

Agglomeration, Misallocation, and (the Lack of) Competition*

Wyatt J. Brooks[†]
Notre Dame

Joseph P. Kaboski[‡]
Notre Dame and NBER

Yao Amber Li[§]
HKUST

This version: January 22, 2020

Abstract

Industrial clusters and policy that promote them may limit the extent of competition. We develop, validate, and apply a novel approach for measuring a lack of competition and evaluating its aggregate consequences. The approach examines markups, which are higher but more equal and comove with combined market share under cooperation. Applying our approach to Chinese manufacturing firms, we find evidence for a lack of competition, especially in certain subsets, and cooperative pricing is almost three times as high inside Chinese special economic zones. Nevertheless, our quantitative estimates of the aggregate consequences of higher, but more equal, markups are positive.

*We are thankful for comments received from David Atkin, Timo Boppart, Hugo Hopenhayn, Oleg Itskhoki, Gaurav Khanna, Daniel Xu, participants in presentations at: Bank of Portugal Conference on Monetary Economics; BREAD; CEMFI; Cornell University; Dartmouth College; Duke University; the Federal Reserve Bank of Chicago; HKUST IEMS/World Bank Conference on Urbanization, Structural Change, and Employment; the London School of Economics/University College of London; London Macro Development Conference; NEUDC; Northwestern University; Princeton University; SED; SITE; the Taipei International Conference on Growth, Trade and Development; the University of Cambridge; the World Bank; and Yale University. We have also benefited considerably from comments of Ben Moll and three anonymous referees. We are thankful to the International Growth Centre and the HKUST IEMS (Grant No. IEMS16BM02) for financial support.

[†]Brooks: Department of Economics, University of Notre Dame. Email: wbrooks@nd.edu.

[‡]Kaboski: Corresponding author. Department of Economics, University of Notre Dame. Email: jkaboski@nd.edu. Research Associate at National Bureau of Economic Research.

[§]Li: Department of Economics and Faculty Associate of the Institute for Emerging Market Studies (IEMS), Hong Kong University of Science and Technology

1 Introduction

While industrial clusters, i.e., geographic concentrations of firms in the same industry, are explicitly promoted through policy and generally regarded as good for productivity, growth, and development, their impact on competition and its consequences for the economy have been generally understudied.¹ Concern that gathering competitors in the same locale and fostering cooperative interaction among firms could instead lead to non-competitive behavior dates back to at least Adam Smith.²

This paper starts with the hypothesis that geographic concentration and cluster policies are associated with non-competitive behavior and examines the prevalence and aggregate consequences of such behavior. We define non-competitive behavior as behavior in either firm sales, hiring, or input purchasing that internalizes pecuniary externalities on other firms. We make three major contributions toward this end. First, we derive a novel, intuitive screen for measuring the extent of internalization among firms competing in the same industry. Independent firms consider their own market share but not the market shares of other firms when setting markups. When firms cooperate by internalizing the impacts of their behavior on other firms, however, their markups tend to depend on the aggregate market share of the cooperating firms. They are therefore higher but more equal across firms. Second, using panel data on Chinese manufacturing firms, we validate our screen by confirming that affiliate plants of the same parent company are not behaving independently, which we would expect from firms with the same owner. Third, we show evidence of non-competitive behavior at the level of organized industrial clusters in the Chinese economy. Although we find limited levels of non-competitive behavior in the economy overall, it is almost three times as high in China's "special economic zones" (SEZs) than outside of them. Furthermore, we find that the levels of non-competitive behavior are also high in a set of industry-geography pairs that we pre-identified using the theory. Finally, we quantify the aggregate impact of this noncompetitive behavior, which – perhaps surprisingly – nets out to be positive. The tradeoff is between the cost of higher average

¹There are currently an estimated 1400 global initiatives fostering industrial clusters, and many studies find positive productivity benefits. ?, ?, and ? are examples of recent evidence. In contrast, ? finds little evidence of agglomeration economies in Detroit's Motor City, however.

²?’s famous quote: “People of the same trade seldom meet together, even for merriment and diversion, but the conversation ends in a conspiracy against the public, or in some contrivance to raise prices. It is impossible indeed to prevent such meetings, by any law which either could be executed, or would be consistent with liberty and justice. But though the law cannot hinder people of the same trade from sometimes assembling together, it ought to do nothing to facilitate such assemblies; much less to render them necessary. (Book I, Chapter X).” However, another strand of influential work, including Marshall, has also viewed industrial clusters as productivity-enhancing through the pro-competitive pressures they may foster (e.g., ?).

markups under firm cooperation and the gains from lower dispersion in markups. Ultimately, the benefits of reduced variation in markups that the macro misallocation literature has emphasized appear larger than the costs of higher distortions from market power emphasized by current antitrust policy.³ Indeed, abstracting from other potentially important considerations, a planner might want to have firms form a syndicate purely for the purpose of equalizing markups.

Why might proximity and frequent interaction lead to non-competitive price behavior? Close proximity and frequent interaction facilitate easy communication and observation that can enable cooperative behavior among firms, reducing the extent of price competition.⁴ There are certainly important cases of non-competitive behavior within industrial clusters. Historically, the most famous industrial clusters in the United States have all been accused of explicit collusion.⁵ In China, our empirical focus, our own interviews with firm owners and administrators of industrial clusters uncovered explicit cooperation on sales and pricing, as we discuss. Such smoking guns for particular cases exist, but what we lack is a sense of the overall prevalence of such non-competitive behavior in the economy, the extent to which it is linked to development policy, and the aggregate impacts of non-competitive behavior. Instead, we derive a screen to quantify the level of non-competitive behavior across an economy.

We derive our screen from a standard nested, constant-elasticity-of-substitution (CES) demand system with a finite number of competing firms and with a higher elasticity of substitution within an industry than across industries. As is well known in this setup and empirically confirmed (e.g., ?, ?), the gross markup that a firm charges is increasing in its own market share. Our theoretical contribution is to show that when a subset of firms internalize their impact on the profits of the other firms, it leads to convergence in markups across these firms, and each firm's markup depends on the total market share of the cooperating firms rather than its own firm-specific market share.

Following this logic, we regress the reciprocal of the firm's markup on the firm's own mar-

³Importantly, these calculations preclude any distortions in the input markets arising from market power. Our related work shows that these can be substantial (see ?).

⁴See, for example, ?, a theoretical case where easy observation helps support tacit collusion, or ? and ?, who document the behavior of actual cartels. Firm cooperation can also be beneficial, however. For example, firm associations have been shown to foster cooperation and information sharing, and increase the level of trust among managers while also increasing profits (?).

⁵See ? for evidence of Detroit's Big 3 automakers in the 1950s, and ? for Wall Street in the 1990s. The major Hollywood production studios were convicted of anti-competitive agreements in the theaters that they owned in the Paramount anti-trust case of the 1940s. Ongoing litigation alleges non-compete agreements for workers among Silicon Valley firms.

ket share and the total market share of its potential set of fellow syndicate members.⁶ From this we compute an index of the lack of competition as a simple function of the relative size of the coefficients on group market share vs. own market share. This screen is similar in spirit to the standard risk-sharing regression of ?, focusing on a syndicate of local (cooperating) firms rather than a syndicate of local (risk-sharing) households.⁷ It has similar strengths, in that it allows for the two extreme cases of independent decision-making and perfect joint maximization, but it also allows intermediate cases. As in Townsend, we can be somewhat agnostic about the actual details of how non-competitive behavior occurs; we instead focus on the outcomes, i.e., whether increased concentration among a set of firms (measured by market share) is associated with increased market power of those firms (measured by markups). The screen is also robust along other avenues. Importantly, our theoretical results, and so the validity of the screen, depend only on the constant elasticity demand system. They are therefore robust to arbitrary assumptions on the cost functions and geographical locations of the individual firms. Moreover, we use simulations to show that our screen performs well for plausible levels of firm uncertainty, including correlated demand or cost shocks, and when we relax our strong assumptions on the demand system. Indeed, simulations calibrated to our empirical exercise show only small biases when departures from our assumptions are in the empirically plausible range.

Empirically, we use the screen to assess the lack of independent competition in Chinese industrial clusters and SEZs. SEZs are generally considered key in China's growth miracle, and we have a high quality panel of firms with a great deal of spatial and industrial variation. The panel structure of the Annual Survey of Chinese Industrial Enterprises (CIE) allows us to estimate markups using the cost-minimization methods of ? and implement our screen using within-firm variation.

Our screen both identifies non-competitive pricing in simple validation exercises and rejects it in simple placebo tests. Specifically, we test for joint profit maximization among groups of affiliates with the same parent company and in the same industry. Similarly, we test for joint profit maxi-

⁶Throughout the paper we consider several different possibilities for sets of firms that are cooperating, such as firms with a common owner, firms in the same geographic region, and firms in the same special economic zone.

⁷An important difference between our context and that of ? is the potential confounding effect of measurement error. ? argue that idiosyncratic measurement error potentially biases the measure of risk-sharing upward when measurement error in the dependent variable is unrelated to that in the independent variable. In our context, measurement error in revenue affects both our measure of market shares and of markups. As discussed in Section ??, that implies that idiosyncratic measurement error in sales actually biases our measure of internalization *downward*. Hence, idiosyncratic measurement error cannot explain our results.

mization among state-owned firms in the same industry. Consistent with the theory, in our validation tests we estimate a highly significant relationship between markups and combined market share, but an insignificant relationship with the individual firms' own market share. This is exactly what the theory predicts for firms that maximize their joint profits. In our placebo tests, we find no response in markups to industrial cluster market shares and no influence of SEZs on markup behavior among these sets of firms.

In the broader sample of Chinese firms, the level of competitive behavior appears high, but as we move to smaller geographic definitions of a cluster the level of independent competition falls. Moreover, we find stronger evidence in subsets of clusters: SEZs and clusters pre-screened as having low initial cross-sectional variation in markups. SEZs target firms in specific industries and locations, giving them benefits such as special tax treatment or favorable regulation.⁸ They also attempt to foster cooperation through industry associations, trade fairs, and coordinated marketing, but such venues can be used to reduce competition. We find that the intensity of cooperative pricing is nearly three times higher for clusters in SEZs than for those not in SEZs. Moreover, we apply our pre-screening criteria, focusing on clusters in the lowest three deciles of cross-sectional markup variation, and find that only the cluster market share is a significant predictor of the panel variation in markups. That is, this subsample appears to be dominated by jointly cooperative, syndicate-like behavior. These clusters are characterized by disproportionately higher concentration industries, have lower export intensities, and contain a greater proportion of private domestic enterprises (as opposed to foreign or state-owned ventures).

