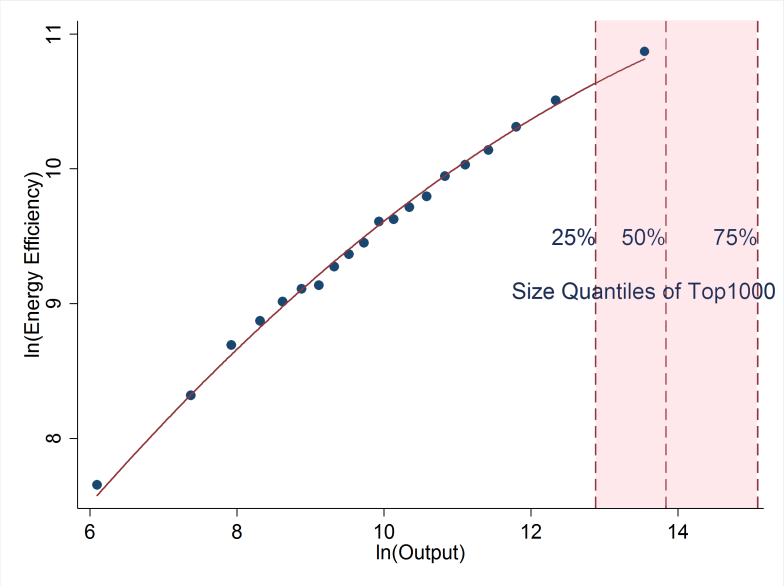
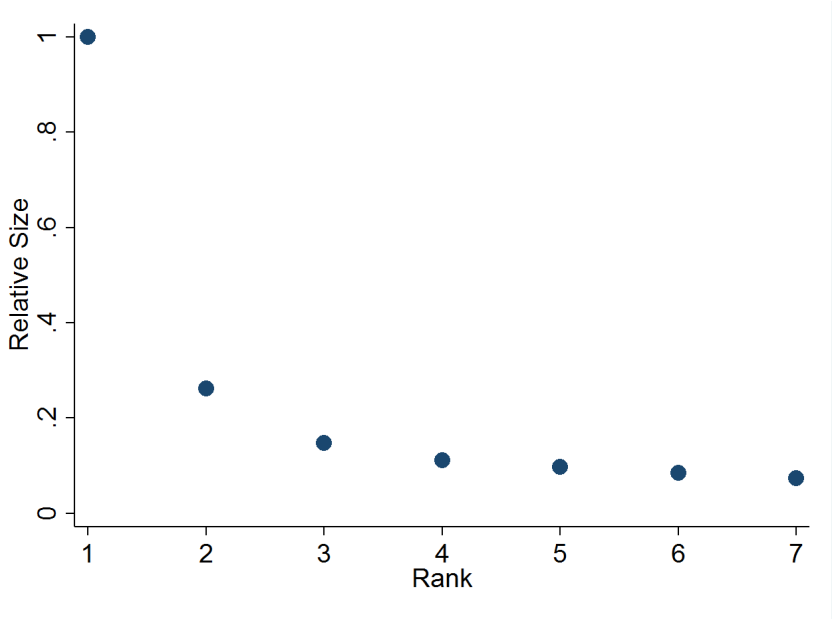


Figure 3: Conglomerate Size and Production Allocation

A. Relative Firm Size



B. Output and Conglomerate Size

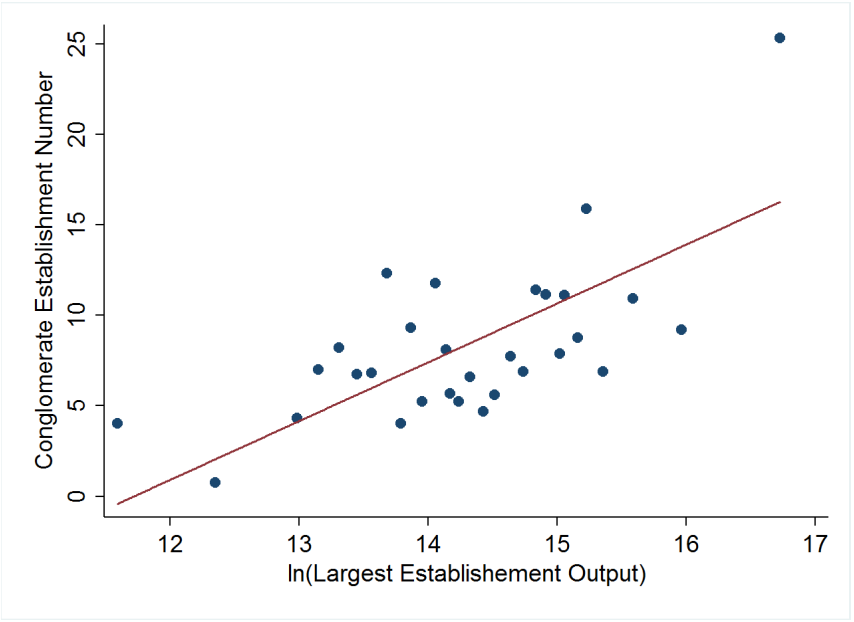
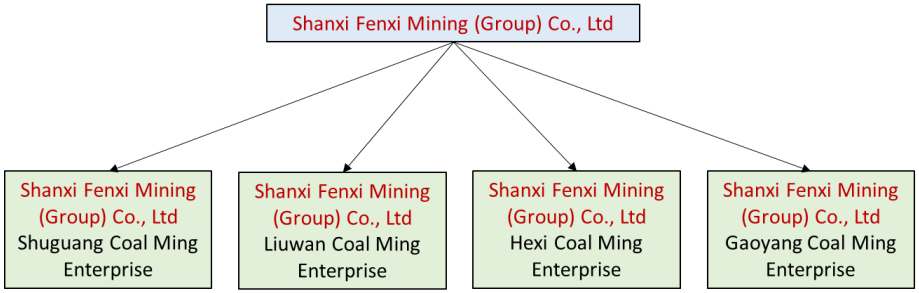
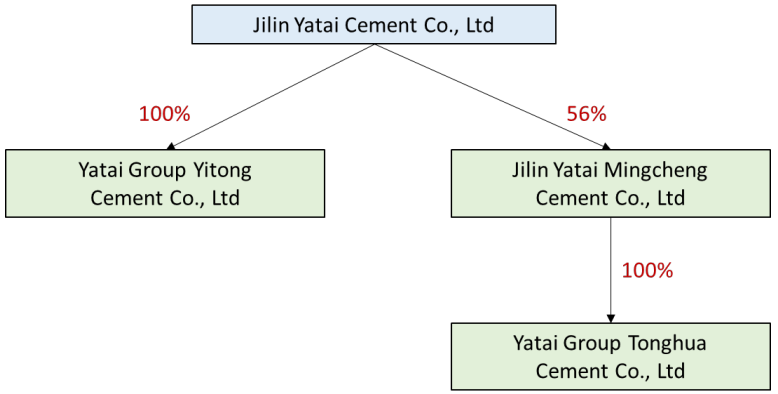


