Market Structure, Oligopsony Power, and Productivity

Michael Rubens*

March 9, 2023

Abstract

I examine the effects of oligopsony power on allocative efficiency and income redistribution by studying a size regulation in the Chinese tobacco industry that led to ownership consolidation. I show that separate identification of input price markdowns, goods price markups, and productivity is challenging when a subset of inputs is non-substitutable, which often holds for materials, and construct and estimate a model to overcome this challenge. I find that the regulation increased input price markdowns by 37% on average. This increase in oligopsony power led to a decline in allocative efficiency and redistributed income away from rural households. (*JEL* L13, J42, O25)

A growing literature finds evidence for market power of firms on their input markets, such as labor or intermediate input markets.¹ Most of this literature focuses on the distributional effects of such 'buyer power', but much less is known about its effects on allocative inefficiency. Moreover, the identification of the level of buyer power and of its allocative effects is empirically challenging. Changes in input prices in response to changes in market structure, such as mergers, could be due to an increase of the 'markdown' of input prices below marginal revenue products, but could just as well be due to changes in these marginal revenue products, for instance because of scale economies or the exertion of market power downstream.

^{*}University of California, Los Angeles. E-mail: rubens@econ.ucla.edu.

This paper previously circulated under the title 'Ownership consolidation, monopsony power and efficiency: Evidence from the Chinese tobacco industry'. I am grateful to Jan De Loecker and Jo Van Biesebroeck for their advice. I also thank Frank Verboven, Chad Syverson, Otto Toivanen, Dan Ackerberg, John Asker, Myrto Kalouptsidi, Mingzhi Xu, and Benjamin Vatter for helpful suggestions, and participants at various conferences and workshops. I thank the editor and anonymous referees for comments, which significantly improved the paper. Financial support from the Research Foundation Flanders (Grant 11A4618N), the ERC Consolidator grant 816638 MPOWER, Fulbright Belgium and the Belgian American Educational Foundation is gratefully acknowledged.

¹I refer to Manning (2011) for an overview of the broader monopsony/oligopsony literature. Recent work on labor market power includes Naidu et al. (2016); Goolsbee and Syverson (2019); Jarosch et al. (2019), among many others.

In this paper, I fill this gap by empirically examining the effects of ownership consolidation on both input price markdowns and total factor productivity, and use these estimates to quantify the allocative and distributional consequences of oligopsony power on intermediate input markets. I study this question in the setting of the Chinese tobacco industry, which is particularly interesting because of two reasons. First, it underwent a large consolidation as a result of a government policy that forced manufacturers below certain production thresholds to exit the market after 2002. Second, a prohibition to transport tobacco leaf across local markets ensures that leaf markets are isolated, which allows to partition them into treatment and control groups. I compare the evolution of both input price markdowns and total factor productivity between these different leaf markets throughout the consolidation episode. Such quasi-experimental variation in market structure is rare, and very useful to study the drivers and consequences of buyer power. Despite these special characteristics, the structure of the Chinese tobacco industry, which features around 4 million farms who sell to around 150 manufacturers, who in turn sell to a monopsonistic wholesaler, is reminiscent of many other vertically structured industries in which there are concerns of buyer power. Prominent examples that have recently received the attention of competition policy-makers in the U.S. include book publishing and beef processing, among others.² Moreover, the Chinese tobacco industry is also interesting due to its sheer size: annual industry revenue exceeds \$7 billion, and 40% of the world's cigarettes are made in China.³

I start the empirical analysis of the paper by providing descriptive evidence for the effects of the manufacturers' consolidation on both input and product prices. I find that tobacco leaf prices at firms in consolidated markets fell by 50% compared to the other firms after 2002, whereas wages of factory workers did not change significantly. Factory-gate cigarette prices fell as well, by 31% on average. Although this evidence suggests the exertion of buyer power on tobacco leaf markets in response to the consolidation, other mechanisms could be at play. For instance, input and product prices could have changed due to changes in productive efficiency, or due to the exertion of market power downstream. In order to understand the underlying mechanism of these price changes, I construct a structural model to separately identify input prices, from total factor productivity. I build on the 'cost-side' approach to markup identification of Hall (1986) and De Loecker and Warzynski (2012), which has been extended to allow for endogenous input prices by De Loecker et al. (2016) and Morlacco (2017). I show that this class of models, which imposes only a model of production and input demand, fails to separately identify markups and markdowns as soon as a subset of inputs is non-substitutable. This is often the case for intermediate inputs in industries

²See this press release for book publishing, and this press release for U.S. meatpacking.

³Moreover, public health externalities are obviously an idiosyncratic aspect of the tobacco industry, although I will abstract from these in the paper.

such as beer (hop), consumer electronics (rare earth metals) and the earlier mentioned examples of meatpacking and book publishing. I show that this identification challenge can be overcome either by imposing more structure on competition and preferences upstream, or downstream. Rather than proposing a one-size-fits-all solution, I argue that this trade-off between assumptions needs to be tailored to the specific institutional setting of the industry at hand. In the context of Chinese tobacco manufacturing, I exploit the institutional feature of monopsonistic price-setting by the cigarette wholesaler, which permits assuming that manufacturers are price-takers on the wholesale market, to recover markdowns and productivity. I show that alternatively, the opposite approach of modeling input supply could also be taken, using a discrete choice model of leaf supply with differentiated manufacturers à la Berry (1994). This approach leads to similar findings, but imposes less assumptions about downstream demand and conduct.

I estimate this model by combining product- and firm-level production and cost data. The estimates reveal that cigarette manufacturers hold considerable buyer power over farmers: the median cigarette manufacturer paid its tobacco farmers merely half of the marginal revenue product of leaf. This high markdown level is consistent with the high degree of concentration on local leaf markets, and migration and mobility frictions faced by farmers. Using the estimated leaf price markdown and manufacturing productivity levels, I examine how the ownership consolidation affected both buyer power and productivity. I find that the consolidation policy led to the increased exertion of oligopsony power on leaf markets: markdowns in consolidated leaf markets increased on average by 37% compared to the control group between 2002 and 2006. Manufacturing productivity fell by 5% on average, although this drop was not significant. However, the main effect of oligopsony power on productivity is through input misallocation: aggregate productivity fell by 42% in provinces affected by the consolidation compared to unaffected provinces. This demonstrates that oligopsony power can be an important driver of resource misallocation. Finally, I use the model to quantify the extent to which the consolidation contributed to rural-urban income inequality in the tobacco industry. This margin of inequality has increased rapidly in China since the early 1990s (Yang, 1999; Ravallion and Chen, 2009). By increasing markdowns on tobacco leaf markets, but not on manufacturing labor markets, the consolidation accounts for 56% of the increase in income inequality between rural farmers and urban manufacturing workers between 2002 and 2006.

These results have two important policy implications. First, the finding that oligopsony power can lead to substantial input misallocation implies that competition policy-makers should take into account upstream competition effects in their merger and antitrust policies, even under the consumer welfare standard. Second, I shed new light on the consequences of large-scale industry consolidation programs, which are increasingly common forms of industrial policy in countries such as China and Indonesia. China recently consolidated, for instance, many of its state-owned enterprises (SOEs) into industrial giants in various important industries such as energy, transport

utilities, telecommunication and military equipment. These policies are also known as "*Grasping the large and letting the small go*" (Naughton, 2007). I show that such policies, which are usually motivated as a means to spur productivity growth, can actually lead to decreased aggregate productivity growth due to the exertion of oligopsony power. This is especially a concern for agricultural markets in developing countries, which often feature internal trade restrictions similar to those in the Chinese tobacco industry (Chatterjee, 2023).

This paper makes three contributions to the literature. First, I examine how ownership consolidation affects both input price markdowns and productivity. Prior work on the effects of ownership consolidation on productivity (Braguinsky et al., 2015; Grieco et al., 2017), on downstream market power (Nevo, 2001; Miller and Weinberg, 2017), or on upstream market power (Prager and Schmitt, 2021), typically assume either exogenous input prices and/or exogenous productivity in response to mergers. I find that allowing for endogenous input prices is crucial to correctly infer the productivity effects of ownership consolidation, and vice-versa.

A second contribution of this paper is to document the effects of oligopsony power on both allocative efficiency and income redistribution. Whereas the allocative effects of oligopoly power have been well-documented (Harberger, 1954; Edmond et al., 2022; Asker et al., 2019), much less is known about the allocative effects of oligopsony power. Berger, Herkenhoff, and Mongey (2022) examine the allocative effects of labor market power in a general equilibrium oligopsony model. In contrast to their paper, I directly estimate the productivity effects of oligopsony power using production and cost data, do not have to rely on a specific conduct assumption upstream, make use of a large observed shock to input market structure, and focus on intermediate input rather than labor markets. Therefore, I consider the micro-evidence provided in this paper to be complementary to the macro-economic framework of Berger et al. (2022).

A third contribution of this paper is methodological. Markdown identification approaches that rely on production functions are increasingly used in the literature (Morlacco, 2017; Brooks et al., 2021; Hershbein et al., 2022; Lamadon et al., 2022) I show that the presence of non-substitutable inputs, as is often the case for intermediate inputs, leads to non-identification of markdowns from markups using only the production function approach. This identification challenge can be overcome by either combining the production model with an input supply model, or by relying on conduct and preference assumptions downstream. I find that this markdown identification challenge has important implications for the inference of both the level and changes of input price markdowns and productivity. Using a substitutable leaf model, the consolidation seems to increase productivity, but this is merely due to falling intermediate input prices. This potentially sheds new light on prior evidence that documented productivity-increasing effects of Chinese consolidation programs in other industries, such as Hsieh and Song (2015) and Chen et al. (2021). If input prices are endogenous to firms, these effects could be partly due to the exertion of oligopsony power,

rather than to efficiency gains.

The remainder of this paper is structured as follows. Section I. contains the industry background, data, and stylized facts. Section II. presents the model. Section III. discusses estimation and identification of the model. Finally, Section IV. examines the consequences of oligopsony power for allocative inefficiency and income redistribution.