Finally, to quantify the aggregate impact of this reduced price competition, we incorporate the same nested preferences, together with our estimated elasticities, into a general equilibrium framework with endogenous labor supply. Markups are wedges. Although higher markups distort labor supply, reduced variation in them allows for resources to be reallocated toward more efficient producers. For all reasonable values of the labor supply elasticity, the impacts of firms locally cooperating on pricing are positive on utility-based welfare and output. This holds even for extreme values (perfect local collusion) and even though we assume no within plant productivity benefits from agglomeration or cooperation. We note two important caveats, however. First, these estimates are valid for the current levels of geographic and industrial concentration in China, and for the

⁸We use SEZ in the broad sense of the term. See ? for a summary of SEZs, their history, and their policies.

estimated elasticities. Second, our calculations abstract from any distortions in the input markets arising from market power. Our related work shows that these can be substantial (?). Nevertheless, our results suggest that antitrust policy may underestimate the benefits that larger firms or cartels yield from more equalized markups.

Our paper contributes and complements the literatures on both industrial clusters and competition. First, we contribute to an emerging literature examining the role of firm competition – markups in particular – on macro development, including ?, ?, ?, and ?. ? study pro-competitive industrial policy in China. Similarly, a recent literature has looked at firm networks and firm cooperation and the productive benefits they may foster (e.g., ?, ?). In our companion paper, ?, we study monopsonistic cooperation among firms in China and India. Among this literature, this paper is unique in emphasizing the reduced variation in markups that can result from firm cooperation.

The local growth impact of Chinese SEZs has been studied in ?, ?, and ?, and they have been found to have sizable positive effects using panel level data at the local administrative units. Our firm-level evidence of non-competitive behavior suggests that the growth from these policies may at least partially reflect important, unintended consequences.⁹ Measured value added may be higher among firms within SEZs in part because cooperation allowed them to achieve higher markups, which is an important caveat when interpreting the previous results.

Finally, several papers have examined explicit collusion in cooperative industry associations, industrial clusters or agglomerations. The 19th century railroad associations in the U.S., originally formed to cooperate on technical (e.g., track width) and safety standards to link the various rails, soon turned to an explicit cartel designed to manage competition (see, e.g., ?). Colluding clusters in the 20th century have also been studied. ? studied collusion of the Big 3 automakers in Detroit, and ? examine NASDAQ collusion on Wall Street. More recently, ? shows that hotels near one another effectively collude.

Methodologically, the recent industrial organization literature tends toward “smoking gun” analysis of explicit collusion: detailed case studies of particular industries, making less stringent assumptions on demand or basing them on deep institutional knowledge of the industry.¹⁰ Our ap-

⁹While a lack of competition is likely an unintended consequence of agglomeration it is not obvious that the effect is negative. In a second best world, reduced competition may be welfare improving over high levels of competition. See, for example, ? or ? for the case where financial frictions are present. In this paper we do not need to take any stand on whether the welfare consequences of cooperation are negative or positive.

¹⁰? give an excellent review of the rationale for moving away from cross-industry identification. Our screen also

proach is different but complementary, developing a screen of effective competition and applying the entire economy of a developing country that has actively promoted industrial clusters. Thus, our screen can be used to guide broad industrial policy *ex ante* and considering competition more broadly, rather than focusing on a case study of a the extreme case of a cartel *ex post*.

2 Model

We develop a simple static model of a finite number of differentiated firms that yields different relationships between firm markups and market shares under independent competition and under syndicate behavior, and we show the robustness of these results to various assumptions. We assume a nested CES demand system of industries and varieties within the industry, which we assume is independent of location. Whereas the structure of demand is critical, we make minimal assumptions on the production side, allowing for a wide variety of determinants of firms costs, such as location choice, arbitrary productivity spillovers, and productivity growth for firms.¹¹

2.1 Firm Demand

A finite number of firms operate in an industry i . The demand function of firm n in industry i is:

$$y_{ni} = D_i \left(\frac{p_{ni}}{P_i} \right)^{-\sigma} \left(\frac{P_i}{P} \right)^{-\gamma}, \quad (1)$$

where p_{ni} is the firm's price, and P_i and P are the price indexes for industry i and the economy overall, respectively. Thus, $\sigma > 1$ is the own price elasticity of any variety within industry i , while $\gamma > 1$ is the elasticity of industry demand to changes in the relative price index of the industry.¹² Typically, $\sigma > \gamma$, so that products are more substitutable within industries than industries are with one another. The parameters D_i captures the overall demand at the industry level. For exposition, we define units so that demand is symmetric across firms in the same industry, but this is without loss of generality. As each firm in the industry faces symmetric demand, the industry price index

relies on within-industry (indeed, within-firm) identification.

¹¹Our assumption that demand is independent of location implicitly assumes negligible trade costs in output, which is important for allowing for agglomeration based on externalities rather than local demand. Empirically, we will focus on manufactured goods.

¹²We analyze highly disaggregated industries, so the assumption $\gamma > 1$ is natural.

within industry i is:

$$P_i = \left(\sum_{m \in \Omega_i} p_{mi}^{1-\sigma} \right)^{1/(1-\sigma)}, \quad (2)$$

where Ω_i is the set of all firms operating in industry i .

As we show in the online appendix, this demand system can be derived as the solution to a household's problem that has nested CES utility.

One can invert the demand function to get the following inverse demand:

$$p_{ni} = P \left(\frac{y_{ni}}{Y_i} \right)^{-1/\sigma} \left(\frac{Y_i}{D_i} \right)^{-1/\gamma}, \quad (3)$$

where:

$$Y_i = \left(\sum_{m \in \Omega_i} y_{mi}^{1-1/\sigma} \right)^{\frac{\sigma}{\sigma-1}}. \quad (4)$$

To establish notation that will be used throughout this paper, we define market shares as:

$$s_{ni} = \frac{p_{ni} y_{ni}}{\sum_{m \in \Omega_i} p_{mi} y_{mi}} = \frac{y_{ni}^{1-1/\sigma}}{\sum_{m \in \Omega_i} y_{mi}^{1-1/\sigma}}, \quad (5)$$

where the second equality follows from substituting in (??) for prices and simplifying.

This demand system implies that the cross-price elasticity is given by a simple expression:

$$\forall m \neq n, \frac{\partial \log(y_{in})}{\partial \log(p_{im})} = (\sigma - \gamma) s_{im}. \quad (6)$$

which allows for simple aggregation in the results that follow. Our structure of demand, which implies a this cross-price elasticity restriction and a constant elasticity of demand, allows us to be very general in our specification of firm costs. The cost to firm n of producing y_{ni} units of output is $C(y_{ni}; X_{ni})$, where X_{ni} represents a general vector of characteristics such as capital, technology, firm productivity, location, externalities operating through the production levels of other firms, and any other characteristics that are taken as given by the producer when making production choices. For example, a special case of our model would be one in which an initial stage involves a firm placement game in which each firms' productivity is determined by the placement of each other firm through external spillovers, local input prices, or other channels. Then the results from that

first stage determine X_{ni} that firms take as given when production choices are made, which is a special case of our framework.¹³

2.2 Imperfect Syndicate

We now consider the case of an imperfect syndicate, in which firms place a positive weight $\kappa \in [0, 1]$ on other firms' profits relative to its own, so that each firm maximizes:

$$\max_{y_{ni}} p_{ni}y_{ni} - C(y_{ni}; X_{ni}) + \kappa \sum_{m \in S/\{n\}} p_{mi}y_{ni} - C(y_{mi}; X_{ni})_{mi}.$$

We are agnostic about the precise reason that a firm might internalize the profits of other firms, since our concern is instead the behavior of markups and the misallocation they may cause.¹⁴ It is clear to see that the extreme case of $\kappa = 0$ the problem of a firm who independently maximizes its own profits, whereas the opposite extreme of $\kappa = 1$ captures a perfect syndicate: a subset of firms within an industry jointly maximizing the sum of their profits.

Using our definition of market shares again, we can express the first-order condition as:

$$\forall n \in S, \quad C'(y_{ni}; X_{ni}) = p_{ni} \frac{\sigma - 1}{\sigma} + p_{ni} \left(\frac{1}{\sigma} - \frac{1}{\gamma} \right) s_{ni} + p_{ni} \kappa \sum_{m \in S/\{n\}} \left(\frac{1}{\sigma} - \frac{1}{\gamma} \right) s_{mi}. \quad (7)$$

Then rearranging (??) gives the relationship between markups and market shares:

$$\frac{1}{\mu_{ni}} = \frac{\sigma - 1}{\sigma} + (1 - \kappa) \left(\frac{1}{\sigma} - \frac{1}{\gamma} \right) s_{ni} + \kappa \left(\frac{1}{\sigma} - \frac{1}{\gamma} \right) \sum_{m \in S} s_{mi}. \quad (8)$$

Examining the above equation, we can learn a lot from the extremes. In the extreme of firms operating independently ($\kappa = 0$), this equation implies that the only information that is needed to predict a firm's markup is that firm's market share. In particular, while factor prices, productivity, and local externalities captured by X_{ni} would certainly affect quantities, prices, costs, and profits, markups are only affected by X_{ni} through their impact on market shares. For $\sigma > \gamma$, the empirically

¹³However, note that the fact that firms maximize static profits below implicitly limits the way the vector X_{ni} can relate to past production decisions, such as dynamic learning-by-doing, sticky market shares, or dynamic contracts.

¹⁴In the case of actual cartels, ? document the importance of side payments among members. The ability to make such side payments could justify firms attempting to maximize total profits within the syndicate, i.e., the extreme case of $\kappa = 1$.

relevant case, additional sales that accompany lower markups come more from substitution within the industry than from growing the relative size of the industry itself. Firms with larger market shares have more to lose by lowering their prices, so they charge higher markups. On the other hand, in the perfect cartel extreme ($\kappa = 1$), the markup of a firm within the set S depends only on the total market share of all firms within the group. While the independent firm considered only its own market share, the syndicate internalizes the costs to its own members of any one firm selling more goods, and these cost depends on the total market shares of the member firms. In this extreme case of a perfect syndicate, the firm's own market share influences its markup only to the extent that it affects the syndicate's share. For intermediate values of κ , a firm partially internalizes the effects of its actions on the profits of other firms, and so its markup depends to both its own market share and the market share of its syndicate.

