

Online Appendix for Market Structure, Oligopsony Power, and Productivity

Michael Rubens*

March 9, 2023

A Model with flexible markup variation

In the main text, it was assumed that manufacturing markups were constant and exogenous to the manufacturers. In this section, I relax this assumption by adding a model of leaf supply and leaf market competition to the paper.

A.1 Leaf supply model

Buyer differentiation

I rely on a discrete-choice leaf supply model with differentiated manufacturers, in an factor supply version of Berry (1994), which has been applied to labor markets by Card et al. (2018) and Azar et al. (2022). There are two reasons to model manufacturers as being differentiated, rather than homogeneous, from the farmer’s perspective. First, there is substantial cross-sectional leaf price variation within leaf markets (counties) that is explained by firm fixed effects. Regressing log leaf prices on county and year dummies yields an $R^2 = 0.669$, adding firm fixed effects increases this to $R^2 = 0.859$. In a model where manufacturers are homogeneous from the farmer’s perspective, we would expect little to no within-market leaf price variation.¹ Second, some of the observable non-price manufacturer characteristics that will be included in the model yield large and statistically significant coefficients, as discussed in Section A.3.

Farmer utility and choices

At time t , every farmer j in leaf market i makes an occupational choice $f \in \mathcal{F}_{it}^g$, with \mathcal{F}_{it}^g being the choice set. Every farmer belongs to exactly one leaf market, and are atomistically small, which is a natural assumption given the large number of tobacco farmers and small farm sizes. There are

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

¹A Cournot model of the leaf market with homogeneous manufacturers would be an example of such a model.

two nests of occupational choices g . The first nest $g = 0$ contains the outside option of not selling tobacco leaf to a manufacturer in the leaf market, meaning that $f = 0$. In Section A.2, I will explain in detail how I define and compute this outside option. The second nest contains the ‘inside good’, meaning that the farmer grows tobacco and sells it to a manufacturer $f > 0$. Farmers get a mean utility δ_{ft} from choosing option f , with the mean utility of the outside option being normalized to zero, $\delta_{0t} = 0$. Farmers obtain a nest-specific utility γ_{jg} , with the dummy d_{fg} being one if the farmer makes a choice within nest g . The advantage of putting the outside option in a separate nest is that the farmer’s substitution elasticity between choosing to be a tobacco farmer or being employed in their outside option, parametrized by the nesting parameter σ , is now different from the substitution elasticity between the different cigarette manufacturers.

Potential farmers that are employed in the outside option can enter farming. Minale (2018) finds that farmers flexibly switch between farming and non-farm work. Land ownership is quite complicated in rural China. Agricultural land is always collectively owned, and usage rights are allocated between households (Minale, 2018). Hence, as a non-farming worker who would move into tobacco farming, you would either need to be allocated a usage right, or would already possess this right through another family member.

Farmer utility depends on the leaf price W_{ft}^M , observed firm characteristics \mathbf{X}_{ft} , latent characteristics ξ_{ft} , and a firm-farmer specific utility term ν_{jft} . Further below, I explain which observable manufacturer characteristics are included in \mathbf{X}_{ft} , and why. The utility derived from the outside option is normalized to zero. I assume that both the observable and latent firm characteristics are exogenous to the manufacturers.

$$U_{jft} = \underbrace{\gamma^W W_{ft}^M + \gamma^X \mathbf{X}_{ft} + \xi_{ft}}_{\equiv \delta_{ft}} + \sum_g (d_{fg} \gamma_{jg}) + (1 - \sigma) \nu_{jft}$$

The preference shocks ν_{jft} are assumed to follow a type-I extreme value distribution. Denote the number of tobacco farmers who grow tobacco and sell to firm $f > 0$ at time t as N_{ft} . The number of potential tobacco farmers who do not farm tobacco in market i is N_{it}^0 . The combined number of actual and potential tobacco farmers in market i is $N_{it} = N_{it}^0 + \sum_{f \in \mathcal{F}_{it}^1} N_{ft}$.

The input market share of a manufacturer f is equal to $S_{ft} \equiv \frac{N_{ft}}{N_{it}}$. Assuming that farmers periodically choose which manufacturer to sell to by maximizing their static utility, the leaf market share is also equal to:

$$S_{ft} = \frac{\exp(\frac{\delta_{ft}}{1-\sigma})}{D_{gt}^\sigma [\sum_g D_{gt}^{1-\sigma}]}$$

with $D_{gt} \equiv \sum_{f \in \mathcal{F}_{it}^g} \exp(\frac{\delta_{ft}}{1-\sigma})$.

Farmers observe the price-grading schedules, but only learn the offered price when visiting a

purchasing station. In the model, I assume that farmers observe all leaf prices and manufacturer characteristics, which implies that farmers visit all purchasing stations in their county to learn about prices, or talk to each other. The assumption of linear price contracts follows from the institutional setting: leaf prices are posted as a linear price. The market share of the outside option in market i is denoted $S_{it}^0 \equiv \frac{N_{it}^0}{N_{it}}$. The market share of firm f in the inside nest of actually growing tobacco is $S_{ft}^1 \equiv \frac{N_{ft}}{N_{it} - N_{it}^0}$. The leaf supply equation that needs to be estimated is Equation (1), with $s_{ft} \equiv \ln(S_{ft})$.

$$(1) \quad s_{ft} - s_{it}^0 = \gamma^W W_{ft}^M + \sigma s_{ft}^1 + \gamma^X \mathbf{X}_{ft} + \xi_{ft}$$

The markdown expression becomes:

$$(2) \quad \psi_{ft}^M \equiv \left(\frac{\partial S_{ft}}{\partial W_{ft}^M} \frac{W_{ft}^M}{S_{ft}} \right)^{-1} + 1 = \left(\gamma^W W_{ft}^M \left(\frac{1}{1 - \sigma} - \frac{\sigma}{1 - \sigma} S_{ft}^1 - S_{ft} \right) \right)^{-1} + 1$$

Timing and information

Compared to the model in the main text, a number of additional assumptions need to be made in the leaf supply model. In time period t , the manufacturers receive their exogenous characteristics $(\xi_{ft}, \mathbf{X}_{ft})$. They make a labor choice L_{ft} , through which they control output Q_{ft} and hence the leaf price W_{ft}^M . The model assumes perfect information by both the farmers and the manufacturers, so the manufacturers know the farmers' utility function and the characteristics $(\xi_{ft}, \mathbf{X}_{ft})$ of themselves and of their competitors. In turn, farmers observe all leaf prices W_{ft}^M and manufacturer characteristics $(\xi_{ft}, \mathbf{X}_{ft})$, on the basis of which they either choose to sell their output to a manufacturer f , or to be employed in the outside option. Manufacturers transform leaf into cigarettes, which they sell to the wholesaler at a price P_{ft} . At time $t + 1$, everything restarts.