I. Key facts on the Chinese tobacco industry

A. Industry setting

Farming

The value chain of the Chinese tobacco industry is visualized in Figure 1a. At the start of the panel in 2003 there were around 4 million tobacco farms in China (Hu et al., 2006), which were mostly organized at the household level and operated small plots of around 0.3-0.4 ha (FAO, 2003). Tobacco plants are cultivated as annual crops rather than perennials, which implies that planting decisions are made on an annual basis. After being harvested and dried, tobacco leaf needs to be 'cured'. Farmers sell cured tobacco leaf to cigarette manufacturers or their agents through local 'purchasing stations' (State Council of the People's Republic of China, 1997). Hence, there are sometimes intermediaries between farmers and manufacturers, but they are vertically integrated with the manufacturers and hence will not be modeled separately. Tobacco farmers need to decide which purchasing station to bring their leaf to, and get paid upon arriving at the purchasing station. Choosing the purchasing stations within their own county borders (State Council of the People's Republic of China, 1997).

Local tobacco leaf procurement prices are posted annually by the industry regulator, the STMA, across 40 quality levels or "grades". Despite this regulatory environment, cigarette manufacturers can nevertheless set leaf prices flexibly, in two ways. First, the 40 quality grades lack a scientific benchmark, so cigarette manufacturers have a lot of flexibility to choose leaf prices through subjective grading (Hu, 2008; Peng, 1996). This gives manufacturers the ability to adapt leaf prices to supply and demand shocks. Second, at the local level, the STMA boards that decide on the price-grading schedules consist of the managers of the local cigarette manufacturer (Wang, 2013), effectively allowing these manufacturers to set their own prices. Anecdotal evidence suggests that cigarette manufacturers compete against each other for leaf (Hu, 2008; Wang, 2013).

Chinese tobacco farms became less profitable throughout the sample period. While tobacco was the median cash crop in terms of farm profitability in 1997, it dropped to the last place in 2004. (FAO, 2003; Hu et al., 2006). Tobacco farmers can switch to other crops, but this entails switching costs. A policy intervention in which Chinese tobacco farmers were helped to substitute

crops in 2008 found that substituting away from tobacco increased annual revenue per acre by 21% to 110% (Li et al., 2012). The fact that farmers do not substitute despite these potential gains suggests substantial switching costs. Some sources also mention coercion of tobacco farmers into not switching crops by local politicians, due to the importance of tobacco for local fiscal revenue (Peng, 1996). Besides switching to other crops, farmers can also exit agriculture altogether, but rural emigration is constrained due to the Hukou registration system. Land tenure insecurity also makes migration more costly. Because rural land is the property of villages or collectives, farmers lose their exclusive land use rights when moving (Minale, 2018).

Figure 1: Tobacco industry structure



Notes: Panel (a) gives a schematic overview of the cigarette value chain in China. The manufacturers, in bold in panel (a), are the entities observed in this paper. Panels (b)-(c) show cigarette manufacturing locations in 1999 and 2006.

Manufacturing

Cigarette manufacturers turn tobacco leaf and other intermediate inputs, such as paper and filters, into cigarettes using labor and capital. Variable input expenditure consists for 90% of intermediate inputs, and for 10% of labor. Tobacco leaf accounts for around two-thirds of intermediate input

expenditure, so I will refer to intermediate inputs as 'tobacco leaf' in the rest of the paper.⁴ Almost all Chinese cigarette manufacturers are formally subsidiaries of the *Chinese National Tobacco Corporation* (CNTC). In practice, however, they operate as separate enterprises responsible for their own losses and profits (Peng, 1996). As was mentioned above, they are autonomous in how they operate and compete against each other (Wang, 2013). A map of tobacco manufacturing locations in 1999 and 2006 is in Figures 1b and 1c.

Wholesaling

Manufacturers sell their cigarettes to a monopsonistic wholesaler, which is controlled by the State Tobacco Monopoly Administration (STMA) through its commercial counterpart, the Chinese National Tobacco Trade Corporation (CNTTC). STMA and CNTTC share most of their leadership (Wang, 2013). The CNTTC is centrally controlled and operates a monopoly on the cigarette market. This wholesaler unilaterally sets factory-gate cigarette prices (Wang, 2013; Nargis et al., 2019). In contrast to tobacco leaf, cigarette markets are not isolated; they are sold across all of China (State Council of the People's Republic of China, 1997). The distinction between centrally controlled wholesaling and decentralized manufacturing has been at the core of the STMA system since its inception in the early 1980s. Even after China joined the World Trade Organization in 2001, the Chinese tobacco industry has been shielded from international competition. Aggregate trade statistics show that industry-wide exports and imports were only 1.0% and 0.2% of total industry revenue between 1998 and 2007 (United Nations, 2019). The fiscal importance of the tobacco industry may be an important reason for this protection. In 1997, tobacco taxes and monopoly profits made up for 10.4% of central government revenue. In 2015, tax revenues from the cigarettes industry amounted to ¥840 B, which is 6.2% of China's total tax revenue, according to the 2015 annual report of the State Administration of Taxation (State Administration of Taxation, 2015).

B. Data

Production and cost data

I use production and cost data on cigarette manufacturers between 1999 and 2006 from the *An-nual Survey of Industrial Firms* (ASIF), which is conducted by the National Bureau for Statistics (NBS). The above-scale survey includes non-SOEs with sales exceeding 5 million Chinese Renminbi (RMB) and all SOEs irrespective of their size. I refer to Brandt et al. (2012) and Brandt

⁴The Chinese data do not break down intermediate inputs into more detailed categories, but US census data from 1997 show that tobacco leaves make up for 60% of all intermediate input costs in tobacco manufacturing firms (U.S. Census Bureau, 1997). Other intermediate inputs, such as filters and paper, fit the assumptions made for tobacco leaf, as they are likely to be non-substitutable as well.

et al. (2017) for a comprehensive discussion of this data set. Unique firm identifiers provided by the NBS are made consistent over time in order to avoid false entries and exits. Firms that exit the market also exit the dataset. The unit of observation in the NBS data is the 'establishment', which also includes subsidiaries. As mentioned earlier though, cigarette manufacturing establishments can be considered to be independent firms, and will therefore be referred to as 'manufacturing firms' in the rest of the paper. I retain all manufacturers in the sector "Tobacco and Manufactured Tobacco Substitutes", which includes cigar and cigarette substitute producers, as well as 'pure' cigarette producers. The product-level descriptions in the data show, however, that firms in the former categories often produce cigarettes as their main product as well, which is why they are included, even though they represent less than 5% of total revenue. The resulting ASIF sample consists of 470 firms and 2,025 observations.

I supplement the ASIF data with production quantity data at the product-firm-month level during the same time period, which is collected by the NBS as well. Quantities are observed for a subset of 1,215 observations. Combining both data sets reduces the sample size to 1,132 observations and 257 firms. This sub-sample covers 78% of total revenue in the raw data.

Additional datasets

I merge multiple other datasets to the firm-level production and cost dataset. County-level demographical information is obtained from the 2000 census of population (China Data Lab, 2020). For a subset of the firms, I obtain brand-level cigarette characteristics from O'Connor et al. (2010), which are used in the robustness checks section. The NBS product-level dataset provides a breakdown of production into four quality grades and information on subsidies for a subset of the years. Agricultural price data are obtained from the Food and Agriculture Organization (FAO, 2019). Aggregate trade flows are from United Nations COMTRADE database (United Nations, 2019). Finally, county-level weather data are obtained from the Chinese Meteorological Agency (China Meteorological Agency, 2018). More details on all these data sources and selected summary statistics are in Appendix B.

C. Ownership consolidation

In May 2002, the STMA ordered all state-owned firms producing less than 100,000 cigarette cases per year to be closed down, whereas state-owned firms with an annual production below 300,000 cases were 'encouraged' to merge with larger firms (State Tobacco Monopoly Administration, 2002). This reform intended to "enable China's cigarette industry to achieve scale and efficiency" (State Tobacco Monopoly Administration, 2002, p2). The number of cigarette manufacturers dropped from 340 in 1999 to 167 in 2006. Figure 2a compares the number of firms which produce less and more than 100,000 cases per year, for those firms at which production quantities

are observed. As quantities are observed for only a subset of firms, the annual number of firms reported is lower than total. Of the 98 firms that produced below the exit threshold in 2002, only 11 survived by 2006. Of these, one existed on paper only but no longer produced output, and 7 firms were not state-owned, so they could not be forced to close down. That leaves just three 'non-complier' firms that kept existing while being below the exit threshold. In contrast, of the 99 firms that produced more than 100,000 cases in 2002, 48 survived. The firms producing less than 100,000 cases represented half of all firms in 2002 and generated 7% of industry revenue.



Figure 2: Market structure

Notes: Panel (a) shows the evolution of the number of cigarette manufacturing firms below and above the exit threshold of 100,000 cases per year. This graph excludes firms for which quantities are unknown, which is why the total number of firms is below 470. Panel (b) plots the evolution of the aggregate revenue share of labor and intermediate inputs in the Chinese cigarettes industry.

Factor revenue shares

Figure 2b plots the evolution of the ratio of total labor and intermediate input expenditure over total revenue in the industry (all deflated). The aggregate labor share of revenue fluctuated at around 3%, whereas the aggregate intermediate input share of revenue fell from 41% to 28% between 1999 and 2006. Hence, the variable cost share of tobacco leaf dropped sharply. One explanation for this could be that less tobacco leaf was needed to produce a cigarette compared to labor. However, this is unlikely; there is very limited variation in the required amount of tobacco leaf per cigarette across firms.⁵ The amount of labor needed per cigarette could have changed due to mechanization, but this would result in a falling cost share of labor, which is the opposite of the evolution shown in

⁵Evidence for this is presented in Appendix D.

Figure 2b. A second, more plausible, explanation for this pattern is that leaf prices fell compared to labor wages.