A number of corollary results follow. First, clearly $\sigma > \gamma > 1$ implies that an independent firm's markup is increasing in its own market share. Second, for a firm in a syndicate, the firm's markup is increasing in the total market share of the syndicate. That is, the firm's own market share plays no role except to the extent that it affects the syndicate market share. Third, syndicate members all charge the same markup, since their markup is based on the sum of their market shares. In our empirical work later we interpret this to mean that there is less variation in markups when firms behave cooperatively than they would have if they operated independently. Fourth, if any member of a syndicate were instead operating independently, that firm's markup would be lower and its market share would be higher. Finally, the market shares of any set of cooperating firms exhibit more variation than if the same set of firms was operating independently.

We summarize the above characterization in the following proposition.

Proposition 1. *Given $\sigma > \gamma > 1$:*

1. *When operating independently, firm markups are increasing in the firm's own market share.*
2. *When maximizing joint profits, firm markups are increasing in total syndicate market share, with the firm's own market share playing no additional role.*
3. *Syndicate firm markups are more similar under perfect syndicate than independent decisions.*
4. *Firm markups are higher under perfect syndicate decisions than independent decisions.*

5. *Firm market shares are less similar under perfect syndicate decisions than independent decisions.*

Each of these claims is addressed in our empirical results that follows. We will use the first two claims to derive our screen in Section ??, while the third and fourth claims will be used to pre-identify potential collusive clusters. Finally, we will use the fifth claim as additional testable implication. We have intentionally written Proposition 1 in general language. In the subsection below, we will show that, while the precise formulas vary, these more general claims are robust to several alternative specifications.

2.3 Alternative Models

We present related results below for the cases of firm-specific price elasticities, Bertrand competition rather than Cournot, a more general demand structure, and monopsonistic internalization.

2.3.1 Firm-specific price elasticities

To allow for markups to vary among competitive firms with the same market share, we allow for a firm-specific elasticity of demand. In particular, suppose that inverse demand takes the form:

$$p_{in} = D_i^{1/\gamma} P y_{in}^{-1/\sigma + \delta_{in}} Y_i^{1/\gamma - 1/\sigma}. \quad (9)$$

Here δ_{in} captures the firm-specific component of demand, and we think of these as deviations from the average elasticity, σ , i.e., $\sum_{n \in \Omega_i} \delta_{in} = 0$. Proceeding as before to derive markup equations, the first-order conditions of the firm imply:

$$\frac{1}{\mu_{ni}} = \delta_{ni} + \frac{\sigma - 1}{\sigma} + \left(\frac{1}{\gamma} - \frac{1}{\sigma} \right) s_{ni} + \kappa \left(\frac{1}{\gamma} - \frac{1}{\sigma} \right) \sum_{m \in S} s_{mi}. \quad (10)$$

Firm markups are again increasing in a combination of the firm and syndicate's market share, where the weight on the latter depends on the degree of internalization, and the magnitude of these relationships are governed by the difference between the within- and across-industry elasticities. In addition, however, the presence of δ_{ni} shows the level of markups may be firm-specific, even when

market share is arbitrarily small or firms are members of the same syndicate. This could explain why firms in the same syndicate have differing markups.

2.3.2 Bertrand competition

Now we consider the case where firms take competitors' prices as given instead of quantities when making production choices. From the demand function (??), we can write the problem of a firm operating independently as:

$$\begin{aligned} & \max_{\{p_{ni}, y_{ni}\}} p_{ni} y_{ni} - C(y_{ni}; X_{ni}) \\ \text{subject to: } & y_{ni} = D_i \left(\frac{p_{ni}}{P_i} \right)^{-\sigma} \left(\frac{P_i}{P} \right)^{-\gamma}. \end{aligned}$$

Taking first-order conditions with respect to both choice variables and dividing them yields the following equation, which is analogous to (??), respectively:

$$\frac{\mu_{in}}{\mu_{in} - 1} = \sigma - (\sigma - \gamma)s_{in} - \kappa(\sigma - \gamma) \sum_{m \in S} s_{im}. \quad (11)$$

In Equation (??), $\kappa = 0$ corresponds to the case where firms operate independently, and $\kappa = 1$ to the case where firms are in a perfect syndicate. Again, given elasticity parameters we see that firms' market shares are sufficient to solve for the firms' markups. As before, higher markups coincide with higher market shares, and the magnitude of this increasing relationship depends on the gap between the two elasticity parameters, and the weight on the syndicate's market share depends on the degree of profit internalization.

2.3.3 General Demand

Our result is not true for all demand systems, but it is useful to consider the extent to which it may hold for other demand systems, and what are the chief characteristics of demand driving this relationship. To examine this, we start with a very general demand system $p_{in}(y_{in}; y_{im})$. Denoting the inverse price elasticity $\frac{y_{im} \partial p_{in}}{p_{in} \partial y_{im}}$ as ε_{nm} , we can solve the Cournot problem to derive the following

general relationship for the perfect syndicate:

$$\frac{1}{\mu_{ni}} = 1 - \varepsilon_{nn} - \kappa \sum_{m \in S} \varepsilon_{nm} \frac{s_{mi}}{s_{ni}}. \quad (12)$$

In order for this to approximate equation the equation (??) above, we need to assume $\varepsilon_{nn} = \varepsilon_{1,nn}^* + \varepsilon_{2,nn}^* s_{ni}$ and $\varepsilon_{nm} = \varepsilon_{nm}^* s_{ni}$, where the starred elasticities are (approximately) constant. That leads to

$$\frac{1}{\mu_{ni}} = \frac{\varepsilon_{1,nn}^* - 1}{\varepsilon_{1,nn}^*} - \varepsilon_{2,nn}^* \left(s_{ni} + \kappa \sum_{m \in S} \frac{\varepsilon_{nm}^*}{\varepsilon_{2,nn}^*} s_{mi} \right). \quad (13)$$

In this expression, inverse markups involve a constant and an elasticity weighted sum of own and syndicate market shares.

Interpreting the above assumption, the inverse own price elasticity has both a component that is independent of market share and a component that increases in market share, while the inverse cross price elasticity is inversely related to market share. The components of these elasticities that are increasing in market share capture the idea that the impact on a price of a percentage output increase of a firm depends positively on the relative size of that firm in the market overall. The precise summation result depends on the inverse cross-price elasticities being equal to the second component of the inverse own price elasticity.

3 Empirical Approach

In this section, we present our empirical screen for non-competitive pricing, assess the robustness of the screen with simulations, and discuss our application to China, including the data and methods of acquiring markups.

3.1 Screen for Non-Competitive Pricing

The model of the previous section yielded the result that the markups of competitive firms depend on the within-industry elasticity of demand and their own market share, while the markups of fully internalizing firms depend on the total market share of the firms in the syndicate. This motivates the

following single empirical regression equation for inverse markups:

$$\frac{1}{\mu_{nit}} = \theta_t + \alpha_{ni} + \beta_1 s_{nit} + \beta_2 \sum_{m \in S} s_{mit} + \varepsilon_{nit} \quad (14)$$

for firm n , a member of (potential) syndicate S , in industry i at time t . While the relationships in equation (??) holds deterministically, the error term ε_{nit} could stem from (classical) measurement error in the estimation of markups (as discussed in Section ??), from uncertainty, or from other model specification error (as discussed in Section ??). We can easily estimate, κ using from equation (??):

$$\hat{\kappa} = \frac{\hat{\beta}_2}{\hat{\beta}_1 + \hat{\beta}_2}. \quad (15)$$

So $\hat{\kappa}$ is a measure of the intensity of internalization or cooperation.¹⁵

Moreover, we have clear null hypotheses that correspond to the model. In the case of purely independent optimization, the hypothesis is $\beta_2 = 0$ and $\beta_1 < 0$. For the case of a pure syndicate, we have the inverted hypothesis of $\beta_2 < 0$ and $\beta_1 = 0$.

Finally, equation (??) implies that we can use the regression in equation (??) to estimate the elasticity parameters. These equations imply that:

$$\frac{1}{\hat{\sigma}} - \frac{1}{\hat{\gamma}} = \hat{\beta}_1 + \hat{\beta}_2 \quad (16)$$

$$\frac{\hat{\sigma} - 1}{\hat{\sigma}} = \frac{1}{N} \sum_i \sum_{n \in \Omega_i} \left(\frac{1}{\mu_{ni}} - \hat{\beta}_1 s_{ni} - \hat{\beta}_2 \sum_{m \in S_{ni}} s_{mi} \right)$$

where N is the number of firms. It is then immediate to solve these equations simultaneously to generate estimates of the elasticity parameters.

Equation (??) has strong parallels with the risk-sharing test developed by ?. In that family of risk-sharing regressions, household consumption is regressed on household income and total (village) consumption in the risk-sharing syndicate. Townsend solves the problem of a syndicate of households jointly maximizing utility and perfectly risk-sharing, and contrasts that with households in financial autarky. We solve the problem of a syndicate of firms jointly maximizing profits in

¹⁵An alternative interpretation of $\hat{\kappa}$ is a measure of the fraction of firms that are cooperating. That is, if it is unknown ex ante which firms are cooperating and the sample pools together some firms that are perfectly cooperating and some that are not, then $\hat{\kappa}$ is a measure of the proportion that are cooperating. This interpretation is discussed in the online appendix.

a perfect syndicate and contrast with those independently maximizing profits. Townsend posited that households in proximity are likely to be able to more easily cooperate, defining villages as the appropriate risk-sharing network. We posit the same is true for firms and examine local cooperation of firms. Our screen also shares another key strength of risk-sharing tests: we do not need to be explicit about the details of how this cooperation arises.¹⁶ Instead, we directly address the effects of less competition that are of most concern: an ability of firms to use their collective market power to raise markups. Finally, as discussed in Section ??, firms could compete as in Cournot or Bertrand, and the essential elements of the screen hold in each.