Figure 4: Examples of Firm Relations

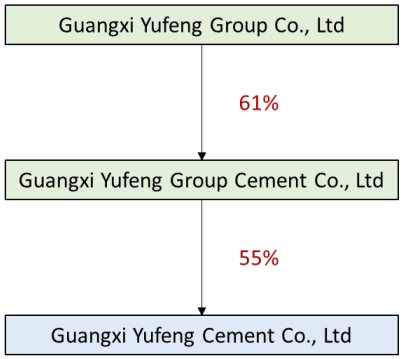
A. Affiliates



B. Investment



C. Shareholders



D. Shareholders' Investment

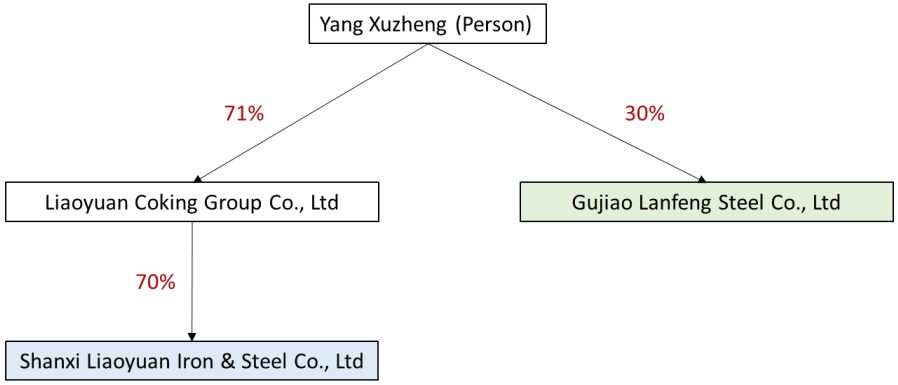
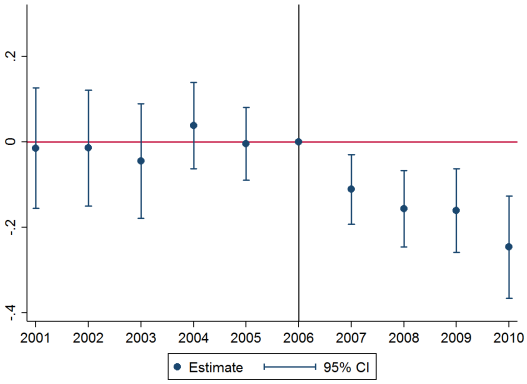
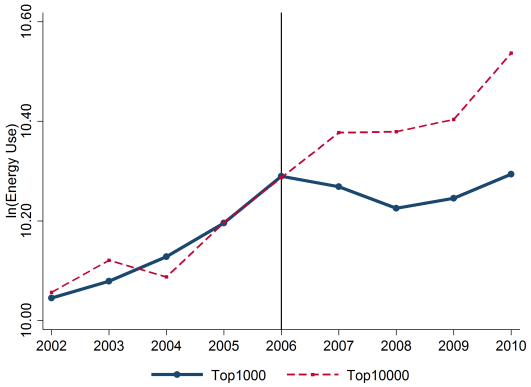