Figure 3: Relative cost shares and consolidation

Notes: Panels (a)-(b) compare the average and median ratio of labor expenditure over intermediate input expenditure over time between the consolidation treatment and control groups.

The relative fall in leaf prices compared to labor costs could be due to rising oligopsony power on leaf markets, but it could also be due to other reasons, such as general equilibrium effects from strong wage growth in other Chinese manufacturing sector over the same time period. In order to isolate the effects of increased market concentration, I make use of the size thresholds in the consolidation policy. Let \mathcal{F}_{it} be the set of firms f in market i in year t. Each firm produces a number of cigarette cases Q_{ft} . The number of firms producing less than 100,000 cigarette cases in market i and year t is denoted N_{it} , using the indicator function I:

$$N_{it} = \sum_{f \in \mathcal{F}_{it}} (\mathbb{I}[Q_{ft} < 100, 000])$$

The policy forced firms producing less than 100,000 cases prior to 2002 to exit from 2002 onwards. I construct a consolidation treatment variable Z_f , which is a dummy indicating whether firm f is located in a county in which there was at least one firm producing below the exit threshold in 2001, just before the reform started: $Z_f = \mathbb{I}[N_{i,2001} > 0]$.

In Figures 3a and 3b, I compare the average and median leaf-to-labor cost ratio between the treatment and the control group. This average is a weighted average by labor usage. Between 1999 and 2002, the average and median leaf-labor expenditure ratios were very similar between both groups, and moved in parallel. In contrast, between 2002 and 2006, the average leaf-labor ratio dropped by half for the firms in the treatment group, while it fell by only 15% for firms in the control group. The relative drop in the leaf-labor ratio was even more pronounced at the median

firm. Hence, the consolidation policy seems to have contributed to the drop in the cost share of leaf after 2002.

Difference-in-differences model

In order to provide more evidence on the effect of ownership consolidation on both input and output prices, I specify a difference-in-differences model in Equation (1). I compare firms with and without competitors below the exit threshold before and after 2002 in terms of an outcome variable y_{ft} , which is the same comparison as was shown in Figure 3. I use the log ratios of labor costs, leaf costs, and revenue over output as the left-hand side variables, as these ratios contain information about input and product price variation. The consolidation dummy Z_f itself is not included on the right-hand side, as it is subsumed into the firm dummy θ_f . I include a linear time trend, with coefficient θ_3 . The coefficient of interest that quantifies the consolidation effects is θ_2 . The residual ε_{ft} contains time series variation in the left-hand variables of interest that is not explained by the consolidation.

(1)
$$y_{ft} = \theta_0 + \theta_1 \mathbb{I}[t \ge 2002] + \theta_2 Z_f \mathbb{I}[t \ge 2002] + \theta_3 t + \theta_f + \varepsilon_{ft}$$
with $y \in \left\{ \log\left(\frac{\text{Leaf cost}}{\text{Cigarette}}\right), \log\left(\frac{\text{Labor cost}}{\text{Cigarette}}\right), \log\left(\frac{\text{Revenue}}{\text{Cigarette}}\right) \right\}$

Leaf market definitions

To estimate the effects of ownership consolidation in a difference-in-differences framework, a leaf market definition is required to partition manufacturers into treatment and control groups. As was explained above, tobacco leaf cannot be transported across county boundaries without official approval of the provincial State Tobacco Monopoly Administration (STMA), so counties are the natural leaf market definition. Throughout the paper, I keep counties as leaf market definitions, and discuss the robustness of the results to using different market definitions in Appendix C7. The average county has 1.24 cigarette manufacturers, and 99% of counties have three or less manufacturers. In counties with just one or just two firms, leaf prices are 69% and 57%, lower compared to markets with more than two firms.⁶

Assumptions

This difference-in-differences model implies three assumptions. First, the evolution of leaf and labor costs per cigarette, and of cigarette prices need to be parallel for both the treatment and the

⁶This evidence is presented in Appendix E1.

control group in the absence of the treatment. Hence, there could have been no policy changes or shocks to the business environment that led to changing relative prices and affected the treatment group differently from 2002 onwards, other than the consolidation. One element in favor of this assumption is that other policy interventions, such as tax reforms, did not use size thresholds (Goodchild and Zheng, 2018). Tests for whether the pre-trends in the dependent variables y_{ft} are parallel between the treatment and the control group will be discussed together with the results. The second assumption is that the assignment of firms into control and treatment markets before 2002 should be independent from the subsequent evolution of input prices, output prices, and input requirements per cigarette. Firms cannot control the output levels of their competitors. They could have self-selected into operating in markets with firms below the exit threshold, but this is in contrast with how this industry operates. Cigarette manufacturers are controlled by local governments and operate in their own jurisdiction, so they are not mobile. Firms could also selfselect into one of the three size groups by adjusting their production, if they had ex-ante knowledge of the consolidation policy. In this case, we would expect some 'bunching' of firms just above the exit threshold, but this is not the case.⁷ The final assumption is that there can be no spillover effects from the treatment to the control group throughout the panel. For leaf prices, this assumption is subsumed into the isolated markets assumption made earlier, which follows from the leaf transport restrictions.

Results

The estimates of Equation (1) are in Table 1a. The change in the average labor cost per cigarette was not significantly different between firms in treatment and control markets. However, leaf costs per cigarette fell by 50% on average,⁸ and factory-gate cigarette prices by 31%. The estimates in Table 1b show that the trends in all three dependent variables were not significantly different before 2002. Therefore, increasing market concentration seems to have mainly led to lower leaf costs per cigarette, and to a lesser extent to lower cigarette prices, while not changing wages.

II. Model: markdowns, markups, and productivity

The evidence that leaf costs per cigarette dropped in response to the consolidation, whereas labor unit costs did not, is not sufficient to draw conclusions about the underlying mechanism. Generally speaking, equilibrium input price changes after mergers can be due to increased markdowns, or to changes in the marginal revenue product of that input, for instance due to markup or productivity

⁷I refer to Figure A1 in the Online Appendix for the annual firm size distributions, which do not have any discontinuity around 100,000 and 300,000 cigarette case (the exit and merger thresholds).

 $^{^{8} = \}exp(-0.686) - 1$

panel A: Treatment effects	log(Labor cost/output)		log(Leaf cost/output)		log(Revenue/output)	
	Est.	S.E.	Est.	S.E.	Est.	S.E.
$T \qquad (*1(X) > 2002)$	0.075	0.100	0.606	0 1 4 0	0.264	0.116
Treatment*1(Year \geq 2002)	-0.075	0.109	-0.686	0.148	-0.364	0.116
R-squared	0.83		0.85		0.86	
Observations	1132		1132		1132	
	log(Labor cost/output)		log(Leaf cost/output)			
panel B: Pre-2002 trends	log(Labo	r cost/output)	log(Leaf	cost/output)	log(Reve	nue/output)
panel B: Pre-2002 trends	log(Labo Est.	r cost/output) S.E.	log(Leaf Est.	cost/output) S.E.	log(Reve Est.	nue/output) S.E.
panel B: Pre-2002 trends	log(Labo Est.	r cost/output) S.E.	log(Leaf Est.	cost/output) S.E.	log(Reve Est.	nue/output) S.E.
panel B: Pre-2002 trends Treatment*Year	log(Labo Est. 0.089	r cost/output) S.E. 0.069	log(Leaf Est. 0.074	cost/output) S.E. 0.076	log(Reve Est. 0.007	nue/output) S.E. 0.070
panel B: Pre-2002 trends Treatment*Year R-squared	log(Labo Est. 0.089	r cost/output) S.E. 0.069 0.17	log(Leaf Est. 0.074	cost/output) S.E. 0.076 0.11	log(Reve Est. 0.007 0	nue/output) S.E. 0.070 .07

Table 1: Consolidation, unit costs, and prices

Notes: Panel (a) reports the average treatment effects from Equation (1), with the left-hand variables being labor cost per cigarette, leaf cost per cigarette and revenue per cigarette, all in logs and at the firm-year level. Panel (b) reports the interaction term of the time trend and treatment dummy prior to 2002, which is a test of whether the pre-trends are parallel. This second regression does not include firm fixed effects. Standard errors are clustered at the county level.

changes. Hence, a model is necessary to separately identify these variables of interest. In this section, I construct such a model that is tailored to the Chinese tobacco manufacturing industry.

A. Input demand

Production

Cigarette manufacturers f produce Q_{ft} cases of cigarettes using tobacco leaf M_{ft} , labor L_{ft} , which are both expressed in quantities, and fixed assets K_{ft} . I assume that tobacco leaf cannot be substituted with either labor or capital. The amount of tobacco leaf needed to produce a case of cigarettes, $\frac{1}{\beta_{ft}^M}$, is assumed to be a scalar that is common across firms, $\beta_{ft}^M = \beta^M$. Brand-level data reveals very little variation in leaf contents per cigarette across manufacturers, as well as other observable characteristics.⁹ The average manufacturer uses 686 mg of tobacco leaf per cigarette of 1000 mg, and the standard deviation of this content is only 30 mg.¹⁰ Let the production function be given by Equation (2):

(2)
$$Q_{ft} = \min\left\{\beta_{ft}^M M_{ft}, \Omega_{ft} H(L_{ft}, K_{ft}, \boldsymbol{\beta})\right\}$$

⁹This evidence is shown in Table A3 in the Online Appendix.

¹⁰I refer to Appendix D for a discussion of the implications of leaf content heterogeneity and to Appendix C3 for an estimate of the elasticity of input substitution between labor and leaf.

Assumption 1. — Cigarette manufacturers do not differ in terms of the physical leaf content per cigarette: $\beta_{ft}^M = \beta^M$.