We also note the presence of time and firm dummies in our screening equation. The time dummies, θ_t , capture time-specific variation, which is important since markups have increased over time, as we show in the next section. In principle, firm-specific fixed effects are not explicitly required in the case of symmetric demand elasticities.¹⁷ Nevertheless, we add α_{ni} to capture fixed firm-specific variation in the markup, stemming perhaps from firm-specific variation in demand elasticities, as discussed in Section ?. Together, these time and firm controls assure that the identification in the regression stems from within-cluster and within-firm variation over time in markups and market shares.¹⁸

3.2 Simulation Results of Robustness

We derived our screen from the model in Section ??, which assumed that (i) all relevant information is known to the firm before it makes its production or pricing decisions, (ii) demand is nested-CES, and (iii) there is no measurement error. In reality firms face unanticipated shocks to production costs and demand, and they take this uncertainty into account when making decisions. Indeed we require such unanticipated shocks in order to identify our production functions used in our empirical implementation. Moreover, demand may not be CES, and there may be measurement

¹⁶For example, we do not need to distinguish between implicit or explicit cooperation.

¹⁷Here the parallel with Townsend breaks, since risk-sharing regression require household fixed effects, or differencing, in order to account for household-specific Pareto weights. In contrast, syndicates maximize profits rather than Pareto-weighted utility, and as long as profits can be freely transferred – an assumption needed for a perfect syndicate – all profits are weighted equally.

¹⁸Firm-level fixed effects also strengthen the precision of our estimates by lowering the variation in the error term. Estimates without firm level fixed effects show a similar qualitative pattern, but the standard errors are larger and we lose some significance in our interaction of our variables with policy variables, i.e., dummies for SEZs. The main difference, however, is that the coefficient on own share, $\hat{\beta}_1$ is much larger. This would be consistent with larger firms facing less elastic demand curves. These results without firm fixed effects are available in the online appendix.

error with specific levels of correlation. Here we examine the robustness of our screen to relaxing these assumptions by running our regression on simulated data from an augmented model.

We augment demand and technologies for firm n in industry i located in region k in year t according to the following equations:

$$y_{nikt} = \varepsilon_{nikt} D_{nikt} \left(\frac{p_{nikt} + \bar{p}}{P_i} \right)^{-\sigma} \left(\frac{P_i}{P} \right)^{-\gamma}, \quad (17)$$

$$y_{nikt} = \rho_{nikt} z_{nikt} l_{nikt}^\eta$$

The parameter η allows for curvature in the cost function, while the parameter \bar{p} allows for demand elasticities that vary with prices.¹⁹ Here D_{nikt} and z_{nikt} are the known component of (firm-specific) demand and productivity, respectively, while ε_{nikt} and ρ_{nikt} are the unanticipated shocks to demand and productivity, respectively. Note that demand and productivity shocks are not equivalent in this model, since productivity shocks affect marginal cost, while demand shocks do not.

We then augment the firm's problem to allow for partial internalization captured by κ and take into account firm uncertainty:

$$\max_{l_{nikt}} \int_{\varepsilon} \int_{\rho} \left[(1 - \kappa) \pi_{nikt}(l, \varepsilon, \rho) + \kappa \sum_{m \in S_{ikt}} \pi_{mikt}(l, \varepsilon, \rho) \right] dF(\varepsilon) dG(\rho) \quad (18)$$

where the unsubscripted ε, ρ, l are *vectors* of demand shocks, cost shocks, and labor input choices. We assume that each firm belongs to a cluster S_{ikt} , and they jointly solve (??). In later sections we consider different cases for the sets of firms that may be operating as a syndicate, but in this section we refer to them generally as clusters. Notice that F and G are probability distributions over vectors. We will consider covariation of these shocks across firms at the firm, cluster, region-industry, industry, and year levels.

We simulate this model for various parameter values, run our screening regression on the simulated data, and evaluate the bias in κ as measured by equation (??). We overview the results here, and full details are given in the online appendix.

Our first exercise is to measure the bias to our estimates from unanticipated shocks. When shocks are at the level of the individual firm or are correlated at the level of the cluster, we find that

¹⁹In particular, own price elasticity is given by $\frac{d \log(y_{nikt})}{d \log(p_{nikt})} = -\sigma \frac{p_{nikt}}{p_{nikt} + \bar{p}}$, which is constant only if $\bar{p} = 0$.

they can bias our results. The direction of bias depends on the level of the shock. Unanticipated shocks at the individual level push our estimate of κ toward zero, while those at the cluster level push κ toward one. This is because individual shocks cause comovement in markups and individual shares independent of the cluster shares, which causes the coefficient on the individual share to increase in magnitude. The opposite is true for the cluster shock, which causes the coefficient on cluster share to increase in magnitude relative to that on the individual share. These effects can bias our $\hat{\beta}_1$ and $\hat{\beta}_2$ estimates. These estimates can lead to bias in $\hat{\kappa}$ for two reasons. First, biases in $\hat{\beta}_1$ and $\hat{\beta}_2$ feed directly into $\hat{\kappa}$. Second, since $\hat{\kappa}$ is a nonlinear function of $\hat{\beta}_1$ and $\hat{\beta}_2$, variance in the estimates of those coefficients leads to bias in $\hat{\kappa}$.

In all of these exercises, we stress that this bias only results from unanticipated shocks, even if they are correlated spatially or across industries. As discussed in Section ??, anticipated shocks cause no bias regardless of whether they are correlated spatially or across industries.

Indeed, anticipated variation actually lowers the relative importance of unanticipated shocks as we show in our second exercise. In this exercise, we study how large these unanticipated shocks would have to be to generate economically significant bias in our results. We parameterize the simulation to match the regression output from our baseline exercise, which is discussed in Section ?. We select the variance of individual shocks, the variance of cluster shocks as well as values of σ , κ and γ in order to match the point estimates and standard errors on the coefficients on own and cluster shares, the average markup, the estimated value of κ and the adjusted R^2 (when averaged across all simulations) to their counterparts in the Chinese analysis. We find that magnitudes of these shocks are not large enough to substantially bias our estimates of κ . In our parameterized simulation, the true value of κ is 0.29 while the estimated value is 0.28. In general, the quantitative importance of these depend on the magnitude of shocks relative to predictable variation in the data. Hence, large biases in estimates of κ would require extremely low R^2 , substantially less than even the small R^2 levels we observe in the data.

In our third exercise, we simulate a non-CES demand system. Applying the form of non-CES demand given in equation (?), we find that, as the CES-deviating parameter, \bar{p} , moves away from zero, our estimated coefficient on firms' own shares can be biased. In the case of $\bar{p} > 0$, the estimate would be downward biased, since a firm's markup would increase with its output (and firm's market share) simply from the decreasing elasticity. That is, this additional force would

increase the absolute magnitude of $\hat{\beta}_1$, decreasing $\hat{\kappa}$. The converse is true for $\bar{p} < 0$. Nevertheless, the coefficient estimate on the cluster shares are unbiased. This is important because, if we wished to screen for the presence of any cooperation, our model implies that we should test if the coefficient on the cluster share is positive. Thus, the fact that our coefficient on cluster share is unbiased with non-CES demand implies that our screen for the presence of cooperation is unaffected by non-CES demand. However, the fact that the coefficient on firms' own shares could be biased implies that our estimate of the *magnitude* of cooperation/internalization, $\hat{\kappa}$, is biased if demand were non-CES, and the direction of bias would depend on the direction of the deviation from CES demand.

Our final exercise is to consider measurement error in revenues and costs in the model to see how they affect our estimate of κ .²⁰ One might suspect that idiosyncratic measurement error would lead to overestimation of cooperation in a way that it can lead to overestimates of risk-sharing.²¹ However, we find that idiosyncratic measurement error actually leads to a *downward* bias in our estimate of κ .

This bias may seem surprising, but it has a simple explanation. Measurement error in regressors typically biases their coefficient estimates toward zero, so measurement error in a firm's own market shares alone should shrink that regressor's coefficient and push the estimate of κ toward one. However, this intuition relies on market share measurement error being independent of markups, but measurement error in revenue affects both measured market shares and measured markups. If the measured value of revenue is higher than its true value, both measured markups *and* measured market shares are by construction higher than their true values, and therefore idiosyncratic measurement error causes them to positively comove. We therefore *overestimate* the strength of the relationship between the two, increasing our estimate of $\hat{\beta}_1$ and causing a downward bias in $\hat{\kappa}$. Hence, if measurement error is idiosyncratic, we would tend to *underestimate* the extent of cooperation.²²

We find overall quantitatively small biases in our simulated estimates. We also perform several empirical robustness checks in Section ???. Finally, we find that our aggregate implications in Section ??? hold over a wide range of κ estimates. All of these exercises give us confidence that these

²⁰Measurement error is distinguished from the case of model misspecification described above in that unanticipated shocks are taken into account when firms make choices, while measurement error has no effect on firm choices.

²¹See, for example, ?'s critique of ?.

²²By the same argument, however, measurement error that is perfectly correlated at the level of the cluster biases the estimate of κ upward, and the overall bias for a mix of idiosyncratic and cluster-specific measurement error depends on the relative strength of each.

biases do not drive our results or conclusions.

4 Application to Chinese Data

For our empirical analysis, we examine manufacturing firms in China. Manufacturing firms have the advantage of being highly tradable, as is consistent with the assumption in our model that demand does not depend on location or local markets. Our measurement methods are standard and closely follow the existing literature.

4.1 Why China?

China has several advantages. First, it has the world's largest population and second largest economy, which provides wide industrial and geographic heterogeneity. Second, China is a well-known development miracle, and its success is often attributed, at least in part, to its policies fostering special economic zones and industrial clusters.²³ Third, both agglomeration and markups have increased over time as shown in Figure 1, which plots the average level of industrial agglomeration (as defined below) and average markups.

Finally, we have a high quality panel of firms for China: the Annual Survey of Chinese Industrial Enterprises (CIE), which was conducted by the National Bureau of Statistics of China (NBSC). The database covers all state-owned enterprises (SOEs), and non-state-owned enterprises with annual sales of at least 5 million RMB (about \$750,000 in 2008).²⁴ It contains the most comprehensive information on firms in China. These data have been previously used in many influential development studies (e.g., ?, ?).

4.2 Measurement

Between 1999 and 2009, the approximate number of firms covered in the NBSC database varied from 162,000 to 411,000. The number of firms increased over time, mainly because manufacturing

²³For example, a World Bank volume (?) cites industrial clusters as an “undoubtedly important engine [in China’s] meteoric economic rise.”