Figure 5: Effects of the Program on Regulated Firms

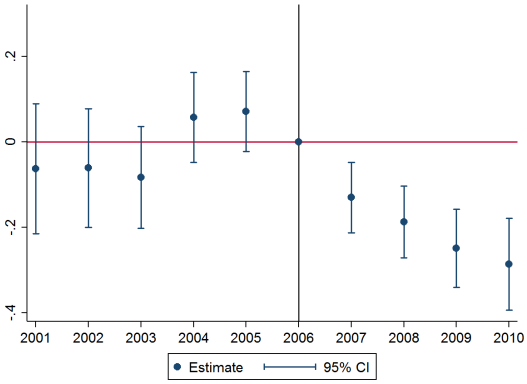
A. Energy Use: Coefficients



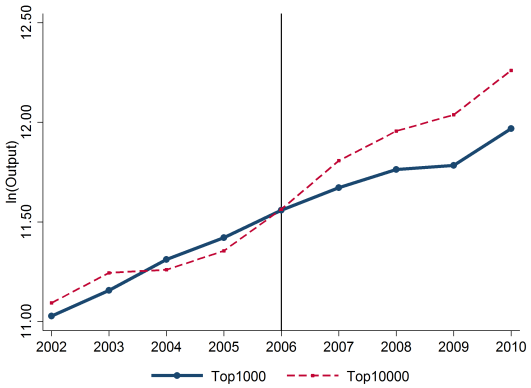
B. Energy Use: Event Study



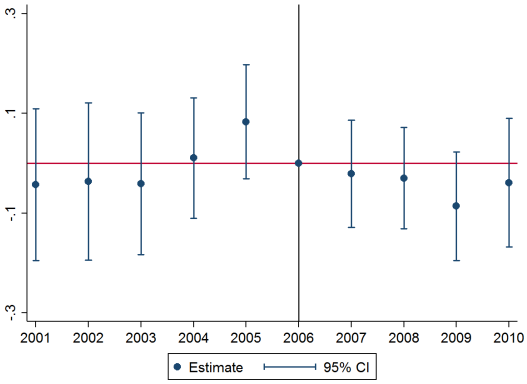
C. Output: Coefficients



D. Output: Event Study



E. Energy Efficiency: Coefficients



F. Energy Efficiency: Event Study

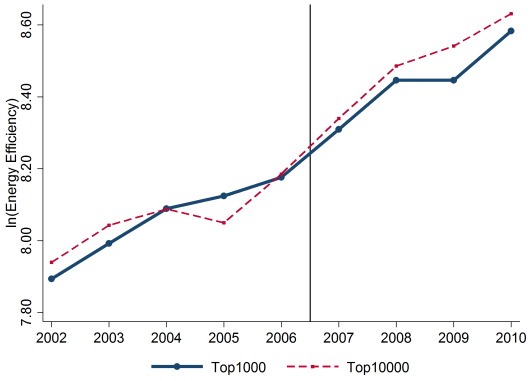


Figure 6: Spillover Effects on Related Firms

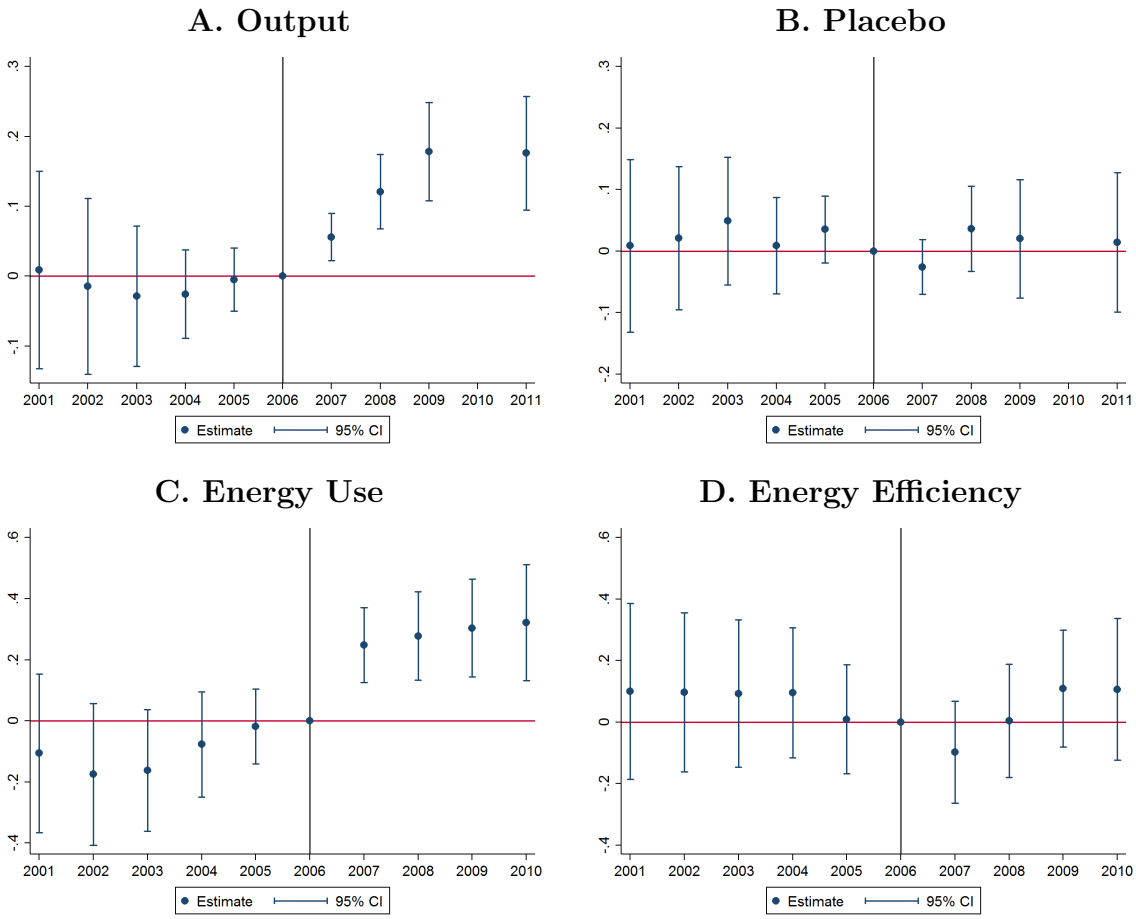


Figure 7: Industry-Level Spillovers

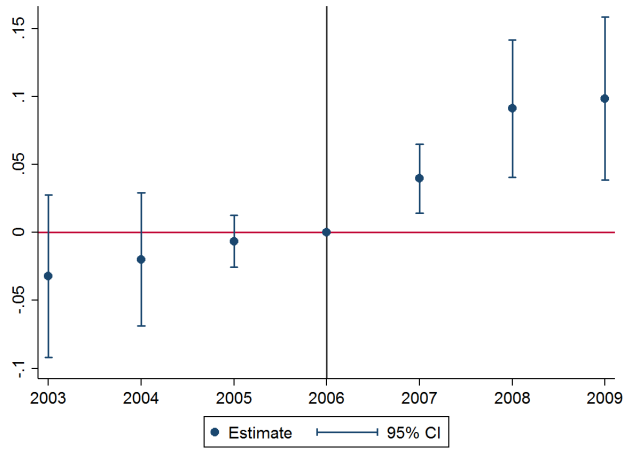
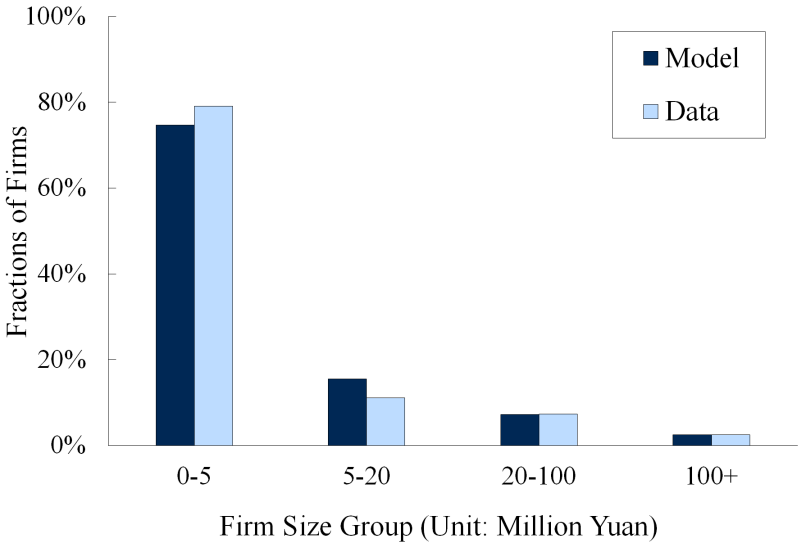