A.2 Measurement, identification, and estimation

Measurement

I define the outside option using the census of population data. As was explained above, the 'inner nest' consists of farmers who choose to grow tobacco leaf and sell it to a manufacturer. The outside option consists of potential tobacco farmers who either (i) grow another crop, rather than tobacco leaf, or (ii) choose to be employed in a non-agricultural industry. First, the cigarette quantity units in the manufacturing data need to be converted into a numbers of farmers. The quantity unit in the manufacturing data is a case of 50,000 cigarettes. In 2003, China produced 1.7 trillion cigarettes, using leaf produced by 4 million farms (Hu et al., 2006), and leaf imports

and exports were negligible. This implies around 425,000 cigarettes per farmer per year, so a case of 50,000 cigarettes corresponds to 0.117647 farmers on average. I assume this ratio is constant.² Using the notation from above, the number of tobacco farmers selling to firm f in year t is hence $N_{ft} = Q_{ft} * 0.117647$. I calculate the combined number of potential and actual tobacco farmers N_{it} to be the county population above 15 years old that is either working or looking for work. On average, the outside option has a market share S_{it}^0 of 90.7% (standard deviation is 17.5%), meaning that in the county of the average firm, 90.7% of the working population is not a tobacco farmer.

As a robustness check, I redefine the outside option differently by assuming that farmers cannot substitute between crops. This implies the outside option only includes non-agricultural workers. The preferred specification includes both farming and non-farming occupational choices in the outside option, because other studies have found that farmers switch to non-agricultural occupations in reaction to cost or demand shocks (Minale, 2018), and because farmers can presumably switch crops.

Identification

I turn to the identification of the input supply function, Equation (1). The manufacturer characteristics \mathbf{X}_{ft} are assumed to be exogenous to each individual manufacturer, but leaf prices and within-nest market shares are endogenous. The manufacturers know that the latent manufacturer characteristic ξ_{ft} affects the utility of the suppliers, which they take into account when setting leaf prices. In order to separately identify input demand from supply, at least two instrumental variables are required.

I rely on three sets of instruments. First, I rely on the manufacturing productivity estimates, ω_{ft} . As productivity enters the input demand function, it is by definition relevant. The exclusion restriction is that the productivity term does not enter the supplier utility function, meaning that it is orthogonal to the supply function residual ξ_{ft} . This instrument implies that farmers do not care about the efficiency of the manufacturing firms they are selling to, conditional on the leaf price and on observable manufacturer characteristics. Productivity differences between manufacturers can have many reasons, such as differences in managerial ability. Given that the farmers are not employed by the manufacturers, but only interact with them through monetary transactions on leaf markets, it seems reasonable to assume that the farmers do not care about how productive their buyers are, conditional on the price they receive. One threat to the validity of this assumption could be that suppliers prefer to sell repeatedly to the same buyers. This is the case in many industries that are characterized by incomplete contracts or weak contract enforceability. The literature on vertical relationships in developing countries has emphasized the importance of relational con-

²Tobacco farm yields does vary over time and across locations, but at least there seems to be not trend (Hu, 2008).

tracts and repeated interaction (Macchiavello and Morjaria, 2015). Search or switching costs on the seller side could be another driver of why repeated interaction would be valuable.³ In these cases, sellers could prefer a more productive buyer, as this reduces the risk of future buyer exit. In the Chinese tobacco industry, this is not likely to be a major concern because leaf markets do not make use of long-term contracts. Moreover, as was mentioned before, exit mainly followed from the government-induced consolidation policy, which is assumed to be exogenous to individual manufacturers. As a robustness check, I also re-estimate the model without own TFP as an instrument.

I construct a second set of instruments using county-year level weather data, which I collect from the China Meteorological Agency. Temperature and rainfall variation act as supply shifters for tobacco leaf, as warmer and more humid weather results in more and better tobacco produce (Li et al., 2021). In the spirit of Berry, Levinsohn, and Pakes (1995), weather shocks in other counties than county i act as demand shifters for county i 's cigarette manufacturers because cigarette markets are integrated beyond the county level. These weather shocks presumably do not affect leaf supply in county i , at least not when controlling for weather conditions in county i . I calculate the average maximum and minimum temperature and rainfall quantity in other counties within the same prefecture and province in each year and use these as additional instruments for leaf prices and within-nest market shares. Second, following a similar reasoning, I calculate the average productivity level of the competitors of firm f in the same county. A productivity shock to a competitor shifts leaf demand at the competitor, and hence also demand at firm f itself through the change in the equilibrium leaf price. Finally, I include the number of active manufacturers in each county as an instrument. This requires the assumption that market structure is exogenous, which was already assumed in the production model. The large observed change in market structure was a consequence of a size-dependent policy intervention, and was hence plausibly exogenous from the individual manufacturers' point of view.

Estimation

I estimate Equation (1) using two-stage least squares (2SLS) using the instruments mentioned above. I include the following manufacturer characteristics in the vector of supply shifters \mathbf{X}_{ft} . First, I control for cigarette prices and CIC codes, as they are proxies for cigarette characteristics. Second, I control for manufacturer ownership types, in order to proxy political pressure; farmers may derive a different utility from selling to manufacturers that are state-owned rather than private. Third, I include province dummies to control for geographical differences in leaf supply. Fourth,

³When applying the same model for manufacturing labor markets, more caution is needed. There are many reasons why employees would prefer to work for highly productive firms, even if these offer lower wages, such as career dynamics or better working conditions.