²⁴We drop firms with less than ten employees, and firms with incomplete data or unusual patterns/discrepancies (e.g., negative input usage). The omission of smaller firms precludes us from speaking to their behavior, but the impact on our proposed screen would only operate through our estimates of market share and should therefore be minimal.

firms in China have been growing rapidly, and over the sample period, more firms reached the threshold for inclusion in the survey. Since there is a great variation in the number of firms contained in the database, we used an unbalanced panel to conduct our empirical analysis.²⁵ This NBSC database contains 29 2-digit manufacturing industries and 425 4-digit industries.²⁶

The data also contain detailed information on revenue, fixed assets, labor, and, importantly, firm location at the province, city, and county location. Of the three designations, provinces are largest, and counties are smallest. We construct real capital stocks by deflating fixed assets using investment deflators from China’s National Bureau of Statistics and a 1998 base year. The “parent id code”, which we use to identify affiliated firms, is only available for the year 2004, but we assume that ownership is time invariant. We construct market shares using sales data and following the definition in Equation (??), where the firm’s sales are the numerator and the denominator depends on the industry classification. (In the analysis in Section ??, we change the sample of firms for various regressions, but the market shares for each firm remain the same: firm sales over the total industry sales in the full dataset.) We also use firms’ registered designation to distinguish state-owned enterprises (SOEs) from domestic private enterprises (DPEs), multinational firms (MNFs), and joint ventures (JVs).

We do not have direct measures of prices and marginal cost, so we cannot directly measure markups. Instead, we must estimate firm markups using structural assumptions and structural methods using method of ?, referred to as DLW hereafter. DLW extend ? to show that one can use the first-order condition for any input that is flexibly chosen to derive the firm-specific markup as the ratio of the factor’s output elasticities to its firm-specific factor payment shares:²⁷

$$\mu_{i,t} = \frac{\theta_{i,t}^v}{\alpha_{i,t}^x}. \quad (19)$$

²⁵The Chinese growth experience necessitates that we use the unbalanced panel. Using a balanced panel would require dropping the bulk of our firms (from 1,470,892 to 60,291 observations), or shortening the panel length substantially.

²⁶We use the adjusted 4-digit industrial classification from ?.

²⁷Specifically, consider a cost-minimization problem of a firm taking the price of factor x as given. The first-order condition with respect to factor x is:

$$p_x = \lambda \frac{\partial q}{\partial x},$$

where λ is marginal cost. Multiplying both sides by the output price p and rearranging to isolate the markup as p/λ , yields the DLW expression. Under the ? production function estimation procedure, unobserved variation in input prices at the firm level still leads to consistent estimates of production elasticities. Because this procedure does not use geographic information, then unobserved variation in input prices by geography (as considered in our main exercises) likewise does not challenge the consistency of our production elasticity estimates.

This structural approach has the advantage of yielding a plant-specific, rather than a product-specific, markup. The result follows from cost-minimization and holds for any flexibly chosen input where factor price equals the value of marginal product. Although the price must be flexibly chosen and price-taking from the point of view of the firm, it can be a firm-specific input price. Importantly, we use materials as the relevant flexibly chosen factor. The denominator $\alpha_{i,t}^x$ is therefore easily measured.

The more difficult aspect is calculating the firm-specific output elasticity with respect to materials, $\theta_{i,t}^v$, which requires estimating firm-specific production functions. The issue is that inputs are generally chosen endogenously to productivity (or profitability). We address this by applying ? (ACF)'s methodology, presuming a 3rd-order translog gross output production function in capital, labor, and materials that is:

$$\begin{aligned}
q_{nit} = & \beta_{k,i}k_{nit} + \beta_{l,i}l_{nit} + \beta_{m,i}m_{nit} + \\
& \beta_{k2,i}k_{nit}^2 + \beta_{l2,i}l_{nit}^2 + \beta_{m2,i}m_{nit}^2 + \beta_{kl,i}k_{nit}l_{nit} + \beta_{km,i}k_{nit}m_{nit} + \\
& \beta_{lm,i}l_{nit}m_{nit} + \beta_{k3,i}k_{nit}^3 + \dots + \omega_{nit} + \epsilon_{nit}. \quad (20)
\end{aligned}$$

Note that the coefficients vary across industry i , but only the level of productivity is firm-specific. This firm-specific productivity has two stochastic components. ϵ_{nit} is a shock that was unobserved/anticipated by the firm (and could reflect measurement error, as mentioned above) and is therefore exogenous to the firm's input choices. However, ω_{nit} is a component of TFP that is observed/anticipated, and is potentially correlated with $k_{i,t}$, l_{nit} , and m_{nit} because the inputs are chosen endogenously based on knowledge of the former. They assume that ω_{nit} is Markovian and linear in $\omega_{ni(t-1)}$. Identification comes from orthogonality moment conditions that stem from the timing of decisions, since lagged labor and materials and current capital (and their lags) are all decided before observing the innovation to the TFP shock. A two-step procedure is used to first estimate ϵ_{nit} and then the production function.²⁸

Production functions are estimated at the industry-level (although the estimation allows for firm-specific factor-neutral levels of productivity). The precision of the production function estimates –

²⁸In our companion paper, ?, we analyze multiple approaches for estimating markups and find that the results are largely robust to alternative methods to measure markups.

and hence the measurement error in markups – therefore depends on the number of firms in an industry. For this reason, we follow DLW and weight the data in our regressions using the total number of firms in the industry. Moreover, estimation of markups is noisy in practice, and within each industry we drop the 3 percent of observations in the tails.

We measure revenues rather than real prices, which can bias our estimates of markups but does not bias our estimate of interest, $\hat{\kappa}$. In particular, ACF’s methods assume that real quantities of output are measured rather than revenues. We follow previous work and deflate by an industry price index, but this does not fully put things into quantity terms because our output prices are firm-specific. Using Monte Carlos and our pricing model from Section ??, we have evaluated the impact of this on our estimated output elasticities, $\theta_{i,t}^v$, by modifying the code by ?. We find that markups themselves are problematic for ACF estimates when only revenues are available, leading to estimates production elasticities that are biased upward by the size of the markup. This in turn biases our markup estimates upward. This holds even in the case of uniform markups, when we set $\gamma = \sigma$, however, and varying the extent or existence of collusion (that is, changing κ or the fraction of firms that are cooperating) has no effect on this bias. Since the estimates are upward biased across the board, it affects our estimated intercept and coefficients by exactly this factor, but since κ is the ratio of coefficients, it leaves that estimate unchanged. Details are in the online appendix.²⁹

Finally, we use information on the geographic industries and clusters that we study. Namely, we merge our geographic and industry data together with detailed data from the China SEZs Approval Catalog (2006) on whether or not a firm’s address falls within the geographic boundaries of targeted SEZ policies, and, if so, when the SEZ started. We use the broad understanding of SEZs, including both the traditional SEZs but also the more local zones such as High-tech Industry Development Zones (HIDZ), Economic and Technological Development Zones (ETDZ), Bonded Zones (BZ), Export Processing Zones (EPZ), and Border Economic Cooperation Zones (BECZ). Since no SEZs were added after 2006, these data are complete. Since our data start in 1999, the broad, well-known SEZs that were established earlier offer us no time variation. We also measure agglomeration at

²⁹Moreover, we can estimate markups as simply sales over costs, which only requires an assumption of constant returns to scale. Using labor market monopsony, ? show that results are robust to various ways of measuring markups. For our estimates, sales show up directly in both the dependent and independent variables, so any measurement error will bias estimates. Instrumenting with lagged market share recovers qualitatively similar patterns with quantitatively plausible results, included in our online appendix, but we are not fully convinced by this instrumenting. In any case, our ACF estimates indicate diminishing returns to scale, even with their upward bias, another argument for using the DLW-based results.

the industry level using the g measure, where 0 indicates no geographic agglomeration (beyond that expected by industrial concentration), 1 is complete agglomeration, and a negative value would indicate “excess diffusion” relative to a random balls-and-bins approach.³⁰

Table 1 presents the relevant summary statistics for our sample of firms.

5 Direct Evidence of Cooperation in Chinese Industrial Clusters

This section details direct evidence on cooperation in Chinese industrial clusters, some of which may be productivity enhancing and some of which may reduce competition. We have direct evidence on the operation of industrial clusters and firm behavior from a small number of field visits to industrial clusters involving qualitative interviews with firm owners, government officials, and other support services in Chinese industrial clusters. Comparison with narrative reports from the field visits of other researchers indicate that the observed cluster behavior appears representative (g).

The clusters we visited were in different regions of the country and different industries. Each of the clusters focused on a unique consumer good industry with products involving a measure of standard automation but differentiated by quality, style and fashion rather than process technology.³¹ Each cluster involved production for both the domestic and export markets – typically each firm had some mix – but some clusters focus disproportionately on the domestic market, while others focus on the export market. Indeed, by government design, China has multiple industrial clusters in the same industry that are located in different regions. Some focus on the domestic markets, while

³⁰Specifically, start by defining a measure of geographic concentration, G :

$$G \equiv \sum_i (s_i - x_i)^2$$

where s_i is the share of industry employment in area i and x_i is the share of total manufacturing employment in area i . This therefore captures disproportionate concentration in industry i relative to total manufacturing. Using the Herfindahl index $H = \sum_{j=1}^N z_j^2$, where z_j is plant j 's share in total industry employment, we have the following formula for the agglomeration index g :

$$g \equiv \frac{G - (1 - \sum_i x_i^2) H}{(1 - \sum_i x_i^2)(1 - H)}$$

³¹Clusters tend to be highly specialized, at a finer level than our industry codes, such as cups, woolen sweaters, or hardware tools, for example.

others on the foreign market, thus partially segmenting the total market across clusters. These field visits uncover several avenues of firm cooperation, including government-firm relationships, industrial associations, coordinated marketing activities, and order sharing. The last reflects an explicit form of anti-competitive firm behavior.

Government cooperation is a common element of industrial clusters, and this government leadership can lead to coordination among firms. Many industrial clusters – though not all – have an official designation as a SEZ (or HIDZ, ETDZ, etc.). In some cases, these official designations and the policies associated with them were implemented at the foundation of the cluster, but typically they have been given to existing clusters to encourage their growth. Special economic zones assist in many ways, including streamlined export processing, preferential regulations, and tax benefits. Much of this is directed by local government officials.