Manufacturers differ in terms of their productivity level Ω_{ft} . In the main specification, this productivity term is assumed to be a scalar, and assumption which I examine in Appendix C4. Firms use a production technology H(.) in labor and capital with common parametrization β . I assume H(.) is twice differentiable in both labor and capital. In the main model, there is no measurement error in output, Appendix E1 relaxes this assumption. Equation (2) nests production functions in which all inputs are substitutable; the leaf parameter β_{ft}^M would then be zero, and leaf would be an additional subtitutable input in H(.).

Product and input differentiation

I assume that manufacturers produce a single product, cigarettes, at a price P_{ft} . Although the model could be generalized to a multi-product setting, for instance using De Loecker et al. (2016), this is not of first-order importance because 88% of firms are single-product firms, and they together account for 91% of sales. Cigarettes are differentiated products, as is tobacco leaf, and high-quality tobacco leaf is required to produce high-quality cigarettes. This is not inconsistent with Assumption 1: quality differentiation does not imply differences in the physical amount of leaf per cigarette, but in the quality of this leaf. I assume that a single quality dimension captures both cigarette and leaf quality variation, and that this quality level is constant within firms over time. In Appendix D, I defend this assumption by relying on more detailed quality information that is available for a subset of the dataset. In Section III.A., I will discuss the implications of cross-sectional quality variation for the identification of the production function and of leaf price markdowns.

Assumption 2. — Cigarette quality differs across cigarette manufacturers f, but is constant over time.

Input markets

I assume that tobacco leaf is a variable input: its amount is chosen by manufacturers in the same time period as when it is used. I also assume that tobacco leaf is a static input, which means that it only affects current profits. These assumptions rule out adjustment costs over cigarettes and inventories. Labor is assumed to be a variable and static input as well. Cigarette manufacturing factories rely mainly on production workers, for which these assumptions are more likely to hold compared to white-collar workers.¹¹ Capital is, in contrast, dynamic and fixed; the capital stock at

¹¹The NBS surveys does not distinguish production from non-production workers, but 70% of US cigarette manufacturing employees and 65% of the wage bill were production workers, and, thus, variable, in 1997 (U.S. Census Bureau, 1997).

time t can only be changed through investment at time t - 1.

The prices of leaf and labor are W_{ft}^M and W_{ft}^L . The extent of oligopsony power of a manufacturer f over an input $V \in \{L, M\}$ is parametrized by the input market power term ψ_{ft}^V , which is equal to one plus the inverse elasticity of input supply:

$$\psi_{ft}^V \equiv \frac{\partial W_{ft}^V}{\partial V_{ft}} \frac{V_{ft}}{W_{ft}^V} + 1 \ge 1$$

If a manufacturer has oligopsony power over input V, the input price W^V increases if more inputs are purchased, meaning that $\psi_{ft}^V > 1$. If the price of an input V is exogenous to a manufacturer, this implies that $\psi_{ft}^V = 1$.

Manufacturer decisions

I assume that in every year, firms choose labor in order to minimize their current variable costs. Given that labor and intermediate inputs are not substitutable, choosing labor implies choosing a quantity of intermediate inputs, and hence also output. Through their input quantity choices, firms also set input prices if input supply is upward-sloping. I assume that manufacturers set input quantities, and hence leaf prices, unilaterally. Hence, tobacco farmers are assumed to be price-takers, which is intuitive given that there are many more farmers than manufacturers.

(3)
$$\min_{L_{ft}} \left(W_{ft}^M M_{ft} + W_{ft}^L L_{ft} - \lambda_{ft} (Q(M_{ft}, L_{ft}, K_{ft}; \beta^M, \beta)) - Q_{ft} \right)$$

Assumption 3. — Firms choose their variable inputs in each period in order to minimize their variable costs.

The cost minimization assumption can be questioned: it is often suggested that SOEs have idiosyncratic objectives, such as generating local employment (Lu and Yu, 2015). In the tobacco industry, however, (Peng, 1996) notes that cigarette manufacturers have "the purpose of making profits" and "often bargain with each other for better deals". In Appendix E2, I discuss the implications for non-cost minimizing behavior for the markdown and productivity estimates.

B. Markups and markdowns

Using the cost minimization problem from Equation (3), marginal costs λ_{ft} can be written as:

$$\lambda_{ft} = W_{ft}^L \psi_{ft}^L \frac{\partial L_{ft}}{\partial Q_{ft}} + W_{ft}^M \psi_{ft}^M \frac{\partial M_{ft}}{\partial Q_{ft}}$$

Given that labor and materials are not substitutable, the marginal cost expression contains both

the marginal labor cost and the marginal material cost: increasing output by one unit implies additional usage of both materials and labor. The input market power terms ψ_{ft}^M and ψ_{ft}^L enter the marginal cost expression because input prices endogenously increase when input quantities increase, if the input supply curves are upward-sloping.

Markups

The markup ratio μ is the ratio of factory-gate prices over marginal costs: $\mu_{ft} \equiv \frac{P_{ft}}{\lambda_{ft}}$. Substituting the marginal cost expression into the markup formula results in Equation (4a), with the revenue shares of each input being denoted as $\alpha_{ft}^V \equiv \frac{V_{ft}W_{ft}^V}{P_{ft}Q_{ft}}$, with $V \in \{L, M\}$ and the output elasticity of labor being $\beta_{ft}^L \equiv \frac{\partial Q_{ft}}{\partial L_{ft}} \frac{L_{ft}}{Q_{ft}}$.

(4a)
$$\mu_{ft} = \left(\frac{\alpha_{ft}^L}{\beta_{ft}^L}\psi_{ft}^L + \alpha_{ft}^M\psi_{ft}^M\right)^{-1}$$

The markup expression in Equation (4a) nests previous markup models. I discuss three special cases that appeared in the prior literature. First, suppose all inputs have exogenous prices and are mutually substitutable. In that case, the non-substitutable input revenue share is by definition zero, $\alpha_{ft}^M = 0$, and the input supply functions are flat, $\psi_{ft}^V = 1$, $\forall V$. The markup expression then simplifies to the formula of De Loecker and Warzynski (2012):

(4b)
$$\mu_{ft} = \frac{\beta_{ft}^L}{\alpha_{ft}^L}$$

Next, consider a setting in which all inputs are substitutable, but in which input prices are endogenous. The markup is now expressed as the output elasticity of a variable input divided by its revenue share *and* its inverse supply elasticity. This corresponds to the expression from Morlacco (2017).

(4c)
$$\mu_{ft} = \frac{\beta_{ft}^L}{\alpha_{ft}^L \psi_{ft}^L}$$

Finally, assume that all input prices are exogenous, $\psi_{ft}^V = 1, \forall V$, but that there is a nonsubstitutable input M: $\alpha_{ft}^M > 0$. The markup is given by Equation (4c), which corresponds to expression from De Loecker and Scott (2022).

(4d)
$$\mu_{ft} = \left(\frac{\alpha_{ft}^L}{\beta_{ft}^L} + \alpha_{ft}^M\right)^{-1}$$

Markdowns

The input market power terms ψ_{ft}^V were defined as one plus the inverse elasticity of input supply, for each variable input V. By rewriting the cost minimization problem as a profit maximization problem, one can see that these input market power terms also have the interpretation of an input price 'markdown ratio'. For instance, taking tobacco leaf M, one can write a profit maximization problem as $\max_{M_f t} P_{ft}Q_{ft} - W_{ft}^M M_{ft} - W_{ft}^L L_{ft}$. Rearranging the corresponding first order conditions in function of the input price elasticity of leaf ψ_{ft}^M gives the following expression.

$$\psi_{ft}^{M} = \frac{\frac{\partial (P_{ft}Q_{ft})}{\partial M_{ft}} - \frac{\partial (W_{ft}^{L}L_{ft})}{\partial M_{ft}}}{W_{ft}^{M}}$$

The numerator of this expression is the marginal revenue product of leaf, minus the additional labor cost required when increasing leaf by one unit. Hence, the numerator can be interpreted as the marginal revenue contribution of tobacco leaf, net of marginal labor costs. The denominator is the unit price of leaf. The parameter ψ_{ft}^M hence indicates the extent to which the marginal benefit of tobacco leaf to the manufacturer exceeds the leaf price. In a competitive leaf market, $\psi_{ft}^M = 1$, which implies that leaf farmers get paid their marginal contribution. In contrast, if $\psi_{ft}^M = 2$, tobacco farmers receive 50% of their marginal benefit to the cigarette manufacturer.

In the literature, the markdown is often also defined as a 'markdown wedge' $\delta_{ft}^M \equiv \frac{\psi_{ft}^M - 1}{\psi_{ft}^M}$, which is the percentage to which the leaf price is marked down below its marginal benefit. For the sake of clarity, I will only report and discuss the markdown ratio ψ_{ft}^M , and refer to it as 'the markdown'. The markdown ratio has the advantage of being scaled similarly to the markup ratio μ , with a support on $[0, \infty]$ and a value of one that corresponds to exogenous prices. The previously made assumption that input suppliers are price-takers implies that $\psi_{ft}^M \ge 1$. The estimation procedure does not force the estimate of ψ_{ft}^M to be larger than one, but the results will be consistent with this range. The product of ψ_{ft}^M and μ_{ft} has the interpretation of a variable profit margin. This means that firms can operate at a positive variable profit even if the markup is below one; there is a wedge both between the product price and marginal costs, and between marginal costs and input prices.

Identification challenge

As is usually the case in production-cost datasets, the revenue shares α_{ft}^M and α_{ft}^L are observed. If all input prices are exogenous, identification of the production function suffices to identify the markup, as can be seen in equations (4b) and (4d). If input prices are endogenous, both markups and markdowns can still be identified by only identifying the production function, if all inputs are substitutable and if there is at least one variable input of which the price is exogenous. This can be seen from Equation (4c). Markdowns can be found by dividing each markup obtained using an input with an endogenous price by the markup expressed using the input with the exogenous price. In the general case with both non-substitutable inputs and endogenous input prices of Equation (4a), however, the unknown parameters are the markup μ_{ft} , the markdowns ψ_{ft}^M and ψ_{ft}^L , and the output elasticity of labor β_{ft}^L . Only knowing β_{ft}^L is insufficient to identify the markup; the wedge between the output elasticity of an input and its revenue share can be due to both market power upstream or downstream.