Figure 8: Distributional Moments and Structural Model Fit

A. Firm Size Distribution



B. Output Distribution

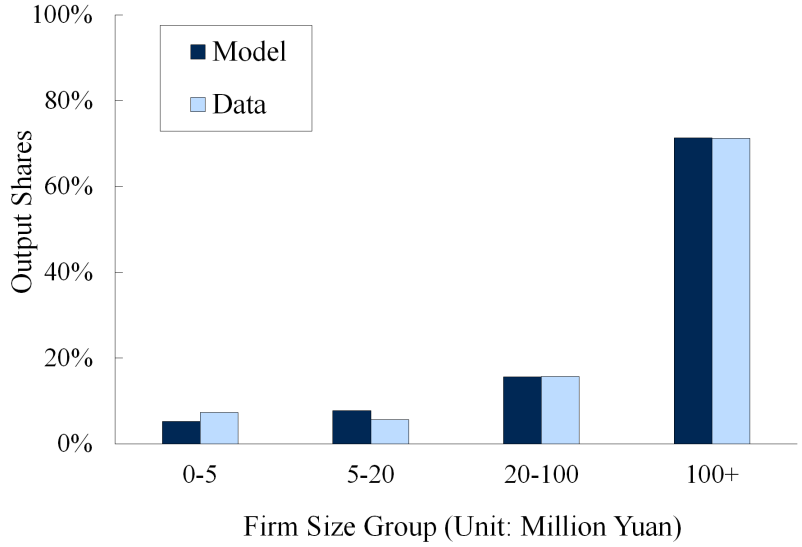
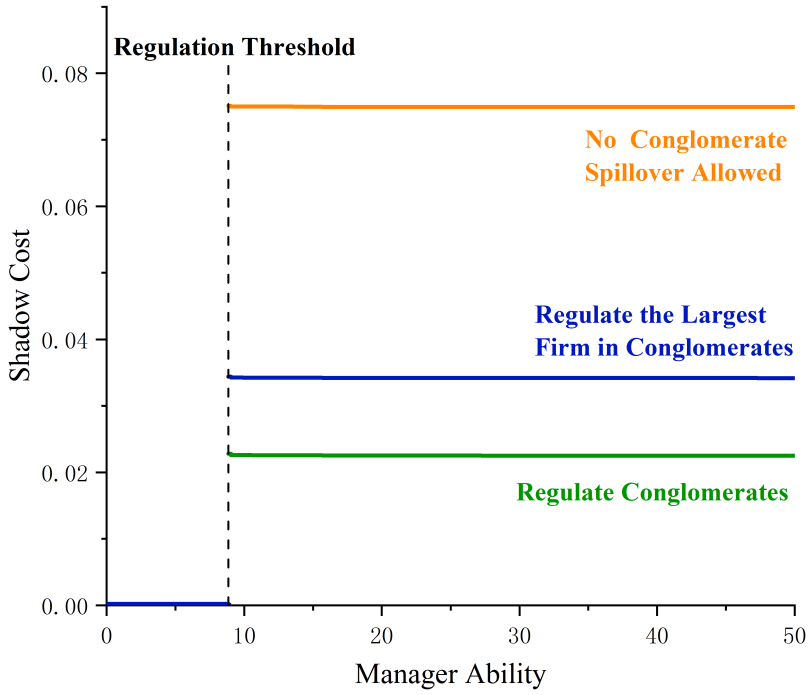
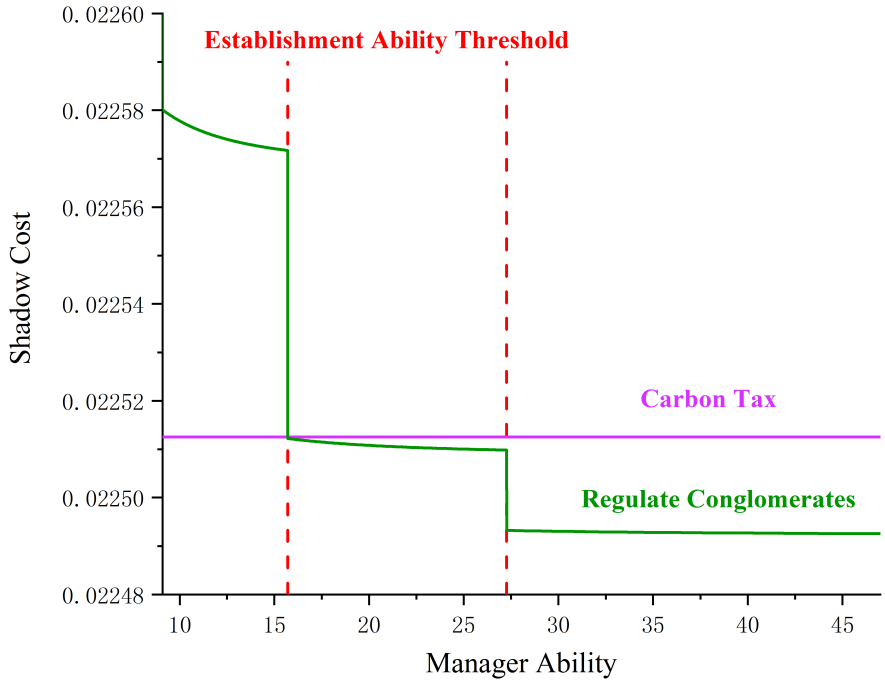


Figure 9: Model-Based Estimates of the Shadow Costs of Regulation

A. Shadow Costs of Alternative Regulations



B. Size Distortions and Equivalent Energy Tax



Tables

Table 1: Energy Consumption of Top1000 Firms in Different Industries

Industry	Energy Consumption $10,000$ ton coal equiv.	Proportion (%)	Firm Number
Iron and Steel	22528.63	30.72	249
Electric Power	16249.64	22.16	144
Chemical	10909.29	14.88	238
Petroleum and Petrochemical	10581.76	14.43	98
Mining	5278.77	7.20	60
Nonferrous	2993.08	4.08	70
Construction Materials	2913.19	3.97	93
Pulp and Paper	961.36	1.31	24
Textile	917.57	1.25	22

Table 2: Summary Statistics

A. Firm-Level Data

Source	Variables	Top1000			Top10000(Exclude Top1000)		
		Obs.	Mean	SD	Obs.	Mean	SD
ASIF	ln(Output)	8,745	14.15	1.58	87,247	12.24	1.61
	Size	8,725	14.44	1.54	87,338	12.18	1.74
	Soe	8,789	0.31	0.46	87,597	0.11	0.31
	Roa	8,557	0.05	0.10	85,402	0.08	0.16
	Age	8,775	23.04	20.32	87,449	12.63	13.82
	Export	8,789	0.34	0.47	87,597	0.25	0.43
	Variety	8,746	2.06	0.88	87,425	1.64	0.84
CESD	ln(Energy)	4,211	12.24	1.61	27,472	9.79	1.69
	ln(Coal)	3,904	12.33	1.55	25,607	9.89	1.69
	ln(Efficiency)	4,167	8.49	1.71	27,282	8.55	1.84
ASIF&ATS	Investment	4,276	0.80	0.40	47,243	0.82	0.39

B. Conglomerate Networks: Related Parties

Datasets	Six Levels	Two Levels		
	20%	20%	25%	51%
CARD	80,341	52,562	47,555	30,737
CARD&ASIF	10,944	8,447	7,811	5,358
CARD&ASIF (same 2-digit industry)	4,887	4,311	4,139	3,055
CARD&ASIF (same 4-digit industry)	2,776	2,561	2,500	2,024

Table 3: Effects of the Policy on Regulated Firms: Difference-in-Differences Results

A. Firm Energy Use

Variables	ln(Firm Energy Use)			
Treat×Post	-0.136***	-0.164***	-0.163***	-0.134***
	(0.042)	(0.045)	(0.045)	(0.047)
Observations	25,295	25,290	24,818	22,060
R^2	0.886	0.888	0.890	0.896
Firm FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Industry×Year FE		Y	Y	Y
Province×Year FE			Y	Y
Firm Control				Y

B. Output

Variables	ln(Output)			
Treat×Post	-0.108***	-0.231***	-0.210***	-0.152***
	(0.040)	(0.040)	(0.041)	(0.042)
Observations	25,113	25,108	24,648	21,930
R^2	0.879	0.884	0.886	0.891
Firm FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Industry×Year FE		Y	Y	Y
Province×Year FE			Y	Y
Firm Control				Y

C. Energy Efficiency

Variables	ln(Firm Energy Efficiency)			
Treat×Post	0.031	-0.066	-0.048	-0.019
	(0.041)	(0.044)	(0.045)	(0.046)
Observations	25,113	25,108	24,648	21,930
R^2	0.835	0.838	0.840	0.846
Firm FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Industry×Year FE		Y	Y	Y
Province×Year FE			Y	Y
Firm Control				Y

Table 4: Spillover Effects on the Output of Related Firms

A. Output

Variables	ln(Output)			
Related×Post	0.138***	0.133***	0.123***	0.133***
	(0.037)	(0.036)	(0.035)	(0.034)
Observations	19,228	19,226	19,226	18,668
R^2	0.893	0.900	0.906	0.914
Firm FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Industry×Year FE		Y	Y	Y
Province×Year FE			Y	Y
Firm Control				Y