Government cooperation also plays an important role in land markets and pollution permitting. In some clusters, the local officials allocate land within the special economic zone to certain firms. In another cluster we visited, the land was owned by a private developer, but the land was purchased by the real estate development company in conjunction with an influential member of parliament who assisted in getting proper regulatory access. In some polluting industries, pollution rights also come from local governments with the help of more influential government leaders at the national level.

Often, local governments organize business associations within SEZs that also foster cooperation. In the clusters we visited, the industry associations met weekly, biweekly, or monthly. The business leaders insisted that one of the key advantages of being in the clusters, in addition to access to specialized suppliers, was sharing information in order to have a pulse on market trends. They were able to differentiate their products from the competition (one way of segmenting the market), coordinate the mass of purchasers in the area (the scale of the market), as well as gather information about prevailing prices.

In many of our interviews, members of clusters discussed order sharing, which can take multiple forms.³² In some cases, a large firm receives a large order, then breaks up the order to be fulfilled by smaller firms. In another case, an industry association would coordinate the bids of its members to allocate orders among firms. Since member firms do not usually compete against one another, this

³²See ? for more discussion of order sharing among firms in industrial clusters in China.

eliminates competitive pressure among members of the association. As the president of an industry association explained, “We do not allow internal competition on pricing. If a firm tried price cutting, we would kick them out.” This president acted as a planner among the firms, allocating orders to member firms.

Other forms of cooperation within SEZs, such as information sharing, discussion of best practices, and entrepreneurship training, are consistent with previous studies showing positive productivity spillovers from firm-to-firm cooperation. In China, ? find that business associations, exogenously organized among medium-sized manufacturing firms, improved revenues and growth among firms by enhancing supplier-client matching and learning from peers. Similarly, ? find that exogenously introducing business owners with more experienced mentor-entrepreneurs in Kenya improved profitability of firms by helping young firms find low cost suppliers. Anecdotally, many entrepreneurs in our interviews reported similar effects. For example, entrepreneurs in one textile SEZ that we visited reported that membership in the SEZ has improved their business, which we can observe directly in our data. Firms inside SEZs enjoy, on average, higher labor productivity (value added per worker 15.4% higher), larger gross output value (6% larger) and sales (8% larger) relative to their counterpart firms in the same industry located outside of SEZs. The differences in means are highly statistically significant.³³ Therefore, there are other potentially important forms of cooperation among firms that are not captured by our screen, and may have positive effects on firm productivity.

Thus, we have direct evidence of both anti-competitive practices and productivity enhancing behaviors from firm cooperation. However, it is *a priori* unclear how quantitatively important these coordinated activities are, how representative these firm patterns are, and the extent to which the higher sales and revenue reflect the internalization of technological or pecuniary externalities. Nonetheless, the levels of cooperation in SEZs do not appear to approach levels of cooperation within large multiplant firms, at least along a few important dimensions. We found no evidence of cross-firm financing or investment coordination, for example. However, the normative implications of this cooperation are potentially important, which motivates more empirical work and aggregate analyses below.

³³To be clear, this is only a comparison of means, and we cannot claim this statement is causal.

6 Empirical Results

We start by presenting the results validating our screen using firms with common ownership. We then present the results for the overall sample (which are mixed), the results for those pre-identified clusters with low variation in markups across firms (which strongly indicate internalization), and some important characteristics of these collusive clusters. Throughout our regression analysis, we report robust standard errors, clustered at the firm level.³⁴

6.1 Validation and Placebo Exercises

We start by running our screen on the sample of affiliated firms. That is, we define our potential syndicates in equation (??) as groups of affiliated firms in the same industry who all have the same parent, and we construct the relevant market shares of these syndicates by summing across these affiliated firms. Note that although the sample is only a subsample of the full set of firms, market shares are the firms' (or syndicates) sales as a fraction of the total market (i.e., including the sales of firms not included in the regression). We know from existing empirical work (e.g., ?) that markups tend to be positively correlated with market share. Our hypothesis is $\beta_1 = 0$ and $\beta_2 < 0$, however, so that own market share will not impact markups after controlling for total market share of the syndicate firms. We estimate (??) for various definition of industries: 2-digit, 3-digit, and 4-digit industries. Note that the definition of industry affects not only the market share of the firm but the set of affiliates in the syndicate, S , and so the market share of the syndicate as well. A broader industry classification incorporates potential vertical cooperation, but it also makes market shares themselves likely less informative of a narrow horizontal market.

Table 2 presents the estimates, $\hat{\beta}_1$ and $\hat{\beta}_2$. (We omit the firm and time fixed effects from the tables.) The first column shows the estimates, where we assume perfectly independent behavior and constrain the coefficient on the internalized share to be zero. In the next three columns, we assume perfect internalization at the cluster level (constraining the coefficient on firm share to be zero), and define clusters at the 2-digit, 3-digit, and 4-digit levels, respectively. The last three

³⁴We cluster at the firm level, since the identification involves within-firm variation, and we can maintain the same clustering for all our analysis. The significance of our main results are robust to clustering at the "cluster" level as well, but such clustering varies from analysis to analysis, while clustering at the firm level allows us to remain consistent throughout, which allows for clearer comparison across results.

columns are analogous in their cluster definitions, but we do not constrain either coefficient. The sample of observations is a very small subset (less than two percent) of our full sample both because we only include affiliates, and because we only have parent/affiliate information for firms present in the 2004 subsample.

Focusing on the last three columns, we see that our hypothesis is confirmed for all three industry classifications with the coefficients on syndicate share being larger and statistically significant, while the coefficients on own share are smaller and not significant. The coefficients are larger for the broader classifications, implying very low elasticities of substitution between broadly defined markets. Since our model is one of horizontal competition, *a priori* we view the 4-digit classification as most appropriate. Applying (??) to the results that constrain $\hat{\beta}_1$ to zero (i.e., column (4)) yields estimates of $\sigma = 4.4$ and $\gamma = 2.9$. The corresponding values implied by column (7) are very similar at 4.4 and 3.1. At this 4-digit level, the implied demand elasticities in all of our results are consistent with those found using other methods, e.g., elasticities based on international trade patterns in ?, which is encouraging given the potential biases discussed in Section ??.

We also run placebo tests that define our syndicates S using all firms in the same region-industry pair (whether affiliated or not), construct syndicate market share by summing across all firms in the syndicate S (whether affiliated or not), but run the screen regression in equation (??) using only the subsample of affiliated firms. Relative to our placebo, which yielded positive cooperation results, the regression sample is the same, but the syndicate definition is changed: both the syndicate definition and syndicate market share values are identical to those used below in Section ??. Relative to our regressions below, which also yield positive cooperation results, the values and definition are the same, but the sample is different. The results are quite strong: we find no significant responses of markups to the syndicate share in the affiliated firm samples, and no effect of being in a SEZ (see Table B1 in the online appendix for full results). Recalling the Monte Carlo simulations in Section ??, a serious challenge to our identification would be correlated and unanticipated productivity or demand shocks that are especially strong locally or within an SEZ. However, these negative results are an important counter-example to the idea that spurious local correlations or something about the construction of our screen or our data automatically lead to false positives in detecting internalization at industry-region levels or SEZs.

In sum, our validation exercise is consistent with firms cooperation within ownership struc-

tures at the disaggregate industry level, and our screen is able to reject cluster-based cooperation in placebo tests.

6.2 Non-Competitive Behavior in Industrial Clusters

We now turn to industrial clusters more generally by defining our potential syndicates as sets of firms in the same industry and geographic location. Again, we change the set of firms included in the regression, and the definitions of a syndicate (i.e., the subset of firms over which we sum up market shares), but the market shares themselves continue to be defined as a fraction of the total market (total sales across all of Chinese producers in an industry). Table 4 presents the results. The first column shows the estimates, where we assume perfectly independent behavior and constrain the coefficient on syndicate share to be zero. In the next three columns, we assume perfect internalization at the cluster level, constraining the coefficient on firm share to be zero, and define clusters at the province, city, and county level, respectively. The next three columns allow for both shares to influence inverse markups, while the final three interact firm market share and cluster market share with an indicator variable for whether the firm is in a SEZ.

Focusing on columns 1 through 7, we note several strong results. First, all of the estimates are highly significant indicating that both firm share and market share are strongly related to markups. Because all estimates are statistically different from zero, we can rule out either perfectly independent behavior or perfect internalization at the cluster level. Second, all the coefficients on market shares are negative, as we would predict if output within an industry are more substitutable than output between industries. Third, the magnitudes are substantially larger for own firm share. Fourth, as we define clusters at a more local level, the coefficient on cluster share increases in magnitude, while the the coefficient on own share decreases. This suggests that cooperation is indeed more prevalent among firms that are in proximity to one another.

The $\beta_2 < 0$ estimates indicate some level of cluster-level collusion in the overall sample.³⁵ Again, applying equation (??), we can interpret the magnitude of the implied elasticities and the extent of internalization. We estimate $\hat{\kappa} = 0.29$ at the county level, while we estimate just $\hat{\kappa} = 0.08$ at the province level. This indicates a relatively low level of non-competitive behavior overall,

³⁵We verify that this is not driven by the affiliated firms in two ways: (i) dropping the affiliated firms from the sample, and (ii) assigning the parent group share within the cluster to firm share. Neither changes affect our results substantially.

especially when examining firms only located within the same province. The implied elasticity estimates are $\sigma = 4.8$ and $\gamma = 2.9$. These implied elasticities are quite similar to those implied in the smaller sample of affiliated firms, even though the level of internalization is greater.

Finally, we examine the role of SEZs examined in columns 8-10 of Table 4. The coefficients on the interaction of the SEZ dummy with firm market share are positive and significant but smaller in absolute value than the coefficient on firm market share itself. Adding the two coefficients, own market share is therefore a less important predictor of (inverse) markups in SEZs. Similarly, the coefficients on cluster market share are negative, so that overall cluster market share is a more important predictor in SEZs. Indeed, using the county-level estimates in the last column, we estimate an internalization index $\hat{\kappa} = 0.42$ for firms within SEZs, nearly three times as high as that of firms not in SEZs, where $\hat{\kappa} = 0.16$. Again, the results for SEZs are strongest, the more local the definition of clusters. Recall that SEZs are essentially pro-business zones, combining tax breaks, infrastructure investment, and government cooperation in order to attract investment. A common goal with industry-specific zones or clusters is to foster technical coordination in order to internalize productive externalities. The evidence suggests that such zones may also facilitate marketing coordination and internalizing pecuniary externalities.