There are three potential identification strategies to still identify markups from markdowns. The first possibility is to identify the markdowns ψ_{ft}^L and ψ_{ft}^M by modeling the supply of each input, which involves taking a stance on supplier utility and behavior, and on competitive conduct of buyers on their input markets. In combination with the output elasticity β_{ft}^L , which can be identified following the production function literature, this leads to identification of the markup μ_{ft} without having to impose a model of demand for cigarettes and of how manufacturers compete downstream. Second, the opposite approach is possible too, which is to impose a model of how firms compete downstream, in order to identify the markup μ_{ft} , and combine this with a production function to identify markdowns. Finally, one could also combine an upstream input supply model with a downstream goods demand model, but remain agnostic about the production function. The optimal identification strategy depends both on data availability and on industry characteristics. In the next section, I will discuss the concrete implementation of this markdown identification approach in the context of the Chinese tobacco industry.

III. Estimation

I now turn to the estimation of the model, which consists of recovering the production function coefficients, input price markdowns, and the effects of the ownership consolidation on markdowns and productivity.

A. Production function

Taking the logarithm of the production function, Equation (2), results in Equation (5a). As tobacco leaf is assumed to be non-substitutable and a linear function of the number of cigarettes, it does

not enter the estimable production function.¹² The production function coefficients β need to be identified.

(5a)
$$q_{ft} = h(l_{ft}, k_{ft}, \boldsymbol{\beta}) + \omega_{ft}$$

As was mentioned before, both cigarettes and tobacco leaf are differentiated products, with important quality differences. The 'output price bias' of Klette and Griliches (1996) is taken into account because output is observed in physical units, rather than in quantities. Labor inputs are observed in units as well, but potentially with error; rather than observing the total hours worked of quality-adjusted labor l_{ft} , I observe the number of employees without quality adjustment, \tilde{l}_{ft} . Capital is measured in monetary values \tilde{k}_{ft} , rather than in physical units k_{ft} , so any variation in capital prices due to differences in technological sophistication is latent as well. If these latent differences in quality and utilization of labor and capital are correlated with cigarette quality, this induces an 'input price bias', as discussed by De Loecker et al. (2016). This is likely to be the case for the tobacco industry. Luxury cigarettes, which are mainly used as gifts, are often handcrafted and hence require more labor hours than cheap cigarettes. As per De Loecker et al. (2016), a function a(.) of wages per worker and cigarette prices was added to the production function to address this input price bias. I refer to De Loecker et al. (2016) for a formal model and discussion of input price bias. Although tobacco leaf is differentiated in terms of quality levels as well, this does not induce input price bias because leaf does not enter the estimable production function.

(5b)
$$q_{ft} = h(\hat{l}_{ft}, \hat{k}_{ft}, \boldsymbol{\beta}) + a(p_{ft}, w_{f,t}^L) + \omega_{ft}$$

Identification

In order to identify the production function, I impose timing assumptions on firms' input choices, as proposed by Olley and Pakes (1996), in combination with a functional form assumption on the productivity transition process. Let $\mathbf{Z}_{ft} = (Z_f, \mathbb{I}(t \ge 2002), Z_f \mathbb{I}(t \ge 2002))$ contain the previously introduced indicator variables of whether the firm is subject to the consolidation or not. Let the productivity transition be given by the first-order Markov process in Equation (6a), with an unexpected productivity shock v_{ft} . Given that achieving productivity gains was he official objective of the consolidation policy, I allow the consolidation indicator \mathbf{Z}_{ft} enter the productivity law of motion, as in De Loecker (2013) and Braguinsky et al. (2015).

¹²The usual caveat applies that it could be optimal for firms to diverge from Equation (5a) by setting intermediate inputs to zero if material prices become too high, or output prices too low (Gandhi et al., 2020). Given that intermediate inputs enter the production function linearly, this would imply that firms do not produce at all, at which point they no longer enter the dataset.

(6a)
$$\omega_{ft} = g(\omega_{ft-1}, \mathbf{Z}_{ft}, \mathbf{Z}_{ft-1}) + \upsilon_{ft}$$

As was explained in Section II.A., labor is assumed to be a variable and static input, whereas capital is fixed and dynamic. Hence, labor is assumed to be chosen at time t, after the productivity shock v_{ft} is observed by the firm, while capital investment is chosen at time t - 1, before the productivity shock is observed. Cigarette and worker quality, which are proxied by cigarette prices and wages, were already assumed to be strictly exogenous from the point of view of the manufacturers. As was assumed earlier in the difference-in-differences model, firms cannot choose whether they are subject to the consolidation or not, so the variables \mathbf{Z}_{ft} are exogenous. These timing assumptions lead to the following exclusion restrictions. The productivity shock is orthogonal to current capital usage, coal prices, wages, and being in a consolidated market, and to lagged labor usage.

$$\mathbb{E}\Big[\upsilon_{ft}|(\tilde{l}_{fr-1},\tilde{k}_{fr},p_{fr},w_{fr}^L,\mathbf{Z}_{fr})\Big]_{r\in[2,\dots,t]}=0$$

The usual approach in the literature is to invert the intermediate input demand function to recover the latent productivity level ω_{ft} , which can be used to construct the productivity shock v_{ft} using the productivity law of motion (Levinsohn and Petrin, 2003; Ackerberg et al., 2015). This approach hinges on productivity being the only latent, serially correlated input demand shifter. However, input demand varies due to markup and markdown variation as well, as pointed out by Doraszelski and Jaumandreu (2021). The approach with input inversion can still be used when making additional parametric assumptions on the distribution of markups and markdowns, as I discuss in Appendix C1. Alternatively, one can also impose more structure on the productivity transition process, with the benefit of allowing for fully flexible heterogeneity in markdowns and markups. Given that leaf price markdowns are the main object of interest, I follow this approach in the main text, and include the input demand inversion approach as a robustness check in Appendix C1.

As shown in Equation (6b), I let productivity be a linear function of the consolidation, and denote the residual productivity term as $\tilde{\omega}_{ft}$. Following Blundell and Bond (2000), I assume that this residual productivity term $\tilde{\omega}_{ft}$ follows an AR(1) process with serial correlation ρ , as parametrized in Equation (6c).

(6b)
$$\omega_{ft} = \boldsymbol{\beta}^Z \mathbf{Z}_{ft} + \tilde{\omega}_{ft}$$

(6c)
$$\tilde{\omega}_{ft} = \rho \tilde{\omega}_{ft-1} + v_{ft}$$

Imposing this AR(1) process comes at the cost of ruling out a richer productivity transition function g(.), and of not coping with selection bias due to endogenous entry and exit (Olley and Pakes, 1996). However, moving to an unbalanced panel usually already alleviates most concerns of selection bias (De Loecker et al., 2016). Moreover, exit in the industry was mainly due to the consolidation, which is assumed to be exogenous to the manufacturers. In Appendix C1, I relax this assumption by allowing for endogenous exit, using the control function approach with input inversion of Ackerberg et al. (2015). In Appendix C4, I rely on a cost shares-based method as another robustness check.

Estimation

In the main specification, I use a Cobb-Douglas specification for both the h(.) and a(.) functions: $h(\tilde{l}_{ft}, \tilde{k}_{ft}) = \beta^L \tilde{l}_{ft} + \beta^K \tilde{k}_{ft} + \beta^0$ and $a(w_{ft}^L, p_{ft}) = \beta^W w_{ft}^L + \beta^P p_{ft}$, with β^L and β^K being the output elasticities of labor and capital.¹³ Rewriting the moment conditions above, and including lags up to one year, the moment conditions are given by Equation (7).

(7)
$$\mathbb{E}\Big[q_{ft} - \rho q_{ft-1} - \beta^0 (1-\rho) - \beta^K (\tilde{k}_{ft} - \rho \tilde{k}_{ft-1}) - \beta^L (\tilde{l}_{ft} - \rho \tilde{l}_{ft-1}) - \beta^W (w_{ft}^L - \rho w_{ft-1}^L) - \beta^P (p_{ft} - \rho p_{ft-1}) - \beta^Z (\mathbf{Z}_{ft} - \rho \mathbf{Z}_{ft-1}) | (\tilde{l}_{f-1}, \tilde{k}_{ft}, \tilde{k}_{ft-1}, w_{ft}^L, w_{ft-1}^L, p_{ft}, p_{ft-1}, \mathbf{Z}_{ft})\Big] = 0$$

B. Input supply

Markdown identification

I assume that factory worker wages are exogenous to manufacturers, which means that $\psi_{ft}^L = 1$. I impose this assumption because labor wages did not adjust much in response to the consolidation, and the market for manufacturing workers does not share the leaf markets' institutional feature of being geographically isolated due to transportation restrictions. In contrast to this, I allow leaf prices to be endogenous to the manufacturers, which implies that the firm-level inverse leaf supply elasticity $\psi_{ft}^M - 1$ can be positive.

As was explained above, separately identifying leaf price markdowns ψ_{ft}^M from cigarette price markups μ_{ft} either requires imposing a model of leaf supply, or a model of cigarette demand. In the context of the Chinese tobacco industry, cigarettes are procured by a wholesaling monopsonist (and monopolist), which is controlled by the central government. The wholesaler unilaterally sets cigarette prices (Wang, 2013; Nargis et al., 2019). Because of this institutional feature, it is natural to assume that individual manufacturers do not have any price-setting power downstream. This does not imply perfect competition downstream, but rather monopsonistic price-setting by

¹³I refer to Appendix C5 for an alternative specification using a translog production function.

the wholesaler. Also, the markup ratio μ_{ft} does not necessarily have to be equal to one. One can allow for a 'cost-plus' based pricing rule imposed by the monopsonistic wholesaler, $\frac{P_{ft}}{mc_{ft}} = \mu > 1$, similarly to De Loecker and Scott (2022). This implies that all manufacturers are allowed to mark up their price at $(\mu - 1)\%$ above their marginal cost. The only required assumption is that the markup term μ is not chosen by the manufacturers and common across manufacturers and over time. Under this assumption, the markdown ratio ψ_{ft}^M simplifies to Equation (8). The leaf price markdown is now identified up to the constant μ using (i) the output elasticity of labor β^L , which can be estimated, and (ii) the material and labor shares of revenue α_{ft}^M and α_{ft}^L , which are observed. This means that we can know markdown levels and changes without having to impose a model of leaf supply and competition on leaf markets.