I control for the observed weather shocks in the own county, as these are cost shifters for the farmers. Fifth, I control for wages, as some firm-level amenities, such as its accessibility, plausibly affect both factory worker and farmer welfare. Sixth, I control for export participation: selling to exporters might entail other benefits for farmers than the offered leaf price, such as learning about international consumer preferences. I include a linear time trend to allow for changing leaf supply over time. In order to estimate the leaf supply function, the productivity residuals from the production function are needed. Therefore, I estimate the production function and leaf supply function sequentially, and bootstrap the entire estimation procedure. I block-bootstrap while resampling within firms over time, using 200 iterations.

A.3 Results

Leaf supply

The results for the nested logit model are in Table A1a when own TFP is included as an instrumental variable, and in Table A2a when not. The first column models the outside option of tobacco farming to include all industries, which is the preferred specification as explained earlier, the second column restricts the outside option to non-agricultural occupations. I start with discussing the estimates when own TFP is included as an instrument, Table A1a . The leaf price semi-elasticity is positive in both outside option specifications. The nesting parameter is close to one when all industries are included in the outside option, which indicates that the outside option is an imperfect substitute to being employed as a tobacco farmer. If only non-agricultural industries form the outside option, the nesting parameter is negative, but not significantly different from zero. In both outside option specifications, farmers prefer to sell to manufacturers who export cigarettes, to manufacturers who pay higher wages to factory workers, and to manufacturers with lower-priced cigarettes. None of the other observed characteristics is statistically significant. When not including own TFP as an instrument, Table A2a, the leaf supply estimates with all industries in the outside option are very similar to the specification that does include own TFP as an instrument, but with wider standard errors. The estimates of the model with only non-agricultural industries in the outside option have a higher inverse semi-elasticity of leaf, which implies a lower markdown.

Markdowns and markups

I continue with the preferred specification that includes all industries into the outside option. The markdown moments are in Table A1b. The markdown average and median are comparable to the levels implied by the model that imposed exogenous cigarette prices. Contrary to the main model, markups can now be estimated too, by substituting the markdown estimate into Equation (4a). The average and median markup are estimated at 0.35 and 0.64 respectively. Both markdown and

markup levels are, however, estimated imprecisely. The 5-95% confidence interval ranges from negative values to extremely high values for both markdowns and markups. Not including own TFP as an instrumental variable gives very similar markup and markdown moments, as can be seen in Table A2b, but the confidence intervals now become even wider.

Consolidation treatment effects

The estimates above show that the full model with a nested logit leaf supply side lacks the statistical power to precisely estimate markdown and markup levels. However, we are not just interested in the markdown and markup level, but mainly in how they change in response to the consolidation. Table A1c shows the difference-in-differences estimates using the logarithms of the estimated markup and markdown levels as the left-hand side variables. Markdowns are estimated to increase on average by 49%, which is a higher increase compared to the estimate from the main model. The markup is estimated to drop by 14% in response to the consolidation, but this change is not significantly different from zero. When dropping own TFP from the instrumental variables, all estimated treatment effects are very similar to the estimates when including own TFP as an instrumental variable. However, due to the loss in statistical power when excluding own TFP as an instrument, we can no longer significantly distinguish any of these changes from zero, even at the 10% confidence level. To sum up, the full model that relaxes exogenous downstream prices by adding a nested logit leaf supply side yields very similar estimates to the main model that imposed price-taking manufacturers on the cigarettes market, but is estimated less precisely.

B Data

B.1 Firm-level data

I rely on the above-scale industrial survey (ASIF) from the Chinese National Bureau of Statistics (NBS) for the firm-level data. I refer to Brandt et al. (2012) for a detailed description of this dataset. I keep all firms with Chinese Industry Classification (CIC) codes 1610, 1620 and 1690. The data were cleaned in accordance with the procedures described in Brandt et al. (2012). The NBS data also reports subsidies in all years except for 2003 and 2004. I deflate all monetary variables (profits, revenues, intermediate input expenditure, wages, and export revenue) using industry-specific output and input deflators. I rely on the updated input price deflators from Brandt et al. (2019).

B.2 Product-level data

Production quantities are recorded at the product-firm-year level by the NBS between 2000 and 2006. 99% of the observations are measured in numbers, with the unit of measurement being cases of 50,000 cigarettes. The remaining observations have tons as their unit of measurement. I re-calculate from tons to cases by using the standard weight of 1 gram per cigarette, which implies that 1 ton is equivalent to 20 cases. In 2004, the NBS changed its measurement unit from cases of 50,000 cigarettes to cases of 10,000 cigarettes for most, but not all firms. Fortunately, the quantity data contains the values of output both during the current and past calendar year. By comparing these lagged quantities with the quantities in the previous year, I bring all values before and after 2004 to the same unit of observation, which is cases of 50,000 cigarettes. Up to and including 2003, production is disaggregated into four quality grades, which I sum to the firm-year level. After 2003, only the total number of goods produced is reported. Using the lagged variables in 2000, I infer the production quantities in 1999, so product quantities are observed from 1999 until 2006.

I use the NBS firm identifiers to merge the quantity data to the ASIF balance-sheet data. I remove outliers in cigarette and leaf prices by winsorizing at the 1st and 99th percentiles, and deleted observations with negative intermediate input expenditure. I restrict the panel to 1999-2006, as quantities are not observed in 1998 and 2007. This cleaning reduces the data set to 2,025 observations, covering 470 firms over 8 years. Combining both data sets and keeping observations for which labor, capital, output, revenue, and material costs are observed reduces the sample size to 1,132 observations and 257 firms. This sub-sample covers 78% of total revenue in the raw data.