We have estimated similar regressions where we differentiate across industries using the ? classification. Rauch classifies industries depending on whether they sell homogeneous goods (e.g., goods sold on exchanges), referenced priced goods, and differentiated goods. Without agriculture and raw materials, our sample of homogeneous goods is limited, but we can distinguish between industries that produce differentiated goods, and those that produce homogenous/reference priced goods. Our estimates of κ range from 0.15 to 0.28 for the former and range from 0.31 to 0.68 for the latter (depending on which Rauch specification is used as shown in Table B3 in the online appendix), indicating stronger cooperation among firms producing more homogeneous goods, consistent with existing arguments and evidence that collusion is less beneficial and common in industries with differentiated products (?). Equally interesting, the coefficients themselves are much larger for these goods, consistent with a larger ρ , which would be expected, since goods should be highly substitutable within these industries. Again, we view this latter consistency as further evidence that our results are driven by the markup-market share mechanism we highlight rather than some other statistical phenomenon.

We have also examined robustness of the (county-industry level, unrestricted) results to various alternative specifications in Table B4 in the online appendix. Although the theory motivates weighting our regressions, neither the significance nor magnitudes of our results are dependent on the weighting in our regressions. We can also use the Bertrand specification rather than Cournot, by replacing the dependent variable with $\mu_{nit}/(\mu_{nit} - 1)$. However, this Bertrand formulation requires us to Winsorize the data because for very low markups the dependent variable explodes. These observations take on huge weight, and very low markups are inconsistent with the model for reasonable values of γ . If we drop all observations below 1.06, a lower bound on markups for a conservative estimate of $\gamma = 10$ (much larger than implied by the Cournot estimates, for example), we get similar results, with implied elasticities $\sigma = 5.4$ and $\gamma = 2.3$ and intensity of internalization, $\kappa = 0.36$. Finally, we can use log markup, rather than inverse markup, as our dependent variable. The log function may make these regressions more robust to very large outlier markups. Naturally, the predicted signs are reversed, but they are both statistically significant, indicating partial internalization, and the implied semi-elasticities with respect to own and cluster share are 11.8 and 5.2 percent, respectively. The details of these robustness studies are in our online appendix.

We next turn to clusters which appear *a priori* likely to be potentially behaving as a syndicate because they have low cross-sectional variation in markups. We do this by sorting clusters into deciles according to their coefficient of variation of the markup. Table 5 presents the coefficient of variation of these deciles, along with other cluster decile characteristics, when clusters are defined at the county-industry level – the most local level, where we found the strongest evidence of cooperation in Table 4. Note that the average markup increases with coefficient of variation of markups over the top seven deciles, but that this pattern inverts for the lowest three deciles, where the average markup is higher as the coefficient of variation decreases. Higher markups and lower coefficients of variation may indicate cooperating behavior, given claims 3 and 4 in Proposition 1. We therefore focus on firms in the these bottom three clusters, and the lowest thirty percent is also consistent with the $\hat{\kappa}$ interpretation that 29 percent of firms collude.³⁶

The other key characteristics of these lowest deciles of clusters are also of interest. First, although they have lower variation in markups, this does not appear to be connected to lower vari-

³⁶These low markup variation deciles contain fewer firms on average, however, and so they constitute only 16 percent of firms.

ation in market shares, as the coefficients of variations in market shares are similar, showing no clear patterns across the deciles. They have fewer firms per cluster, and are in industries with higher geographic concentration (measured by the Ellison-Glaeser agglomeration index) and higher industry concentration (as measured by the Hirschman-Herfindahl index). The firms themselves are somewhat smaller in terms of fewer employees per firm. Fewer firms in these clusters export, and overall exports are a lower fraction of sales. Finally, although there are not sharp differences in the ownership distribution, they are disproportionately domestic private enterprises and somewhat less likely to be multi-national enterprises or joint ventures.³⁷

Table 6 presents the results for this restricted sample of the lower three deciles. The columns follow a parallel structure as in Table 4, but there are three columns even for the regressions that only include firm market share because the set of firms here varies depending on whether we define our clusters at the province, city, or county level. In the results that assume perfectly independent behavior we again find negative and significant estimates at the province and county level.³⁸ In the results that assume perfectly cooperative behavior, we again find negative significant estimates on cluster market share, and the results are again stronger, the more locally the cluster is defined. The most interesting results in the table, however, are those where we do not constrain either coefficient. In this restricted sample, we again find evidence of partially internalizing behavior at the province level.

What is striking, however, is that the internalization appears complete at local levels within these restricted samples: only the $\hat{\beta}_2$ estimates are negative and significant. The *positive* $\hat{\beta}_1$ at the city level is admittedly at odds with the theory, but the coefficient is not statistically significant. The county-level estimate in column (9) implies a within-industry elasticity σ that compares well with that in the full sample (5.0 vs. 4.8), but the between-industry elasticity is somewhat higher than in the full sample (6.6 vs. 2.9).

Once again, we find significant impacts of SEZs when interacted with market share. For counties, the region's share is nearly twice as large for firms in SEZs.

³⁷Moreover, the single most disproportionately overrepresented industry in these clusters is petroleum refining, a classic syndicate in U.S. history.

³⁸The city estimates have fewer observations, since there are fewer firms in the low markup variation deciles of city clusters.

6.3 Robustness

We now examine the robustness of our results to various alternatives. In particular, we attempt to address the issue that the correlation between markups and cluster share may simply be driven by spatially correlated shocks to costs or demand across firms. (Although our Monte Carlo simulations indicated this was unlikely to be problematic quantitatively.) We address this concern in two ways.

First, we add region-time specific fixed effects as controls into our regressions. Our Monte Carlo simulations showed that these effectively control for any general shocks or trends to production or costs at the region level, e.g., rising costs of land or (non-industry-specific) labor from agglomeration economies. Controlling for these, our regressions will only be identified by cross-industry variation in market shares within a geographic location. Table 7 shows these results for the sample of clusters with low initial variation in markups. The patterns are quite similar to those in Table 6, for example, the magnitudes of the coefficients on cluster share are -0.065 vs. -0.064 in column 9. The results are significant at a one percent level. We find very similar results for the overall sample, but since our SEZs show very little variation with counties, we cannot separately run our SEZ regression using these fixed effects. Nonetheless, we view the robustness of our results as evidence that spatially correlated shocks (or trends) do not drive our inference, although in principle, industry-specific spatially correlated shocks could still play a role.

Second, we attempt an instrumental variable approach, since shares themselves are endogenous. Identifying general instruments may be difficult, but in the context of the model and our ? estimation, exogenous productivity shocks affect costs and therefore exogenously drive both market share and markups. We motivate our instrument using an approximation, the case of known productivity z_{in} and monopolistic competition. This set up yields the following relationship between shares and the distribution of productivity:

$$s_{in} = \frac{p_{in}y_{in}}{\sum_{m \in \Omega_i} p_{im}y_{im}} \approx \frac{z_{in}^{1-1/\sigma}}{\sum_{m \in \Omega_i} z_{im}^{1-1/\sigma}} \quad (21)$$

We construct instruments for own market share (I_1) and cluster market share (I_2) using variants of the above formula that exclude the firm's own productivity and the productivities of all firms in the

firm's cluster (S_n), respectively:

$$I_1 = \frac{1}{\sum_{m \in \Omega_i/n} z_{im}^{1-1/\sigma}}, \quad I_2 = \frac{1}{\sum_{m \in \Omega_i/S_n} z_{im}^{1-1/\sigma}} \quad (22)$$

This two-stage estimation yields very similar results (see Table B5 in the online appendix). For example, the coefficient on cluster share in the analog to column (9) is -0.048 and is significant at the one percent level. Again, the patterns we develop are broadly robust.

In sum, we have shown for China that: the screen detects internalization among firms owned by the same parents in the affiliated and SOE samples; the markups of SOEs in a placebo test do not respond to their market share in a local cluster; the estimates are consistent with the model's mechanism based on the Rauch classification; our internalization patterns are stronger in SEZs; the internalization patterns are very strong in clusters that the model pre-identifies as likely syndicate clusters; these patterns are robust to inclusion of time-region specific fixed effects and instrumenting for market share.

7 Aggregate Implications

Having presented empirical evidence for local firm cooperation in pricing, we now study the potential quantitative implications of such cooperation. To do this, we embed the results from previous sections into a general equilibrium model and measure the change in aggregate output and household welfare from a change in the parameter κ . As discussed in greater detail below, the relationship between κ and aggregate output is not obvious. While less cooperation results in lower average markups, it also increases the dispersion of markups. Furthermore, when labor supply is endogenous, firm cooperation to restrict output also reduces labor demand. Therefore, the direction of the change is theoretically ambiguous.

7.1 Households

Consider a representative household that has GHH preferences over a final consumption good and labor. Their problem is:

$$\max_{\{c_t, L_t\}} \sum_{t=0}^{\infty} \rho^t \frac{\left(c_t - \frac{\phi}{1+\phi} \frac{L_t^{1+1/\phi}}{A} \right)^{1-\xi}}{1-\xi} \quad (23)$$

subject to:

$$\forall t, P_t c_t + P_t x_t = w_t L_t + r_t K_t + \Pi_t \quad (24)$$

$$\forall t, K_{t+1} = (1 - \delta) K_t + x_t$$

In most of what follows we suppress the time index to economize on notation. P is the aggregate price level, w is the wage rate, r is the rental rate of capital and Π is aggregate profit. Households own all firms, all capital and sell their labor L . Investment in capital x is in terms of the consumption good.