(8)
$$\psi_{ft}^M = \frac{1}{\alpha_{ft}^M} \left(\frac{1}{\mu} - \frac{\alpha_{ft}^L}{\beta^L} \right)$$

with $\beta^L \equiv \frac{\partial Q_{ft}}{\partial L_{ft}} \frac{L_{ft}}{Q_{ft}}$, $\alpha^M_{ft} \equiv \frac{W^M_{ft}M_{ft}}{P_{ft}Q_{ft}}$, and $\alpha^L_{ft} \equiv \frac{W^L_{ft}L_{ft}}{P_{ft}Q_{ft}}$.

Although we cannot identify μ without imposing more structure on the wholesaler's behavior, it is possible to estimate markdown levels and the change in markdowns due to the consolidation under a variety of plausible values for μ , as a robustness check. In the main text, I will report the leaf price markdown levels and changes due to the consolidation assuming $\mu = 1$, in Appendix C7 I explore alternative calibrations for μ .

Although non-strategic markups from the manufacturer point of view follows from the Chinese tobacco institutional setting, it is not an appealing assumption in many other industries. By adding a input supply model to the production function, one can still allow for flexible markup variation in the model without imposing a model of competition downstream. I present such a model in Appendix A. This model leads to similar findings in terms of markdown levels and changes as the main model, but is estimated less precisely.

Quality and markdown variation

As was mentioned earlier, tobacco leaf is vertically differentiated. High-quality leaf is more expensive, so leaf quality enters the markdown expression through the leaf price W_{ft}^M , in the material expenditure share. Moreover, cigarettes are also differentiated, which is reflected through the cigarette price P_{ft} , which also enters the markdown expression, both through the material and labor expenditure shares. To the extent that leaf quality variation affects leaf and cigarette prices differently, this unobserved quality will enter the markdown parameter, which hence might not uniquely measure the exertion of buyer power on leaf markets. However, as was stated in As-

sumption 2, quality is assumed to only vary cross-sectionally across manufacturers, not over time. Hence, the inclusion of firm fixed effects in the difference-in-differences model below will suffice to trace out the effects of the consolidation on buyer power. In Appendix D, I relax this assumption and examine whether quality might have changed over time, in response to the consolidation.

Overview of the timing and information assumptions

In summary, the timing and information assumptions made are as follows. In time period t, cigarette manufacturers receive their productivity shock ν_{ft} . Based on this shock, and taking into account the firm-level elasticity of the leaf supply curve, they choose the amount of labor to use, L_{ft} , which also determines the amount of leaf they procure, and hence also the leaf price. The cigarette price is set by the wholesaler at a markup μ above marginal costs - the manufacturers hence know what the cigarette price will be in function of their input demand choices. Both demand and costs are hence deterministic: the manufacturers know what their costs and prices are conditional on how many cigarettes they produce when making their input decisions. At time t, manufacturers also choose the capital stock in the next period through their investment level. All these decisions are made with knowledge of all productivity levels and leaf supply elasticities in the market. Still in the same time period, manufacturers turn leaf and labor into cigarettes, sell it to the wholesaler at a price P_{ft} , who then sells it to the retailers. At time t + 1, this entire decision process restarts.

In the model described above, no particular assumptions are imposed on how the tobacco farmers choose the manufacturers. In the extended model in Appendix A, more assumptions are imposed on how farmers choose manufacturers, and on how manufacturers compete on the leaf market. In contrast, this relaxes the assumption that the wholesaler unilaterally sets factory-gate cigarette prices at a fixed markup.

C. Consolidation treatment effects

I re-use the difference-in-differences model from Section I.C. in order to know how ownership consolidation affected leaf price markdowns and total factor productivity. For this purpose, I alternatively use the log leaf price markdown and log productivity as the left-hand side variables in Equation (1). This requires additional assumptions compared to the estimation of the markdown, which were discussed when presenting the difference-in-differences model for prices in Section I.C. These assumptions are internally consistent with the model used to estimate the production function and markdowns: the productivity process that was specified in Equation (6b) contains the difference-in-difference model as part of the productivity equation of motion, and the model allowed for flexible heterogeneity in markdowns across firms and over time.

I did not cluster standard errors as in Bertrand et al. (2004), but block-bootstrap them instead

with replacement within firms. This is to take into account the error when estimating the production function, given that the production function estimates are used to construct both the markdown and the productivity residual, which are the left-hand side variables in the difference-in-differences regression. Given that the number of leaf markets, 276, is relatively large, block bootstrapping should perform well (Bertrand et al., 2004).

D. Results

Production function and markdown levels

The estimated output elasticities are in Table 2a. The estimates using the dynamic panel approach are in the right column, and are 0.532 and 0.630 for labor and capital. These estimates are respectively lower and higher compared to the ordinary least squares estimates, which are 0.563 and 0.569 for labor and capital. Both specifications have a scale parameter that is above one, which implies increasing returns to scale. I estimate markdowns using the simplified markdown expression in Equation (8), relying on the production function estimates and assuming that the monopsonistic wholesaler sets factory-gate prices equal to manufacturers' marginal costs. I refer to Appendix C7 for alternative calibrations of the markup parameter. The markdown moments are in Table 2b. Using the dynamic panel approach markdown ratio ψ_{ft}^M is 2.126 at the median manufacturer, and 2.904 at the average manufacturer. This implies that farmers selling to the median firm receive around half of what they would receive in a perfectly competitive input market, and a third of the competitive price when selling to the average manufacturer.

The markdown estimates are lower compared to other studies who estimate markdowns using production-cost models, but assume substitutable inputs, such as Brooks et al. (2021), who find a markdown ratio around 5 in China and India, and Morlacco (2017), who estimates the markdown ratio for French food industries to be around 3.5. More traditional approaches that do not rely on production function estimation find similar or smaller estimates. Naidu et al. (2016) finds an average markdown ratio of 2 for recently hired immigrant workers in the United Arab Emirates, Azar et al. (2022) find a markdown ratio of around 1.2 for online job board vacancies in the U.S.A., Goolsbee and Syverson (2019) find a markdown ratio of 1.5 for tenured college professors, and Ransom and Sims (2010) a markdown ratio of around 1.4 for grocery clerks. The reason for these differences most likely relates to the level of frictions on local labor markets. As was discussed earlier, rural labor markets are highly frictional in China due to immigration restrictions and crop switching costs. The worse the outside employment options of farmers, the higher markdowns should be. I refer to Appendix E1 for correlations between the markdown estimates and firm and market characteristics. These correlations show, for instance, that markdowns are higher in more concentrated leaf markets, which is consistent with the standard oligopsony model.

panel A: Production function	Ordinar	y least squares	Dynamic panel		
	Est.	S.E.	Est.	S.E.	
Output elasticity of labor	0.563	0.082	0.532	0.147	
Output elasticity of capital	0.569	0.066	0.630	0.105	
Scale parameter	1.132	0.044	1.162	0.060	
R-squared		0.91	0.92		
Observations		1130	849		
panel B: Leaf price markdown	Ordinar	y least squares	Dynamic panel		
	Est.	S.E.	Est.	S.E.	
Average	2.934	0.414	2.904	0.442	
Median	2.134	0.066	2.126	0.079	

Table 2: Model estimates

Notes: Panel (a) reports the estimated output elasticities using both OLS and the dynamic panel estimator. Panel (b) contains the leaf markdown moment estimates. Standard errors are block-bootstrapped with 200 iterations.

Consolidation treatment effects

The estimated treatment effects are in Table 3a.¹⁴ Markdowns increased on average by 37% on average for manufacturers affected by the consolidation compared to those in the control group, and this increase is statistically significant. Therefore, the exit of the smaller manufacturers mainly resulted in an increase in oligopsony power of the surviving manufacturers. The third column of Table 3a shows the estimated change in productivity due to the consolidation. On average, the consolidation is estimated to reduce productivity by 5%, although this decrease is not statistically significant. However, we are mainly interested in aggregate productivity changes, rather than in changes on average, because the main objective of the policy was to boost aggregate productivity by forcing small inefficient firms to exit. I consider these aggregate productivity effects in Section IV..

Figure A2 in the Online Appendix provides a visual check of the parallel trends assumption by estimating average annual log markdowns and productivity for both the treatment and control groups. As is explained in more detail in Appendix E1, I estimate interaction effects between all years and the treatment indicator in the difference-in-differences model. Whereas leaf markdowns were very similar both in terms of levels and trend prior to 2002, this was less the case for produc-

¹⁴The number of observations falls for the markdown regression because in 9 instances, the logarithm of the markdown cannot be taken because the estimated markdown is negative.

panel A: Markdown and productivity	log(Markdown)		log(Productivity)			
	Est.	S.E.	Est.	S.E.		
Treatment*1(Year ≥ 2002)	0.315	0.103	-0.055	0.083		
R-squared	0.72		0.88			
Observations	1123		1132			
panel B: Allocative efficiency	log(Agg. TFP)		log(Avg. TFP)		Reallocation	
	Est.	S.E.	Est.	S.E.	Est.	S.E.
Treatment*1(Year ≥ 2002)	-0.544	0.166	-0.084	0.135	-0.460	0.106
R-squared	0.65		0.33		0.77	
Observations	221		221		221	
panel C: Output	log(Agg. output)		log(Avg. output)		Reallocation	
	Est.	S.E.	Est.	S.E.	Est.	S.E.
Treatment*1(Year ≥ 2002)	-0.485	0.171	0.220	0.154	-0.704	0.090
R-squared	0.65		0.48		0.85	
Observations	221		221		221	

Table 3: Consolidation treatment effects

Notes: Panel (a) reports the estimated treatment effects from Equation (1) with the logarithms of the markdown ratio and productivity as the dependent variables. Controls are firm fixed effects and a linear time trend. Panel (b) estimates the effects of the consolidation on log aggregate productivity, weighted by labor usage, log unweighted average productivity, and a reallocation term, all at the province-year level. Panel (c) reports the effects of the consolidation on log total province-level output, log average output, and the difference between these two. All standard errors are block-bootstrapped with 200 iterations.

tivity. This is not entirely surprising given that the central motivation for the consolidation was to address lackluster productivity growth among the smaller producers.