B.3 Additional data sets

I merge multiple additional data sets into the production-cost data. First, I retrieve county-level population data from the 2000 population census through the *Harvard Dataverse* (China Data

Lab, 2020), which I match based on the county identifiers.

Second, I obtain brand-level cigarette characteristics from O'Connor et al. (2010) for a subset of manufacturing firms in 2009, such as the leaf content per cigarette and other characteristics which affect the smoking experience. This data set is observed for only 13% of the observations, but covers 29% of total revenue. I link the brands in the dataset of O'Connor et al. (2010) to the manufacturers in the data set based on their names. As firm sales are not decomposed into brands, I have to aggregate from the brand to the firm-level. I do so by taking simple averages across brands.

Third, annual meteorological data on average and extreme temperatures and rainfall are obtained from the Chinese Meteorological Agency at the level of individual weather stations (China Meteorological Agency, 2018). I match these weather stations to their nearest county by minimizing the distance between the weather station and the county midpoint. I average the meteorological variables at the county-year level.

Agricultural aggregate data are obtained from the FAO statistics (FAO, 2019). Aggregate trade flows are from UN COMTRADE (United Nations, 2019). These two datasets are merely used to provided aggregate statistics on trade exposure and agricultural productivity growth.

Finally, shapefiles for China are obtained from Hijmans et al. (2004), and are used to visualize the manufacturing locations in Figure 1a in the main text.

B.4 Summary statistics

Table A3 contains a selection of summary statistics on the 1,132 firms in the cleaned dataset. The average manufacturing firm earns a revenue of \$107 million (in 1998 US dollars) and sells 355,000 cases per year. The average factory-gate price for a case of 50,000 cigarettes is \$1589, so the factory-gate price for a pack of 20 cigarettes is on average \$0.64. Using retail price data from Nargis et al. (2019), this means that factory-gate prices were on average around two thirds of retail prices, and the difference between both includes wholesale margins, retail margins, transport costs and sales taxes. The average firm made an accounting profit of \$13 million, and 10% of the firms operate at a loss. One out of four firms export, but these exports account on average for only 1% of their revenue. The average county in the dataset has a population of 570,000. The average firm employs 1204 employees, has a capital stock that is worth \$49 million, and spends \$3.5 million on wages and \$36 million on intermediate inputs.

C Robustness checks

C.1 Control function approach

In the main text, I combined the timing assumptions of Akerberg et al. (2015) with the dynamic panel approach of Blundell and Bond (2000), which relies on differencing out the persistent part

of the productivity residual. An alternative identification strategy is to rely on an inverted input demand function to control for the latent productivity scalar, with the benefit of allowing for endogenous exit and entry (Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Akerberg et al., 2015). In this section, I use this alternative identification approach.

Choice of the flexible input

Intermediate inputs are usually used as the flexible input for the first stage inversion. Oligopsony power over these intermediate inputs is problematic for the input inversion approach if intermediate inputs are substitutable with the other inputs. Persistence in oligopsony power would then induce serial correlation in intermediate input prices, which are not observed separately from input quantities, and, hence, violate the assumptions of Akerberg et al. (2015).⁴ If intermediate inputs are not substitutable and enter the production function with fixed proportions, as is the case in this paper, intermediate input prices can be backed out by taking the ratio of material expenditure over output quantities, as was explained in Appendix A. Therefore, I use tobacco leaf as the flexible input for the first-stage inversion, and include the leaf price as an argument in the first-stage regression, as explained below.

Input demand

I derive the leaf demand function in Appendix E.3. Leaf demand depends on the cigarette price, the leaf price, and all other input prices and quantities. All these variables are either observed or imputed from the data. Leaf demand also depends on the output elasticities of labor and capital, the leaf requirement per cigarette, the markup, the markdown, and on productivity, which are all latent.

$$m_{ft} = m(P_{ft}, W_{ft}^M, W_{ft}^L, L_{ft}, K_{ft}, \beta^L, \beta^K, \beta^M, \omega_{ft}, \mu_{ft}, \psi_{ft}^M)$$

The inverse leaf requirement β^M and output elasticities β^L and β^K are assumed to be the same for all firms and time periods. Markups, markdowns, and productivity vary, however, across firms and time. Without making further assumptions on the distributions of at least two of these three variables, the scalar unobservable assumption is violated, and Akerberg et al. (2015) is not identified. High input demand can be due to high productivity, low markups, and/or low markdowns.⁵

Akerberg et al. (2015) can still be used if we impose additional assumptions on the markup and markdown distributions. I impose a logit demand system for cigarettes and a nested logit supply model for tobacco leaf, as specified in Appendix A. Markdowns are still given by Equation

⁴Moreover, substitutable intermediate inputs would be subject to the identification problem for gross output production functions highlighted by Gandhi et al. (2020).

⁵Doraszelski and Jaumandreu (2021) make a similar point for markups.

(2). Denoting the price elasticity of demand as γ^P , which is assumed to be constant across firms and time, and the cigarette market share as S_{ft}^Q , markups are given by:

$$\mu_{ft} = \mu(\gamma^P, P_{ft}, S_{ft}^Q)$$

All variation in markups and markdowns is captured by the observed cigarette and leaf prices and the cigarette and leaf market shares. A similar assumption was made in De Loecker et al. (2016) for settings in which there is only imperfect competition downstream. Therefore, I include the leaf price W_{ft}^M , cigarette price P_{ft} , and market shares $s_{ft} \equiv (s_{ft}^Q, s_{ft}, s_{ft}^1)$ in the first stage regression, which is given by Equation (3). The market shares that need to be included are the product market share, which I assume to be the province-level cigarettes market share, the total leaf market share, and the within-nest leaf market share. In contrast to the main model, there is now measurement error in output ϵ_{ft}^q , which is assumed to be additive to the function $\phi_t(\cdot)$.

$$(3) \quad q_{ft} = \phi_t(\tilde{l}_{ft}, \tilde{k}_{ft}, w_{ft}^M, w_{ft}^L, p_{ft}, \mathbf{s}_{ft}) + \epsilon_{ft}^q$$