B. Heterogeneous Effects on Output by Firm Size

Variables	ln(Output)			
Related×Post(0%-30%)	0.103*	0.104*	0.049	0.076
	(0.056)	(0.054)	(0.054)	(0.052)
Related×Post(30%-70%)	0.142***	0.139***	0.116***	0.136***
	(0.044)	(0.042)	(0.042)	(0.039)
Related×Post(70%-100%)	0.163***	0.151***	0.168***	0.164***
	(0.052)	(0.051)	(0.049)	(0.047)
Observations	18,174	18,174	18,174	17,666
R^2	0.896	0.903	0.909	0.917
Firm FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Industry×Year FE		Y	Y	Y
Province×Year FE			Y	Y
Firm Control				Y

C. Additional Firm Outcomes

Variables	ln(Sale)	ln(Profit)	ln(Asset)	ln(Fixed Asset)	ln(Employment)
Related×Post	0.120***	0.209***	0.071**	0.110***	0.037*
	(0.030)	(0.065)	(0.029)	(0.035)	(0.021)
Observations	18,200	13,787	18,444	18,308	16,511
R^2	0.929	0.779	0.944	0.932	0.928
Firm FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Industry×Year FE	Y	Y	Y	Y	Y
Province×Year FE	Y	Y	Y	Y	Y
Firm Control	Y	Y	Y	Y	Y

Table 5: Spillover Effects on Related Firms: Placebo Test, Energy Use, and Energy Efficiency

A. Placebo Test on Output

Variables	ln(Output)			
Related×Post	-0.010	-0.009	-0.006	0.004
	(0.041)	(0.040)	(0.040)	(0.039)
Observations	9,391	9,389	9,379	9,169
R^2	0.894	0.899	0.907	0.916
Firm FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Industry×Year FE		Y	Y	Y
Province×Year FE			Y	Y
Firm Control				Y

B. Energy Use

Variables	ln(Energy)			
Related×Post	0.356***	0.355***	0.343***	0.379***
	(0.075)	(0.074)	(0.076)	(0.094)
Observations	3,906	3,906	3,852	2,886
R^2	0.912	0.914	0.933	0.921
Firm FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Industry×Year FE		Y	Y	Y
Province×Year FE			Y	Y
Firm Control				Y

C. Energy Efficiency

Variables	ln(Energy Efficiency)			
Related×Post	-0.031	-0.030	-0.042	-0.008
	(0.079)	(0.078)	(0.080)	(0.099)
Observations	3,836	3,836	3,794	2,829
R^2	0.842	0.848	0.861	0.856
Firm FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Industry×Year FE		Y	Y	Y
Province×Year FE			Y	Y
Firm Control				Y

Table 6: Industry-Level Spillovers

Variables	ln(Output)			
	All Sample		Energy-intensive Industries	
Spillover	0.100***	0.067***	0.140***	0.082***
	(0.036)	(0.015)	(0.051)	(0.019)
Observations	2,026,061	1,979,314	680,826	663,995
R^2	0.837	0.856	0.846	0.865
Firm FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Firm Control		Y		Y

Table 7: Model Calibration and Fit

A. Calibrated Parameters

Parameter		Value	Target
1. Fixed Values			
Elasticity of substitution	$\sigma = \frac{1}{1-\rho}$	4.00	Melitz and Redding (2015)
Return to scale	α	0.90	Burnside et al. (1995)
2. Regression			
Fixed energy cost	f	3.7×10^4(Yuan)	Relationship between energy and output
Efficiency depreciation	$\delta^{\frac{1}{1-\alpha}}$	0.28	Relative firm size in a conglomerate
3. Simulated Model			
Dispersion of ln-ability ϕ	σ_m	1.18	Firm size distribution
Entry cost	f_e	8.8×10^6(Yuan)	
4. Match Policy Size and Impact			
Policy Threshold	$\tilde{\phi}$	8.84	Input share of regulated firms
Input Quota	$1-\xi$	0.13	Regulated firm output response

B. Model Fit

	Model Prediction	Empirical Result	95% CI
Output Response			
Regulated	-15.17%	-15.17%	(-23.34%,-7.00%)
Related	16.11%	13.29%	(6.68%,19.91%)
Other	3.54%	6.69%	(3.70%,9.68%)
Input Response			
Regulated	-16.52%	-13.36%	(-22.56%,-4.15%)
Related	17.92%	37.86%	(19.38%,56.33%)
Efficiency Response			
Regulated	1.53%	-1.87%	(-10.85%,7.12%)
Related	-1.46%	-0.77%	(-20.32%,18.78%)

Table 8: Effects of Simulated Counterfactuals

A. Regulated Conglomerate Response

	Short Run				Long Run
	No Spillover	Conglomerate Spillover	Market Spillover	Conglomerate & Market Spillover	
Output	-8.39%	-4.98%	-6.22%	-2.90%	-4.87%
Input	-9.24%	-5.44%	-6.82%	-3.10%	-5.32%
Efficiency	0.94%	0.49%	0.65%	0.20%	0.47%

B. Aggregate Impact

	Short Run				Long Run
	No Spillover	Conglomerate Spillover	Market Spillover	Conglomerate & Market Spillover	
Output	-6.78%	-4.02%	-3.54%	-1.66%	-0.09%
Input	-5.04%	-2.96%	-2.58%	-1.17%	0.00%
Efficiency	-1.83%	-1.09%	-0.98%	-0.50%	-0.09%
Mean Size	-8.05%	-4.78%	-5.66%	-2.64%	-4.69%
Medium Size	0.00%	0.00%	7.78%	3.54%	0.07%

Online Appendix: Not For Publication

This appendix contains multiple additional analyses.