Sensitivity analysis and robustness checks

The main methodological argument made in the paper was that markdowns and markups are no longer separately identified using the production approach if there are non-substitutable inputs. In Appendix C3, I examine the extent to which this identification challenge matters in practice. I show that the key model estimates change substantially when assuming substitutable tobacco leaf. The most striking difference is that using such a model leads to the finding that the consolidation

increased manufacturing productivity. The reason for this is straightforward: adding material expenditure as a substitutable production input causes leaf prices to enter the productivity residual. The drop in leaf prices due to the exertion of oligopsony power is hence interpreted as an increase in total factor productivity in the substitutable leaf model.

I carry out a series of additional robustness checks in Appendix C7, I re-estimate the production function using different identification strategies, different functional forms, different markup calibrations, and different productivity transition equations. I also allow for labor-augmenting technological change, omit small firms from the difference-in-differences sample. These robustness checks lead to very similar conclusions about the levels and changes of markdowns and productivity.

IV. Misallocation and distributional consequences

The previous section showed that consolidating the cigarette manufacturers increased their oligopsony power on tobacco leaf markets. What were the consequences of this increase in oligopsony power? I focus both on allocative efficiency and productivity growth, in Section IV.A., and on the distribution of income, in Section IV.B.

A. Output and productivity growth

The effects of the consolidation on average manufacturer productivity could be very different from its effects on aggregate productivity. Similarly to other types of market power, oligopsony power leads to deadweight loss, as firms hold back on their inputs in order to reduce input prices. Moreover, it could lead to allocative inefficiency, if highly productive firms with oligopsony power produce less than they should in a competitive equilibrium, similarly to the oligopoly setting of Asker et al. (2019). I test this using the decomposition of Olley and Pakes (1996). In order to take into account input reallocation between firms and between leaf markets, I aggregate productivity to the province level, ι . Provincial aggregate productivity is denoted $\bar{\Omega}_{\iota t}$ and is weighted by the number of workers; $\bar{\Omega}_{\iota t} \equiv \sum_{f \in \mathcal{F}_{\iota t}} \left(\frac{\Omega_{f \iota} \tilde{L}_{f \iota}}{\sum_{f \in \mathcal{F}_{\iota t}} (\tilde{L}_{f \iota})} \right)$. Average province-level productivity is denoted as $\hat{\Omega}_{\iota t}$. I estimate how both aggregate and average productivity were affected by the consolidation by estimating Equation (1) at the province-year level.

Table 3b shows the effect of the ownership consolidation on log aggregate productivity and log average productivity. Average productivity at the province-level decreased by 8% in response to the consolidation, although this decrease was not statistically significant. However, aggregate productivity fell by 42%, and this drop is significant. Input misallocation, which is defined as the difference between log aggregate and log average productivity, following Olley and Pakes (1996), fell by 37% due to the consolidation: the decline in aggregate productivity was entirely due to reallocation of market share towards low-productivity firms. This evidence shows that the exertion

of oligopsony power can have important adverse consequences for aggregate productivity growth through reduced allocative efficiency.

Under the assumptions of the classical oligopsony model used in this paper, higher oligopsony power should lead to a lower equilibrium amount of cigarettes produced. This is a testable implication of the model. I estimate how both total and average cigarette production were affected by the consolidation in Table 3c. Average cigarette output increased by 25%, which is logical given that the small firms were forced to exit. However, total cigarette production fell by 38% in the consolidated provinces compared to provinces unaffected by the consolidation. The surviving large manufacturers hence reduced their output, which is how they exerted oligopsony power: this also decreased their leaf usage, and hence leaf prices. This does not necessarily mean that Chinese consumers consumed less cigarette imports and/or increased illegal cigarette production. As both these variables are unobserved, these mechanisms cannot be verified.

B. Distributional consequences

Under a consumer welfare standard, the allocative inefficiency and deadweight loss documented above are the main outcomes of interest. However, there are valid reasons to also care about the distributional consequences of the ownership consolidation, and especially about its effects on income inequality between farmers and manufacturing workers. For instance, the first-best solution of redistributing income through the tax system might not be possible due to political constraints and other imperfections in the tax system. This is especially relevant in the Chinese context, where tobacco leaf taxes are an important source of fiscal revenue of local governments.

Changes in leaf prices and factory worker wages

Although leaf prices W_{ft}^M and quantities M_{ft} are, as usual, not observed separately in the production and cost dataset, the homogeneous leaf content assumption allows to recover leaf prices up to a constant. Dividing leaf expenditure by the number of cigarettes produced gives $W_{ft}^M = \frac{W_{ft}^M M_{ft}}{Q_{ft}} \beta^M$, which is identified up to the inverse leaf content per cigarette β^M .

The Chinese tobacco industry has seen increased income inequality between its two most important factors, leaf farmers and factory workers. Whereas the average wage of factory workers grew on average by 13.7% per year between 1999 and 2006, tobacco remained almost constant, growing at merely 0.4% per year. This divergence fits within a broader trend of increased income inequality between manufacturing workers and farmers in China over the past two decades (Yang, 1999; Ravallion and Chen, 2009). I quantify the extent to which the consolidation of the cigarette manufacturers contributed to this margin of inequality by increasing oligopsony power on leaf markets, but not on manufacturing labor markets. The difference-in-differences model assumes that

leaf prices would have evolved similarly for firms in the control group and treatment groups from 2002 onwards in the absence of the consolidation. I rewrite the difference-in-differences model, Equation (1), to have interaction terms between the treatment indicator and the year fixed effects:

$$\log(W_{ft}^{M}) = \sum_{n=2000}^{2006} \left(\theta_{n}^{1} \mathbb{I}(t=n) + \theta_{n}^{2} \mathbb{I}(t=n) Z_{ft}\right) + \theta^{3} Z_{ft} + \theta^{4} + \varepsilon_{ft}^{M}$$

The predicted leaf price in year t is $\hat{W}_{ft}^{M,1} = \exp(\hat{\theta}_t^1 + \hat{\theta}_t^2 + \hat{\theta}^3 + \hat{\theta}^4)$ for firms in the treatment group, and $\hat{W}_{ft}^{M,0} = \exp(\hat{\theta}_t^1 + \hat{\theta}^4)$ for firms in the control group. The counterfactual leaf prices without consolidation are denoted as $\tilde{W}_{ft}^{M,1}$ and $\tilde{W}_{ft}^{M,0}$ for firms in the treatment and the control group, respectively. For firms in the treatment group, the leaf price would follow the evolution of leaf prices in the control group from 2002 onwards. For firms in the control group, nothing would change: $\tilde{W}_{ft}^{M,0} = \hat{W}_{ft}^{M,0}$

$$\begin{cases} \tilde{W}_{ft}^{M,1} = \exp(\hat{\theta}_t^1 + \hat{\theta}^3 + \hat{\theta}^4) \text{ if } t \ge 2002\\ \tilde{W}_{ft}^{M,1} = \hat{W}_{ft}^{M,1} \text{ if } t < 2002 \end{cases}$$

Results

I calculate the average predicted leaf price per year both in reality and in the counterfactual without consolidation by weighting the predicted prices for the treatment and the control group by the number of firms in each group. Figure 4 shows the ratio of the average leaf price over the average wage in reality (the solid line) and without consolidation (the dashed line). Both series are normalized at 1 in 2001. In reality, wages increased relatively to leaf prices by 49.4% between 2001 and 2006. In the absence of the consolidation, however, wages would increased relatively by 21.5%. This suggests that increased oligopsony power due to the consolidation explains 56% of the increase in factory-worker to farmer inequality over thie period. The fact that manufacturing wages outgrew farmer prices even in the absence of the consolidation is unsurprising, given the fast productivity growth in Chinese manufacturing industries other than tobacco during the 2000s.

Caveats

The analysis above is a partial equilibrium exercise and, as such, comes with a number of caveats. First, it ignores entry and exit of farmers. Higher entry and/or lower exit of farmers in response to higher leaf prices could have suppressed the rise in leaf prices in the absence of the consolidation. Second, tobacco represents a large share of economic activity in some counties, so changing leaf prices likely affected equilibrium prices and wages in other sectors as well. Finally, besides tobacco





Notes: The solid line represents the change in average manufacturing wages and leaf prices compared to 2002 (normalization: 2002=1). The dashed line represents the counterfactual leaf price evolution in the counterfactual scenario in which the exit thresholds were not enforced.

leaf prices, farm productivity and agricultural input costs matter as well for farm profits. However, aggregate producer statistics from the Food and Agriculture Organization (FAO) show that farm sizes remained constant and yields per acre grew by merely 1.8% per year during this time period, which was not enough to compensate for falling leaf prices (FAO, 2019).

V. Conclusion

In this paper, I examine the consequences of oligopsony power on intermediate input markets for allocative inefficiency and the distribution of income. I find that ownership consolidation in the Chinese tobacco industry following a regulatory reform led to an important rise of input price markdowns. This increased exertion of oligopsony power slowed down manufacturing growth due to increased input misallocation, and explains 56% of the increased urban-rural income gap between tobacco farmers and cigarette manufacturing workers.