Productivity can be recovered as a function of the data and the estimable parameters, using the same functional form assumptions for the production function as were made in the main text.

$$\omega_{ft} = \hat{\phi}_{ft} - h(\tilde{l}_{ft}, \tilde{k}_{ft}, \boldsymbol{\beta}) - a(w_{ft}^L, p_{ft})$$

Exogenous exit model

I contrast two model specifications: one with exogenous exit, and one with endogenous exit. The exogenous exit model retains the AR(1) equation of motion for productivity that was used in the main text, with productivity depending on the consolidation dummies, \mathbf{Z}_{ft} , as in De Loecker (2013) and Braguinsky, Ohyama, Okazaki, and Syverson (2015).

$$\omega_{ft} = \rho\omega_{ft-1} + b(\mathbf{Z}_{ft}) + v_{ft}$$

The productivity innovation v_{ft} is given by the difference between productivity and its expected value from the equation of motion.

$$v_{ft} = \omega_{ft} - \mathbb{E}(\omega_{ft} | \omega_{ft-1})$$

Keeping the timing assumptions from the main text, the moment conditions to identify $\boldsymbol{\beta}$ are given by:

$$\mathbb{E}[v_{ft}(l_{ft-1}, k_{ft}, w_{ft}, p_{ft})] = 0$$

I estimate this model using a second-order polynomial in all the inputs for the first-stage regression $\phi(\cdot)$, and a linear function $b(\cdot)$. I block-bootstrap the standard errors with 200 iterations.

Endogenous exit model

Alternatively, I also estimate the model allowing for endogenous exit. I estimate an exit probability \hat{p}_{ft}^x using a probit regression of observed exit on all the variables in the input demand function, and add this exit probability in the first stage regression above. I allow for this exit probability to be endogenous to productivity, as proposed by Olley and Pakes (1996), and following the implementation of Collard-Wexler and De Loecker (2015). Moreover, I also allow productivity to depend on the consolidation dummies, \mathbf{Z}_{ft} , as in the exogenous exit specification.

$$\omega_{ft} = \rho\omega_{ft-1} + \tilde{b}(p_{ft-1}^x, \mathbf{Z}_{ft}) + v_{ft}$$

Similarly to above, I impose a linear function for $\tilde{b}(\cdot)$. I form the same moment conditions as were used above, and again estimate the production function by block-bootstrapping with 200 iterations.

Results

In Table A4, I present the model estimates using the control function approach outlined above. The first column assumes exogenous exit decisions, the second allows for endogenous exit. The output elasticities are in panel A4a. The output elasticity of labor and capital in both specifications of the control function approach are respectively higher and lower compared to the dynamic panel approach in the main text. The markdown moments, in panel A4(b), are slightly higher than the main specification, but are very similar across exogenous and endogenous exit specifications. The consolidation treatment effects for markdowns and productivity, in panels A4(c)-(d), are also very similar between the endogenous and exogenous exit approaches, and to those in the main text.

ACF (2015) without first stage regression

Under the assumptions that intermediate inputs are not substitutable and come in fixed proportions, that the inverse intermediate input requirement β_{ft}^M is constant across firms and time, and that there is classical measurement error in output, Akerberg et al. (2015) suggest a simpler identification strategy that does not require filtering out measurement error in the first stage. One could simply back out measurement error ϵ_{ft}^q up to a constant as follows:

$$q_{ft} - m_{ft} = \log(\beta^M) + \epsilon_{ft}^q$$

However, this approach is not possible when intermediate input prices are both latent and endogenous due to oligopsony power. Intermediate input prices are not separately identified from measurement error, and are serially correlated as soon as either the observed or latent firm amenities \mathbf{X}_{ft} or ξ_{ft} from the model in Appendix A are persistent, which is usually the case.

$$q_{ft} - m_{ft} - w_{ft}^M = \log(\beta^M) + \epsilon_{ft}^q - w_{ft}^M$$

C.2 Estimating leaf substitutability

Throughout the paper, it has been assumed that tobacco leaf cannot be substituted with labor. Tobacco leaf could be substitutable to a limited extent, for instance if hiring more workers or using better technologies leads to reduced waste.⁶ I test this by estimating the elasticity of substitution between tobacco leaf and labor. Let the cigarette production no longer take the Leontief form from Equation (2), but the following constant elasticity of substitution (CES) production function.

$$Q_{ft} = \left(\left(\beta^M M_{ft}^{\frac{\sigma^M - 1}{\sigma^M}} + \beta^L L_{ft}^{\frac{\sigma^M - 1}{\sigma^M}} \right)^{\frac{\sigma^M}{\sigma^M - 1}} \right)^{\beta^{ML}} K_{ft}^{\beta^K} \Omega_{ft}$$

The substitution elasticity σ^M parametrizes the extent to which labor and tobacco leaf can be substituted. I still impose that the elasticity of substitution between the variable inputs and capital is equal to one, which I relax in Section C.4. Following Doraszelski and Jaumandreu (2018), the first order conditions of the cost minimization problem are given by Equation (4a), which can be used to estimate the elasticity of substitution between labor and leaf. The key difference with Doraszelski and Jaumandreu (2018) is that the leaf markdown now enters the residual of this input ratio regression. Manufacturers use relatively more labor compared to tobacco leaf if wages are lower, if the output elasticity of labor compared to leaf is higher, and if manufacturers have more oligopsony power over tobacco leaf.

$$(4a) \quad l_{ft} - m_{ft} = \sigma^M (w_{ft}^M - w_{ft}^L) + \sigma^M (\ln(\beta^L) - \ln(\beta^M)) + \sigma^M \ln(\psi_{ft}^M)$$

Leaf prices can no longer be recovered from the Leontief production function, and are hence latent. Therefore, equation (4a) needs to be estimated using intermediate input expenditure $W_{ft}^M M_{ft}$, as written in Equation (4b). The only observed variable in the right-hand side of Equation (4b) is