Appendix Figures

Figure A.10: Map of Top 1000 Firms

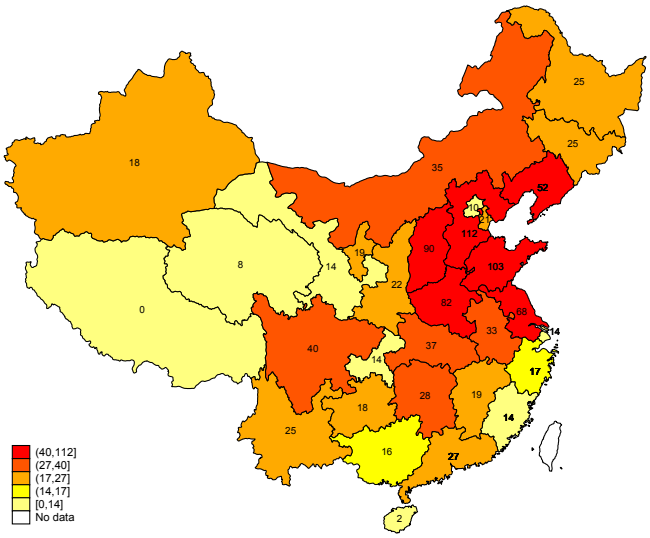


Figure A.11: Types of Related Parties

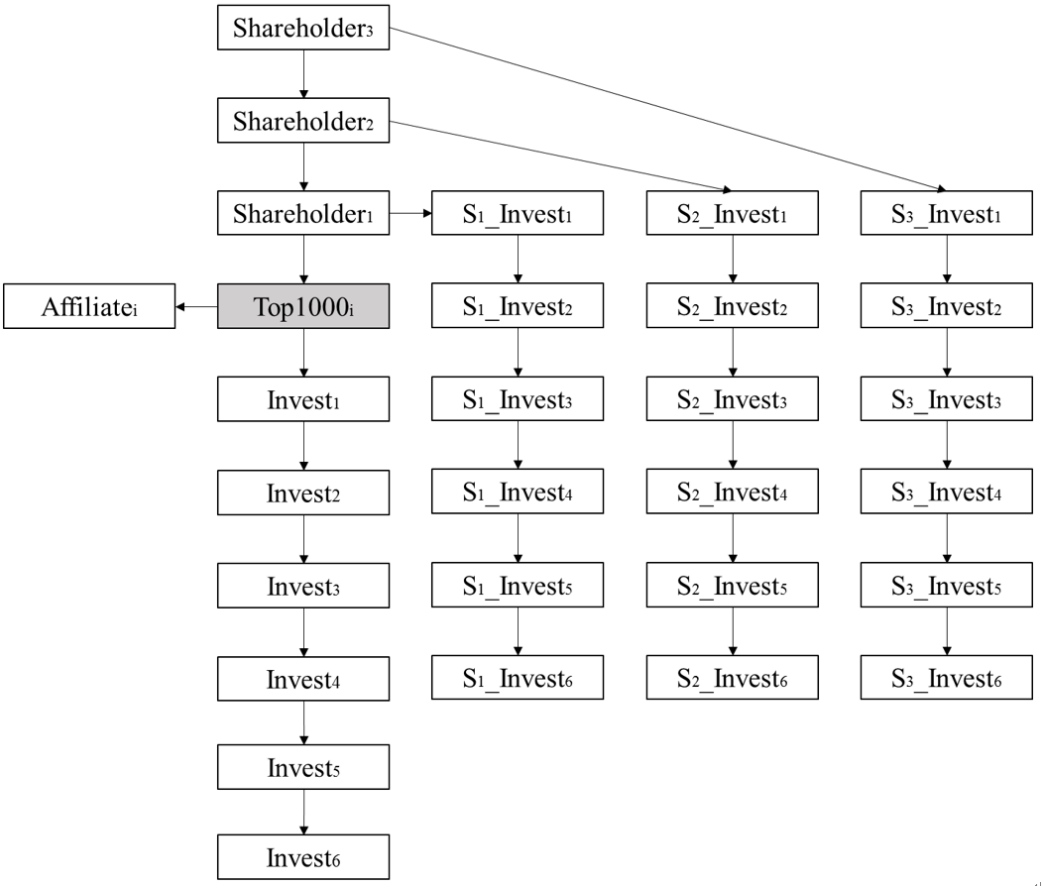


Figure A.12: Spillover Effects on Output by Firm Relation

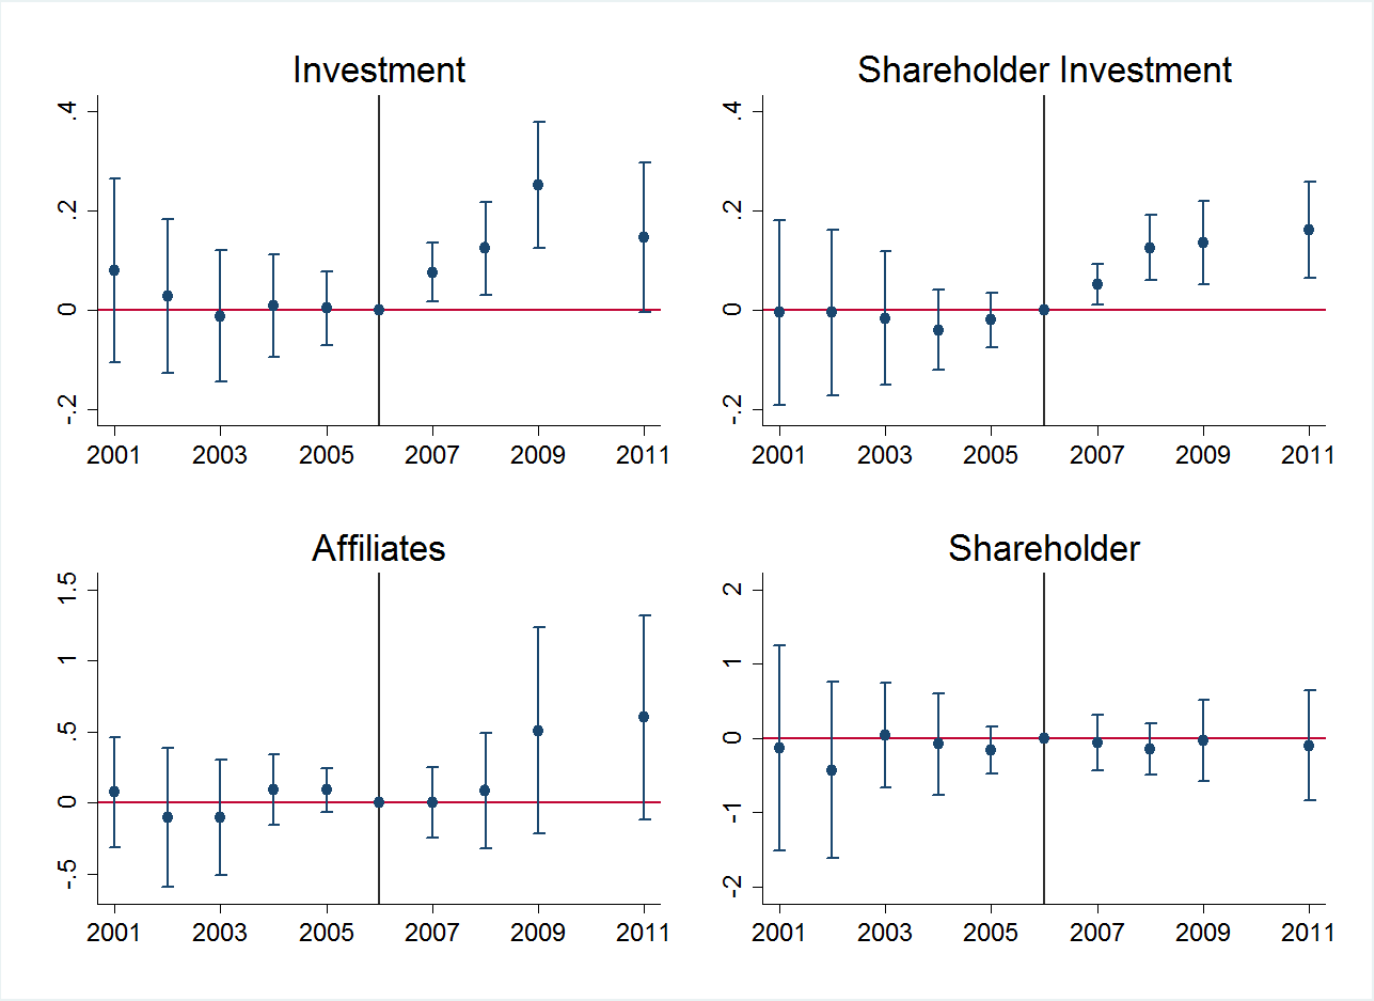