This finding has three important implications beyond the Chinese tobacco industry setting. First, it demonstrates that ownership consolidation in industries with oligopsonistic competition can lead to important allocative inefficiency, even if firms do not have any pricing power downstream. The exertion of such oligopsony power in a vertical chain is not unique to Chinese tobacco manufacturers, but probably applies to many agricultural and non-agricultural industries that feature a large number of input suppliers that sell to a limited set of manufacturers, such as meat processing, publishers, and motor vehicles. This calls for taking into account oligopsony power when conducting merger and antitrust policy, even when downstream markets are competitive and even when policy-makers are only concerned with consumer surplus. Second, the results in this paper show that industrial policies that shut down small inefficient producers in order to increase aggregate productivity, which are commonplace in low- and middle-income countries, can lead to the opposite outcome due to input reallocation in the presence of input market power. Third, on the methodological front, I show that production-cost approaches cannot separately identify markdowns from markups as soon as a subset of inputs is non-substitutable, and propose two empirical solutions to this challenge. This is important when studying buyer power over intermediate inputs, which are less likely to be substitutable than labor.

References

Ackerberg, D., K. Caves, and G. Frazer (2015). Identification properties of recent production function estimators. *Econometrica* 83(6), 2411–2451.

Asker, J., A. Collard-Wexler, and J. De Loecker (2019). (Mis) allocation, market power, and global oil extraction. *American Economic Review 109*(4), 1568–1615.

Azar, J. A., S. T. Berry, and I. Marinescu (2022). Estimating labor market power. Technical Report 30356, National Bureau of Economic Research.

Berger, D., K. Herkenhoff, and S. Mongey (2022). Labor market power. *American Economic Review 112*(4), 1147–93.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, 242–262.

Bertrand, M., E. Duflo, and S. Mullainathan (2004). How much should we trust differences-indifferences estimates? *Quarterly Journal of Economics 119*(1), 249–275.

Blundell, R. and S. Bond (2000). GMM estimation with persistent panel data: an application to production functions. *Econometric reviews 19*(3), 321–340.

Braguinsky, S., A. Ohyama, T. Okazaki, and C. Syverson (2015). Acquisitions, productivity and profitability: Evidence from the Japanese cotton spinning industry. *American Economic Review 105*(7), 2086–2119.

Brandt, L., J. Van Biesebroeck, L. Wang, and Y. Zhang (2017). WTO accession and performance of Chinese manufacturing firms. *American Economic Review 107*(9), 2784–2820.

Brandt, L., J. Van Biesebroeck, and Y. Zhang (2012). Creative accounting or creative destruction? Firm-level productivity growth in Chinese manufacturing. *Journal of Development Economics* 97(2), 339–351.

Brooks, W. J., J. P. Kaboski, Y. A. Li, and W. Qian (2021). Exploitation of labor? Classical monopsony power and labor's share. *Journal of Development Economics*, 102627.

Chatterjee, S. (2023). Market power and spatial competition in india. *Quarterly Journal of Economics Forthcoming*.

Chen, Y., M. Igami, M. Sawada, and M. Xiao (2021). Privatization and productivity in China. *The RAND Journal of Economics* 52(4), 884–916.

China Data Lab (2020). China County Map with 2000-2010 Population Census Data.

China Meteorological Agency (2018). China Meteorological Dataset.

De Loecker, J. (2013). Detecting learning by exporting. *American Economic Journal: Microeconomics* 5(3), 1–21.

De Loecker, J., P. Goldberg, A. Khandelwal, and N. Pavcnik (2016). Prices, markups and trade reform. *Econometrica* 84(2), 445–510.

De Loecker, J. and P. T. Scott (2022). Markup Estimation using Production and Demand Data: An Application to the US Brewing Industry. Mimeo.

De Loecker, J. and F. Warzynski (2012). Markups and firm-level export status. *American Economic Review 102*(6), 2437–2471.

Doraszelski, U. and J. Jaumandreu (2021). Reexamining the De Loecker and Warzynski (2012) method for estimating markups. *Unpublished*.

Edmond, C., V. Midrigan, and D. Y. Xu (2022). How costly are markups? *Journal of Political Economy Forthcoming*.

FAO (2003). Issues in the global tobacco economy: Selected case studies. In *FAO Commodity Studies*, Volume 1810-0783. Food and Agriculture Organization of the United Nations.

FAO (2019). FAOSTAT Statistical Database.

Gandhi, A., S. Navarro, and D. A. Rivers (2020). On the identification of gross output production functions. *Journal of Political Economy* 128(8), 2973–3016.

Goodchild, M. and R. Zheng (2018). Early assessment of China's 2015 tobacco tax increase. *Bulletin of the World Health Organization* 96(7), 506–512.

Goolsbee, A. and C. Syverson (2019). Monopsony power in higher education: a tale of two tracks. Technical report, National Bureau of Economic Research.

Grieco, P., J. Pinkse, and M. Slade (2017). Brewed in North America: Mergers, marginal costs and efficiency. *International Journal of Industrial Organization* 59(2018), 24–64.

Hall, R. E. (1986). Market structure and macroeconomic fluctuations. *Brookings Papers on Economic Activity* 2(1986), 285–322.

Harberger, A. C. (1954). The welfare loss from monopoly. *American Economic Review* 44(2), 77–87.

Hershbein, B., C. Macaluso, and C. Yeh (2022). Monopsony in the US Labor Market. *American Economic Review 112*(7).

Hsieh, C.-T. and Z. Song (2015). Grasp the large, let go of the small: The transformation of the state sector in China. Working Paper Spring, Brookings Papers in Economic Activity.

Hu, T., Z. Mao, M. Ong, E. Tong, M. Tao, H. Jiang, K. Hammond, K. Smith, J. de Beyer, and A. Yurekli (2006). China at the crossroads: the Economics of tobacco and health. *Tobacco Control* 15(1), 37–41.

Hu, T.-w. (2008). Tobacco Control Policy Analysis in China: Economics and Health.

Jarosch, G., J. S. Nimczik, and I. Sorkin (2019). Granular search, market structure, and wages. Technical report, National Bureau of Economic Research.

Klette, T. J. and Z. Griliches (1996). The inconsistency of common scale estimators when output prices are unobserved and endogenous. *Journal of Applied Econometrics* 11(4), 343–361.

Lamadon, T., M. Mogstad, and B. Setzler (2022). Imperfect competition, compensating differentials, and rent sharing in the US labor market. *American Economic Review 112*(1), 169–212.

Levinsohn, J. and A. Petrin (2003). Estimating production functions using inputs to control for unobservables. *Review of Economic Studies* 70(2), 317–341.

Li, V., Q. Wang, N. Xia, S. Tang, and C. Wang (2012). Tobacco crop substitution: pilot effort in China. *American Journal of Public Health 102*(9), 1660–1663.

Lu, Y. and L. Yu (2015). Trade liberalization and markup dispersion: evidence from China's WTO accession. *American Economic Journal: Applied Economics* 7(4), 221–253.

Manning, A. (2011). Imperfect competition in the labor market. In *Handbook of labor economics*, Volume 4, pp. 973–1041. Elsevier.

Miller, N. H. and M. C. Weinberg (2017). Understanding the price effects of the millercoors joint venture. *Econometrica* 85(6), 1763–1791.

Minale, L. (2018). Agricultural productivity shocks, labour reallocation and rural–urban migration in China. *Journal of Economic Geography 18*(4), 795–821.

Morlacco, M. (2017). Market power in input markets: theory and evidence from French manufacturing. Technical report.

Naidu, S., Y. Nyarko, and S.-Y. Wang (2016). Monopsony power in migrant labor markets. *Journal of Political Economy* 124(6), 1735–1792.

Nargis, N., R. Zheng, S. S. Xu, G. T. Fong, G. Feng, Y. Jiang, Y. Wang, and X. Hu (2019). Cigarette affordability in China, 2006–2015: findings from International Tobacco Control China surveys. *International Journal of Environmental Research and Public Health* 16(7), 1205.

Naughton, B. (2007). *The Chinese Economy: Transitions and Growth*. Cambridge, MA: MIT Press.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica* 69(2), 307–342.

O'Connor, R., Q. Li, W. E. Stephens, D. Hammond, T. Elton-Marshall, K. M. Cummings, G. A. Giovino, and G. T. Fong (2010). Cigarettes sold in China: Design, emissions and metals: Dataset. *Tobacco Control 19*(2), 47–53.

Olley, S. and A. Pakes (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica* 64(6), 1263–1297.

Peng, Y. (1996). The politics of tobacco: relations between farmers and local governments in China's southwest. *The China Journal 36*, 67–82.

Prager, E. and M. Schmitt (2021). Employer consolidation and wages: Evidence from hospitals. *American Economic Review 111*(2), 397–427.

Ransom, M. R. and D. P. Sims (2010). Estimating the firm's labor supply curve in a "new monopsony" framework: schoolteachers in Missouri. *Journal of Labor Economics* 28(2), 331–355.

Ravallion, M. and S. Chen (2009). *China's (uneven) progress against poverty*, pp. 65–111. Routledge.

State Administration of Taxation (2015). Annual report. Technical report, State Administration of Taxation.

State Council of the People's Republic of China (1997). Regulation for the implementation of the law on tobacco monopoly of the People's Republic of China.

State Tobacco Monopoly Administration (2002). Implementation opinions of the state tobacco monopoly administration on the organizational structure adjustment of cigarette industry enterprises in the tobacco industry.

United Nations (2019). UN COMTRADE Database. http://comtrade.un.org/.

U.S. Census Bureau (1997). Cigarette manufacturing. 1997 economic census - industry series, US Census Bureau.

Wang, J. (2013). State-Market Interaction in China's Reform Era: Local State Competition and Global Market-building in the Tobacco Industry. Routledge Contemporary China Series. Routledge.

Yang, D. T. (1999). Urban-biased policies and rising income inequality in China. American Economic Review: Papers and Proceedings 89(2), 306–310.