⁶Another reason why intermediate inputs could be substitutable with labor, even if leaf and labor are non-substitutable, would be vertical integration between cigarette factories and farms. However, this is not a feature of the Chinese tobacco industry (Peng, 1996; FAO, 2003; Wang, 2013).

the log wage per worker w_{ft}^L .

$$(4b) \quad l_{ft} - (m_{ft} + w_{ft}^M) = -\sigma^M w_{ft}^L + (\sigma^M - 1)w_{ft}^M + \sigma^M (\ln(\beta^L) - \ln(\beta^M)) + \sigma^M \ln(\psi_{ft}^M)$$

The latent variation in intermediate input prices could reflect quality differences in tobacco leaf or buyer power over farmers. The input price bias that was explained earlier again applies: higher quality cigarettes are likely to require higher quality, high-wage workers, which leads to an endogeneity problem because it makes wages correlated with leaf prices. Thus, an instrument for the labor wage is needed to estimate Equation (4b). I rely on the average export share of revenue and average export participation of other cigarette manufacturers within the same prefecture as instruments for wages. Even if exporting accounts for only a small fraction of cigarette sales, shocks to international demand shift labor demand, but plausibly not labor supply. The exclusion restriction is that export participation and behavior in other manufacturing industries did not affect either leaf market oligopsony power or the production function coefficients in the cigarette manufacturing industry.

I estimate Equation (4b) at the firm-year level using two-stage least squares, with the instruments for the log wage that were mentioned above. I control for the firm's ownership type and CIC code, year dummies, the firm's export status, and the log cigarette price of the firm, in order to control for possible confounding variables that affect both input usage ratios and wages. The estimated elasticity of substitution between labor and tobacco leaf is in the first column of Table A5a, and is very close to the value of 0 assumed in the main model, even if it is estimated imprecisely.⁷

C.3 Comparison to the substitutable leaf model

The main methodological point in the paper was that markdowns and markups are no longer separately identified using the production approach if there are non-substitutable inputs. In this section, I examine to which extent this matters for correct inference about markdown and productivity levels and changes.

Substitutable leaf production function

The non-substitutability of tobacco leaf in the main model was supported by the industry background and the leaf substitution elasticity estimates above. Suppose we would have assumed tobacco leaf to be substitutable, how would this affect the results? Keeping the Cobb-Douglas functional form, the production function is now given by the gross output function in Equation

⁷The elasticity of substitution that is close to zero contrasts with some of the prior literature, such as Sumner and Alston (1987), but these approaches use an input demand approach which rules out oligopsony power on factor markets.

(5a), with the output elasticities of materials, labor and capital being β^{M*} , β^{L*} , and β^{K*} .

$$(5a) \quad q_{ft} = \beta^{L*}l_{ft} + \beta^{K*}k_{ft} + \beta^{M*}m_{ft} + \omega_{ft}$$

However, given that input prices and quantities are not separately observed in the substitutable inputs case, the estimable production function is not Equation (5a), but Equation (5b), with material expenditure on the right-hand side.

$$(5b) \quad q_{ft} = \beta^{L*}l_{ft} + \beta^{K*}k_{ft} + \beta^{M*}(m_{ft} + w_{ft}^M) + \underbrace{\omega_{ft} - \beta^{M*}w_{ft}^M}_{\hat{\omega}_{ft}}$$

If all inputs are substitutable, and maintaining the assumption that labor wages are exogenous to individual firms, the substitutable leaf markup estimate μ^* is given by $\mu_{ft}^* = \frac{\beta^{L*}}{\alpha_{ft}^L}$, using Equation (4b). The estimated leaf price markdown ψ_{ft}^{M*} in the substitutable leaf model is equal to the ratio of the markup of the variable of which the price is endogenous over the markup of the input of which the price is exogenous (Morlacco, 2017):

$$\psi_{ft}^{M*} = \frac{\beta^{M*}}{\beta^{L*}} \frac{\alpha_{ft}^L}{\alpha_{ft}^M}$$

Estimated markdown levels

I estimate the substitutable leaf model using the same dynamic panel identification approach as used in the main text and the same price controls in the production function, but now with three rather than two inputs. The estimates are reported in the first column of Table A6a. Both the output elasticity of labor and capital are lower compared to the non-substitutable leaf model. Table A6b shows that the average markup is estimated to be 5.26 and the median markup at 3.81, whereas the average markdown ratio is estimated to be 0.89 and the median markdown is merely 0.34. Both markups and markdowns are imprecisely estimated using the substitutable leaf model. Nevertheless, these point estimates suggest that cigarette manufacturers do not have market power on leaf markets, but high market power on the cigarette market, the opposite conclusion from the Leontief model. The substitutable leaf model results are implausible, given that manufacturers are faced with many small farmers upstream, but with a monopsonistic wholesaler downstream. This shows that the substitutability assumption imposed for intermediate inputs has important consequences for inference about markups and markdowns.

The reason for these biased markdown and markup estimates when assuming substitutable leaf is as follows. If the Leontief production function is the true model with true labor output elasticity

β^L , the estimated markup in the substitutable leaf model, μ_{ft}^* , is an overestimate of the true markup μ_{ft} if:

$$\frac{\beta^{L*}}{\alpha^L} > \frac{\beta^L}{\alpha^L + \beta^L \psi_{ft}^M \alpha_{ft}^M}$$

If the estimated output elasticity of labor would be the same for the substitutable and non-substitutable leaf model, then the markup from the substitutable leaf model always overestimates the markup from the Leontief model. The reason is that this model does not take into account that the marginal cost of labor also depends on the input price elasticity of materials due to the complementarity between labor and materials. Hence, if we assume that the output elasticity of labor is the same in both specifications, $\beta^L = \beta^{L*}$, marginal costs are underestimated and markups overestimated.⁸ This is also the case in the context of this paper: the substitutable leaf model seems to overestimate cigarette price markups and underestimate leaf price markdowns.

Estimated markdown and productivity changes

The estimated effects of the consolidation on markdowns and productivity in the substitutable leaf model are in the first column of Table A6(c). Markdowns are now estimated to have increased by 90% in response to the consolidation, which is a much larger change compared to the Leontief model in the main text. Interestingly, total factor productivity is estimated to have *increased* by 41% in the substitutable leaf model, whereas it was estimated to fall in the main model.