Figure A.13: Effects of Policy on the Investment of Regulated Firms

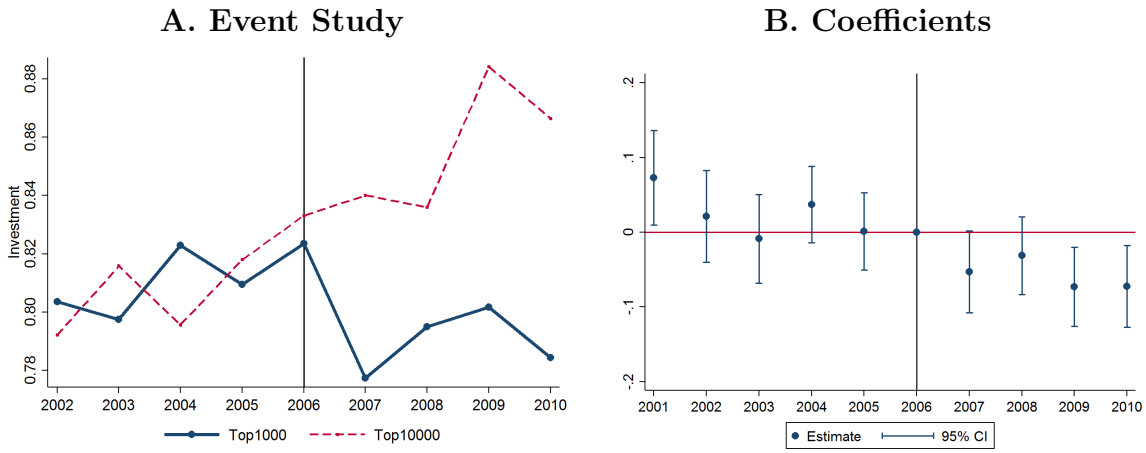
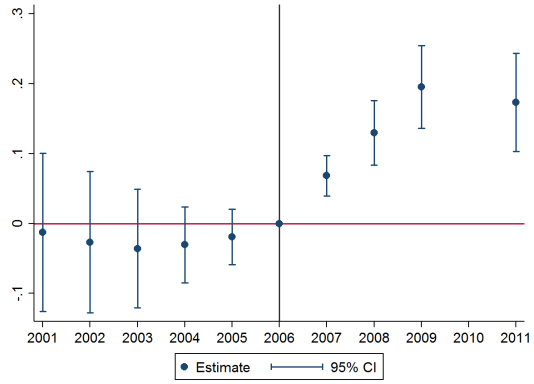
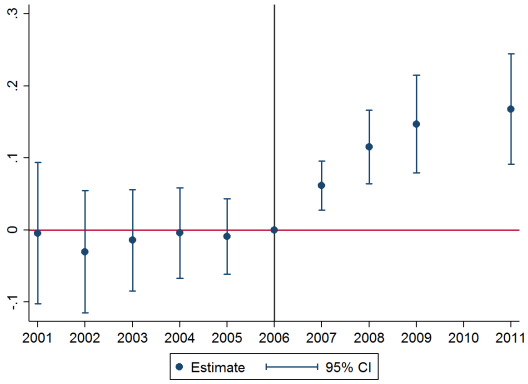


Figure A.14: Additional Spillover Effects of the Program

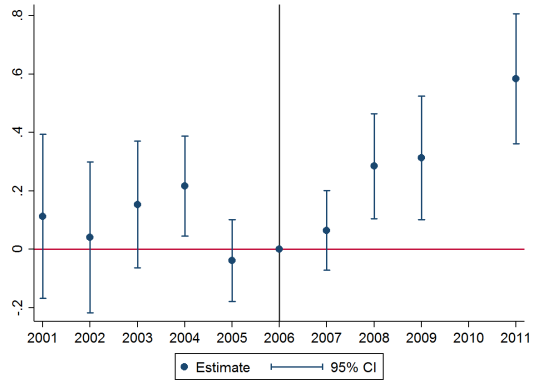
A. Output (Robustness to 3:1 Matching)