This preference specification allows us to manipulate the first order conditions to get a simple function for labor supply given by:

$$L = \left(A \frac{w}{P} \right)^\phi \quad (25)$$

The rental rate of capital is given by:

$$r_{t+1} = \frac{U_1(c_t, L_t)}{\rho U_1(c_{t+1}, L_{t+1})} - 1 + \delta \quad (26)$$

or, in a steady state:

$$r_t = \frac{1}{\rho} - 1 + \delta \quad (27)$$

The final consumption good is generated by aggregating the goods produced across I industries, and N_{jk} firms in each industry j and location k following the implied nested CES aggregator underlying the demand specification in ???. Furthermore, we set the aggregate price level as numeraire ($P = 1$). The demand for each firm j in industry i and location k is given by:

$$y_{ijk} = p_{ijk}^{-\sigma_j} P_j^{\sigma_j - \gamma} \quad (28)$$

where P_j is the industry j price index given by:

$$P_j = \left(\sum_k \sum_{l=1}^{N_{jk}} p_{ljk}^{1-\sigma_j} \right)^{\frac{1}{1-\sigma_j}} \quad (29)$$

7.2 Firms

The set of firms is fixed, and each firm i is characterized by its industry j , its location k and its productivity z_{ijk} . They produce using a Cobb-Douglas production function. Whatever their output choice y_{ijk} , they solve the following cost minimization problem:

$$C_{ijk}(y_{ijk}) = \min_{\{l_{ijk}, k_{ijk}\}} w l_{ijk} + r k_{ijk} \quad (30)$$

subject to:

$$y_{ijk} = z_{ijk} k_{ijk}^{\alpha_j} l_{ijk}^{\beta_j} \quad (31)$$

As in previous sections, the markup determination follows the earlier equation (??). We further assume that every firm in the data corresponds to an exact firm in the model that has the same industry and location. Given the firm's markup μ_{ijk} in the data, we choose its productivity z_{ijk} to match its empirical market share.

7.3 Counterfactual

The nature of the counterfactual is to modify the markup determination equations by changing κ , recompute the set of markups, then recompute prices and aggregates. That is, suppose that we change κ_j to κ'_j . Note that the exercise does not depend on the reason for the particular values of κ , only their impact on the distribution of markups. We solve for three potential measures for welfare: aggregate output, the wage (for a social welfare function that places weight only on the income of workers), and a consumption equivalent welfare measure (which accounts for changes in leisure as well).

Using the markup determination equation above, we can see that:

$$\frac{1}{\mu'_{ijk}} = \frac{1}{\mu_{ijk}} + \left(\frac{1}{\sigma_j} - \frac{1}{\gamma} \right) (s'_{ijk} - s_{ijk}) + \left(\frac{1}{\sigma_j} + \frac{1}{\gamma} \right) [\kappa'_j(\bar{s}'_{jk} - s'_{ijk}) - \kappa_j(\bar{s}_{jk} - s_{ijk})] \quad (32)$$

To solve for the new shares, s'_{ijk} , and new markups, μ'_{ijk} , this difference equation is combined with the relationship between shares and markups:

$$s'_{ijk} = \frac{p'_{ijk} y'_{ijk}}{\sum_m \sum_l p'_{ljm} y'_{ljm}} = \frac{(z_{ijk} / \mu'_{ijk})^{\frac{\sigma_j - 1}{\sigma_j - (\sigma_j - 1)(\alpha_j + \beta_j)}}}{\sum_m \sum_l (z_{ljm} / \mu'_{ljm})^{\frac{\sigma_j - 1}{\sigma_j - (\sigma_j - 1)(\alpha_j + \beta_{lj})}}}. \quad (33)$$

To compute the new values of aggregates, note that labor market clearing is given by:

$$(Aw)^\phi = \sum_j \left[\left(\frac{r}{\alpha_j} \right)^{\alpha_j(1-\sigma_j)} \left(\frac{w}{\beta_j} \right)^{\sigma_j - \alpha_j(1-\sigma_j)} P_j^{\sigma_j - \gamma} \right]^{\frac{1}{\sigma_j - (\sigma_j - 1)(\alpha_j + \beta_j)}} \sum_i \sum_k \left(\frac{z_{ijk}}{\mu_{ijk}} \right)^{\frac{\sigma_j - 1}{1 - \frac{\sigma_j - 1}{\sigma_j}(\alpha_j + \beta_j)}} \quad (34)$$

Given a rental rate r and a set of z_{ij} and μ_{ij} , it is straightforward to compute P_j and w . In particular, we compute w as the wage rate under the original markup determination equation, and w' as that after κ_i has been changed to κ'_i .

The change in total output is given by:

$$\frac{\sum_i \sum_j p'_{ij} y'_{ij}}{\sum_i \sum_j p_{ij} y_{ij}} = \frac{\sum_i (P'_i)^{1-\gamma}}{\sum_i (P_i)^{1-\gamma}} = \sum_i \left(\frac{P'_i}{P_i} \right)^{1-\gamma} \frac{P_i^{1-\gamma}}{\sum_k P_k^{1-\gamma}} \quad (35)$$

Finally, to quantify a consumption-equivalent welfare gain, we solve for η :

$$U(\eta c, L) = U(c', L') \quad (36)$$

where c is aggregate output and L is aggregate labor. Then:

$$\eta = \frac{c'}{c} - \frac{\phi}{1+\phi} \frac{L^{1+1/\phi}}{Ac} \left(1 - \left(\frac{L'}{L} \right)^{1+1/\phi} \right) = \frac{c'}{c} - \frac{\phi}{1+\phi} \frac{w^{1+\phi}}{A^\phi c} \left(1 - \left(\frac{w'}{w} \right)^{1+\phi} \right) \quad (37)$$

7.4 Results

We now use the formulas derived above to conduct several counterfactuals. In particular, we consider a variety of changes to κ_j , and different assumptions on the structure of the elasticities of substitution σ_j , and the production elasticities α_j and β_j . Moreover, we see how the results vary with the elasticity of labor supply ϕ . In each case, we normalize $A = 1$ while initial market shares, markups, industries and locations are exactly as in the data.

Table 8 demonstrates the case where κ_j and σ_j are constant across industries. In addition, the production elasticities are set to $\alpha_j = 0.25$ and $\beta_j = 0.5$ in each industry. We start with our benchmark estimate of $\kappa_j = 0.29$. This uniform case captures the effect of the change in κ in an economy where the only variation across firms is productivity, and location-industry grouping. We relax these assumptions sequentially in the next two sets of counterfactuals. The results vary with the labor supply elasticity ϕ . All cases are qualitatively similar, in that the higher κ , the higher are firm profits and aggregate output (as well as the welfare measure η), but the lower the wage.³⁹ That is, when all inter-firm cooperation is eliminated ($\kappa = 0$), aggregate income is lower than when firms cooperate. This is because decreasing κ causes markups to become more unequal, which decreases allocative efficiency across firms within the industry. The effect on wages captures the countervailing effect: reducing cooperation reduces the incentive to restrict output, which increases labor demand. This increases the total usage of factors. The labor supply elasticity ϕ plays a crucial role in determining the increase in labor usage. When $\phi = 0$, so that labor supply is exogenous, this effect is absent. Yet even when $\phi = 5$, which we think of as an implausibly high level, the change in factor usage does not dominate the changes due to markup equalization.

In Tables 9 and 10 we consider cases where κ_j is allowed to vary across industry. In these versions, we estimate our baseline model industry-by-industry instead of in the pooled data. For a given value of γ , the estimates allow us to identify σ_j and κ_j . In Table 10, we also allow for α_j and β_j to vary by industry and take values equal to those from production function estimation described in earlier sections. In these two cases, the results are in the same directions as in the first case, but now the magnitudes are greatly decreased.

³⁹As discussed in the online appendix, in these simulations we find that higher values of κ are associated with higher profits for firms at every income level. At least in this parameterization, this shows that cooperation could be valuable to all firms even without transfers.

The welfare magnitudes are not large, but they are also not tiny when compared to other policies, especially in environments without dynamic inefficiencies or externalities. Harberger-style estimates of the costs of monopoly pricing in the US are of the order of 1-4% (?). Estimated static gains from autarky to trade are of the order of 2-4% (?). The estimates of the costs of business cycles are of the order of 0.1 percent (?).

In any case, we interpret this quantitative exercise to demonstrate that, for empirically plausible values of parameters disciplined by the identification strategy proposed in this paper, cooperation does not reduce aggregate income. This is also robust to reasonable variations in our elasticities, given the inherent difficulties in measuring markups. Overall, the change in the consumption-based welfare measure, which also includes the value of leisure, is also positive. However, increases in cooperation among firms does have the effect of greatly increasing profits at the expense of reduced wages. Therefore, while cooperation does not reduce aggregate income, it does hurt real wages. In a world, where workers and owners of capital were different groups of people, this would lead to winners and losers.

8 Conclusion

We have measured the prevalence and consequences of non-competitive behavior in industrial clusters for the aggregate economy. Toward the first end, we developed a simple, intuitive and robust screen for identifying non-competitive behavior for subsets of firms competing in the same industry. Using this screen we have found evidence of a lack of competition in Chinese industrial clusters. These results are strongest within narrowly-defined clusters in terms of narrow industries and narrow geographic units. A small but non-negligible share of firms and clusters appear to exhibit non-competitive behavior. This behavior is disproportionately strong – nearly three times greater – in special economic zones. Our theory shows that such cooperation has ambiguous impacts, since it increases markups but also lowers the variation across markups in an industry. Our quantitative exercises imply that the latter dominates aggregate welfare calculations for Chinese manufacturing, which is perhaps surprising, since this beneficial aspect of price cooperation has not been emphasized in the literature.

The results open several avenues for future research. In this paper we have focused exclusively

on China. However, the fact that it satisfied our validation exercises means it could easily be applied more generally to other countries and contexts where firm panel data are available. Furthermore, the potential normative importance of our results is compelling with respect to evaluating industrial policies that promote clustering, such as local tax breaks, subsidized credit, or targeted infrastructure investments. They motivate more rigorous evaluation of various normative considerations, including productivity gains from external economies of scale vs. monopoly pricing losses from syndicates, local vs. global welfare implications and incentives, distributional consequences within societies, and the potential benefits and costs of allowing firms to merge, forming large firms rather than simply cooperative clusters. by current antitrust policy.⁴⁰

One final caveat is that our analysis has abstracted from any consequences of agglomerations and agglomeration policy that don't involve competition. So, for example, while our analysis was robust to local externalities, we did not attempt to measure any external economies or diseconomies. These may be quite important but are left for future work.

⁴⁰Importantly, these calculations preclude any distortions in the input markets arising from market power. Our related work shows that these can be substantial (see ?).