This difference is due to the fact that input prices and quantities are not observed separately. As is shown in Equation (5b), the estimated productivity residual $\hat{\omega}_{ft}$ now includes leaf prices in addition to the true TFP level ω_{ft} . A drop in latent intermediate input prices due to increased oligopsony power will hence be interpreted as rising productivity in the substitutable leaf model.⁹ In the Leontief model, intermediate inputs do not enter the estimated production function, and hence unobserved leaf prices do not enter the productivity residual. Prior work on SOE privatization and consolidation policies found that they led to large increases in profitability (Brown et al., 2006; Hsieh and Song, 2015; Chen et al., 2021). These profitability gains could be due to both increased oligopsony power or actual TFP growth.

⁸In the application above, the estimated output elasticity of labor in the substitutable leaf model, β^{L*} , diminishes compared to the Leontief model, which counteracts the positive bias on the markup estimate, but does not neutralize it.

⁹De Loecker et al. (2016) discussed how unobserved input quantities led to biased production function coefficients when inputs differ in terms of quality. The source of bias in this paper is, in contrast, oligopsony power rather than input quality variation.

C.4 Labor-augmenting productivity

The productivity residual ω_{ft} was assumed to be Hicks-neutral throughout the paper, and capital and labor to be Cobb-Douglas substitutable. However, there could be unobserved heterogeneity in the output elasticities of labor and capital due to labor-augmenting productivity, and labor and capital may have a different substitution elasticity. In order to examine these two issues, I re-define the production function to a CES form in Equation (6). Labor and capital are substitutable at a rate σ^K , and the parameters β_{ft}^L and β_{ft}^K vary across firms and time in order to capture labor-augmenting productivity differences.

$$(6) \quad Q_{ft} = \min \left\{ \left((\beta_{ft}^K K_{ft}^{\frac{\sigma^K-1}{\sigma^K}} + \beta_{ft}^L L_{ft}^{\frac{\sigma^K-1}{\sigma^K}})^{\frac{\sigma^K}{\sigma^K-1}} \Omega_{ft} \right), \beta_{ft}^M M_{ft} \right\}$$

Capital substitutability

Similarly to the previous section, and still following Doraszelski and Jaumandreu (2018), the first order conditions of the cost minimization problem can be written as Equation (7), which can be used to estimate the elasticity of substitution between labor and capital. There are two differences with Equation (4b), which was used to estimate the elasticity of substitution between leaf and labor. First, the capital market is assumed to be perfectly competitive, so there is no markdown over capital prices that enters the residual. Second, there is heterogeneity in the output elasticities of labor and capital, which reflects differences in labor-augmenting productivity.

$$(7) \quad l_{ft} - (k_{ft} + w_{ft}^K) = -\sigma^K w_{ft}^L + (\sigma^K - 1)w_{ft}^K - \sigma^K (\ln(\beta_{ft}^K) - \ln(\beta_{ft}^L))$$

The only observable variable on the right-hand side of Equation (7) is the wage, all the other parameters are unknown and enter the residual. This introduces endogeneity bias, as labor and capital choices of firms are a function of their factor-augmenting productivity levels. I estimate Equation (7) using the same BLP instruments for log wages and control variables that were mentioned in Section C.2. The estimated elasticity of substitution between labor and capital is in the right column of Table A5a, and it is estimated to be 0.915. This is very close to the value of one imposed in the main model.

Directed technical change

Next, I test whether the consolidation induced factor-biased technical change. If the factor-augmenting productivity levels β_{ft}^L and β_{ft}^K changed in response to the consolidation, for instance because consolidated firms upgraded their production technology, this would threaten the interpretation of the

markup and markdown results. From Equation (7), it is clear that the capital-labor ratio would then have to change in response to the consolidation. I test this by estimating the difference-in-differences Equation (1) with the log of the capital stock per employee, $k_{ft} + w_{ft}^K - l_{ft}$, as the left-hand side variable. The results are in Table A5b. The change in the capital stock per employee was not significantly different between the firms in the treatment and the control group.

C.5 Translog production function

The labor-capital substitution elasticity that was estimated in Appendix C.4 confirms the Cobb-Douglas assumption for the labor-capital term $H(\cdot)$ in the production function. Nevertheless, I use a more flexible translog specification for $H(\cdot)$ as an additional robustness check. The corresponding functional form of $h(\cdot)$ in logarithms is given by:

$$h(L_{ft}, K_{ft}) = \beta^L l_{ft} + \beta^K k_{ft} + \beta^{LK} l_{ft} k_{ft} + \beta^{2L} l_{ft}^2 + \beta^{2K} k_{ft}^2$$

I use the same identification approach as in the main specification. The moment conditions from Equation (7) are now adapted to:

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

I still include the instruments only up to one time lag, as in the main model. The resulting estimates are in the first column of Table A7. The output elasticities and markdowns are very similar to the estimates in the main text, except that the average markdown now has a much higher standard error. The estimated treatment effects for both markdowns and productivity are very similar to the Cobb-Douglas results in the main text.

C.6 Cost shares approach

As a robustness check with respect to the production function identification strategy, I rely on a cost shares approach similar to the one proposed by Foster, Haltiwanger, and Syverson (2008). Assuming labor and capital are both variable inputs, and assuming constant returns to scale, the output elasticity of labor is equal to its cost share of labor and capital costs: $\beta_{ft}^l = \frac{W_{ft}^L L_{ft}}{W_{ft}^L L_{ft} + I_{ft}^K}$, with capital expenditure being denoted I_{ft}^K . In order to take measurement error into account, I estimate the output elasticity of labor as the yearly median cost share of labor, and analogously for capital. The resulting output elasticities are in the second column of Table A7. Compared to the main model, the capital coefficient is smaller. The estimated markdown levels and treatment effects for both markdowns and productivity are very similar to those in the main model.