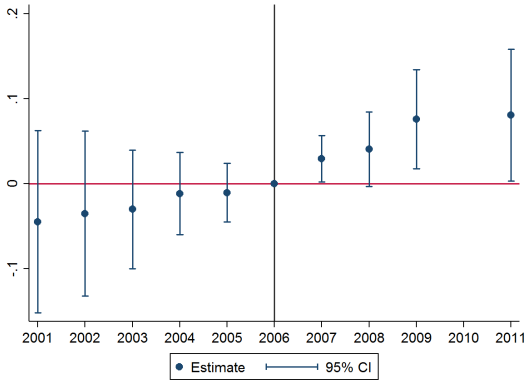
B. Sales



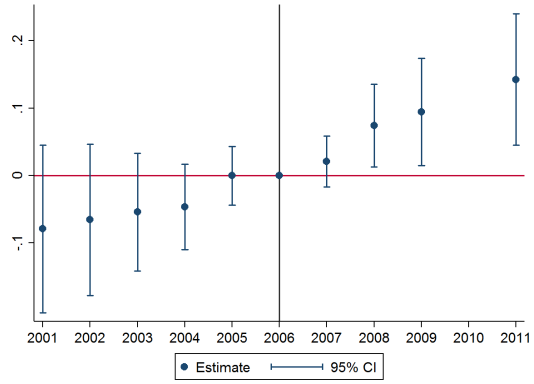
C. Profit



D. Assets



E. Fixed Assets



F. Employment

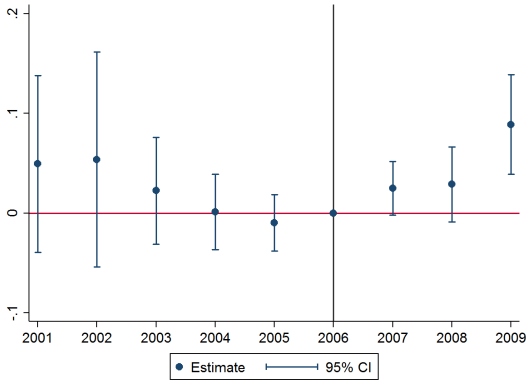


Figure A.15: Innovation of Top1000 Firms: Energy-saving Patent Application

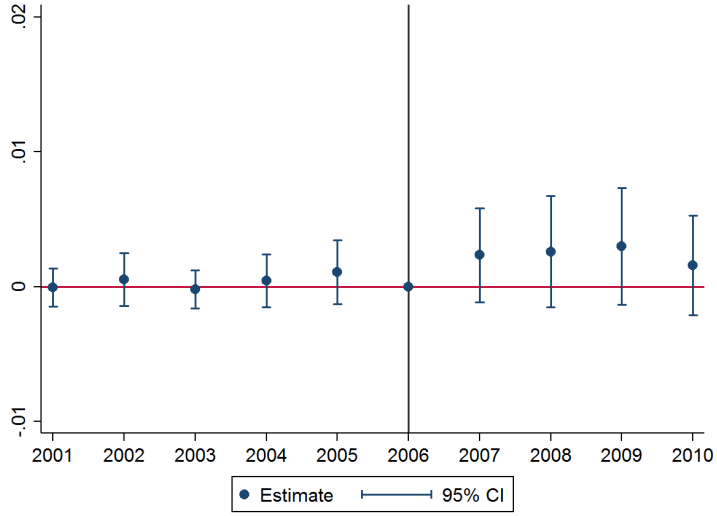


Figure A.16: Relationship between Firm Energy Consumption and Revenue

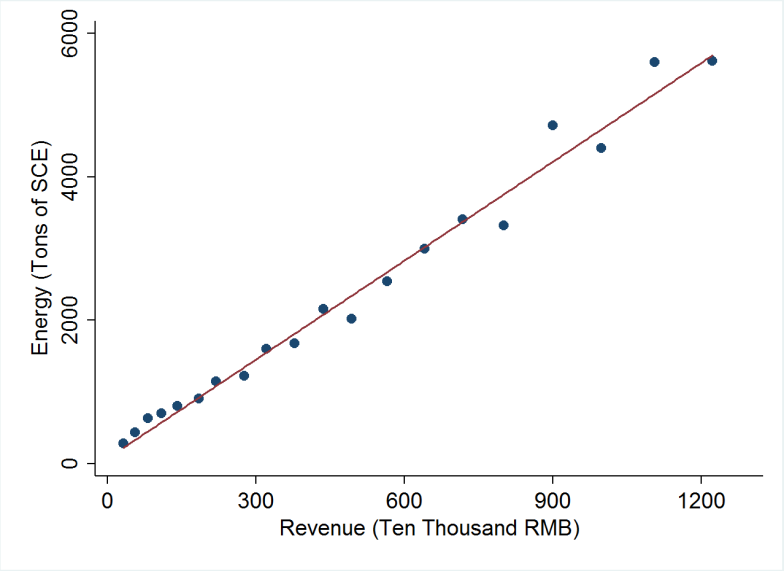
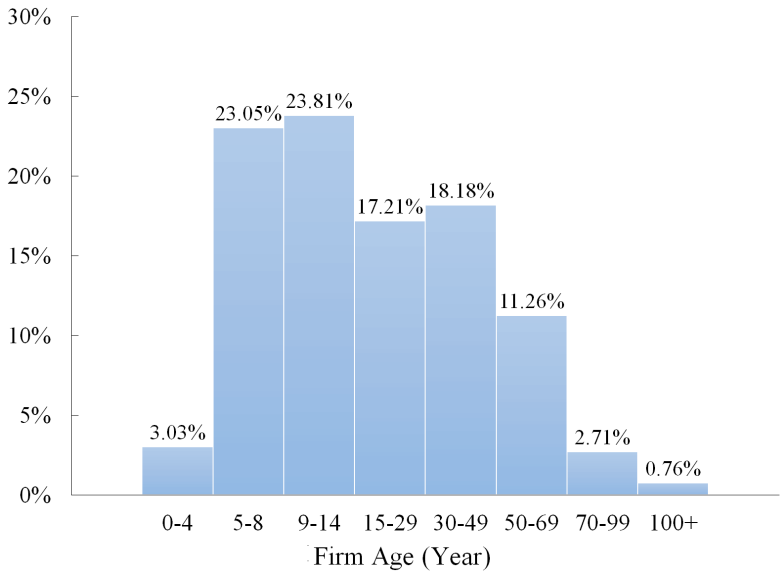


Figure A.17: Firm Age Distribution



Appendix Tables

Table A.9: Policy Compliance

Type	Orig.list	Evaluation			
Year		2007	2008	2009	2010
Firm Number	1008	953	922	901	881
Non-compliant firms	-	74	36	28	15
Non-compliant rate	-	7.76%	3.90%	3.11%	1.70%

Table A.10: Dataset Matching

Datasets	Top1000		Top10000	
	Number	Ratio	Number	Ratio
List	1008	-	14639	-
ASIF	1001	99.31%	14300	97.68%
CESD	818	81.15%	10723	73.25%
ASIF & CESD	809	80.26%	9482	64.77%
ASIF & ATS	447	44.35%	6623	45.24%

Table A.11: Determinants of Overfulfillment

Variables	Overfulfill			
	All Top1000 Firms		Completion Rate < 200%	
ln(related num.)	0.0467***	0.0635***	0.0732***	0.0909***
	(0.0148)	(0.0178)	(0.0215)	(0.0258)
ln(complement num.)		-0.0307*		-0.0354
		(0.0180)		(0.0287)
Observations	512	512	318	318
R^2	0.167	0.172	0.228	0.232
Firm Control	Y	Y	Y	Y
Industry FE	Y	Y	Y	Y
Province FE	Y	Y	Y	Y

Table A.12: Conglomerate Groups

Number of Top1000 Firm in Group	Number of Groups
1	341
2	44
3	26
4	5
5	3
6	4
7	2
9	1
10	1

Table A.13: Effects of Policy on the Investment of Regulated Firms

Variables	If Firm Invest			
Treat×Post	-0.057***	-0.071***	-0.072***	-0.071***
	(0.013)	(0.014)	(0.014)	(0.014)
Observations	50,987	50,967	50,967	49,346
R^2	0.191	0.200	0.208	0.212
Firm FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Industry×Year FE		Y	Y	Y
Province×Year FE			Y	Y
Firm Control				Y

Table A.14: Number of Establishments in a Conglomerate

Related Parties	Conglomerate Fraction	Establishment Fraction
0	53.73%	28.68%
1	22.04%	23.54%
2	13.28%	21.26%
3	6.80%	14.52%
4	2.85%	7.61%
5	0.97%	3.10%
6	0.27%	0.99%
7	0.06%	0.25%
8	0.01%	0.05%
9	0.00%	0.01%

Table A.15: Definition of Counterfactuals

	Short Run				
	No Spillover	Related Spillover	Market Spillover	Related&Market Spillover	Long Run
Regulated Establishment	$k_1 \leq \bar{k}_1$	$k_1 \leq \bar{k}_1$	$k_1 \leq \bar{k}_1$	$k_1 \leq \bar{k}_1$	$k_1 \leq \bar{k}_1$
Related Establishment	$k_i \leq k_i^*$	No belief Update	$\Delta q_r/q_r^* \leq \Delta q_o/q_o^*$	-	-
Other Establishment	$k_i \leq k_i^*$	$k_i \leq k_i^*$	-	-	-
New Entry	NO	NO	NO	NO	YES