C.7 Other robustness checks

Alternative leaf market definitions

Defining leaf markets is crucial to know which firms are subject to the consolidation and which are not. In the main text, leaf markets were defined at the county level, which is in line with the legal restrictions on tobacco leaf trade that were discussed in Section 1 of the main text. In Table A8a, I show the estimated consolidation treatment effects when leaf markets are re-defined at the province and prefecture levels, which are both wider than the county-level market definition used in the main text. The markdown treatment effects decrease as leaf markets are defined to be larger. Some firms are now defined to be in the treatment group, although they are in reality not affected by the exit of small manufacturers in the wider market. Whereas markdowns are estimated to increase by 37% due to the consolidation when using county-level markets, this effect diminishes to 24% using prefectural market definitions, and 9% using provincial leaf market definitions.

Keeping only large firms in sample

The difference-in-differences model contains firms that produce below the exit threshold of 100,000 cases per year in the pre-treatment period. As was explained earlier, a small subset of SOEs producing below the exit threshold should have exited after the consolidation, but managed to stay in the market. The main model kept these firms in the sample, but one could also exclude them and only focus on the large firms. When doing so, the treatment effect on markdowns remains very similar, as can be seen in Table A8b.

Different calibrated markups.

In the main text, it was assumed that the monopsonistic wholesaler sets cigarette prices equal to manufacturer marginal costs, which implies $\mu = 1$. In this Appendix, I examine the robustness of the results under different calibrated markup values. In particular, I examine two markup values of interest. First, I set the markup equal to the median markup found in the extended model with a nested logit leaf supply side, which is $\mu = 0.644$. Second, I set the markup equal to the upper bound of the 95% confidence interval for the median markup in Appendix A, which is $\mu = 1.502$. This results in a lower bound for the leaf price markdown.

The resulting markdown levels and consolidation treatment effects are in Table A9. The average markdown increases to 4.813 when taking the markup to be the median markup in the nested logit model $\mu = 0.644$, and decreases to 1.750 when the markup is equal to $\mu = 1.502$. In both cases, the average markdown is significantly above one. The consolidation treatment effects are very similar across markup calibrations: markdowns are estimated to increase by 36.9% on aver-

age if $\mu = 0.644$ and by 29.3% if $\mu = 1.502$, compared to 37.0% under the main specification of $\mu = 1$. The markdown changes in response to the consolidation are statistically significant for any of these markup calibrations.

D Quality

Both cigarettes and tobacco leaf are differentiated products. In the main text, two assumptions were made with respect to this differentiation: (i) the leaf content of cigarettes is constant across firms and over time, and (ii) cigarette quality does not change within firms over time. In this appendix, I rely on additional information about cigarette characteristics to defend the first assumption, in Section D.1, and the second assumption, in Section D.2. I also discuss how violations of these assumptions would affect the level and changes of the estimated leaf price markdowns.

D.1 Leaf content of cigarettes

Given that intermediate input usage consists mostly of tobacco leaf, the parameter $\beta_{ft}^M \equiv \frac{Q_{ft}}{M_{ft}}$ is interpreted as the inverse of the leaf content per cigarette. The homogeneous leaf requirement assumption, $\beta_{ft}^M = \beta^M$, was defended by pointing to low cross-sectional leaf content heterogeneity. As shown in Appendix Table A3, there is little cross-sectional heterogeneity in the leaf content per cigarette for a sample of cigarette brands from 2005-2007, which covered 30% of the market in those years. This data was collected by O'Connor et al. (2010). The average tobacco content per cigarette was 686 mg per 1000 mg, with a standard deviation of 30 mg. The leaf content also did not differ significantly between the control and treated groups. Interestingly, O'Connor et al. (2010) also took repeated samples of cigarettes between 2005 and 2007, and found very little change in leaf contents over time. Admittedly, both these sample were obtained *after* the large consolidation wave of 2003. However, there was still some consolidation after 2004 through mergers of the medium-sized firms (producing between 100,000 and 300,000 cases per year), so the limited cross-sectional and time-series variation in leaf contents per cigarette is reassuring.

Of course, there is still *some* variation in leaf contents per cigarette; the standard deviation of the leaf content is not exactly zero. Any such variation will feed into the estimated markdown level. If a firm has a below-average leaf content per cigarette, $\frac{M_{ft}}{Q_{ft}} < \frac{1}{\beta^M}$, this results in a lower expenditure share of tobacco leaf, and hence a higher estimate for the leaf price markdown ψ_{ft}^M . However, as long as this variation in leaf contents is exogenous with respect to the consolidation, meaning that manufacturers do not adjust their leaf contents in response to the change in market structure, the reported treatment effects should still be correct. The fact that leaf contents did not vary much throughout the post-2003 consolidation episodes supports this assumption. In addition, I compare cigarette leaf contents across leaf markets, by regressing the average leaf content per cigarette in each leaf market on the number of cigarette manufacturers per leaf market, in Table

A8c. This yields very small and insignificant coefficients, so more concentrated markets do not differ significantly in terms of leaf contents from less concentrated markets.

D.2 Cigarette differentiation

Cigarettes are differentiated along other dimensions than their leaf content. In particular, vertical differentiation bears some implications for the empirical model. Producing high-quality cigarettes requires high-quality tobacco leaf (Li et al., 2021). The implications of this differentiation for the identification of the production function and of leaf price markdowns were discussed in the main text. The key assumption here is that cigarette quality can differ across firms, but is not allowed to change within firms over time. If quality would vary over time, leaf quality differences feed into the estimated markdown level, and are not captured by firm fixed effects. A similar reasoning applies as in the previous section: a decrease in leaf quality results in lower leaf prices, and hence lower leaf expenditure. To the extent that such a quality decrease is not entirely passed through to cigarette prices, this will result in a higher leaf markdown estimate, even if buyer power might have stayed constant. This is mainly a problem if such quality changes were the result of the ownership consolidation: the increase in leaf price markdowns documented after the consolidation could be due to decreased leaf quality as well. However, the broad consensus about Chinese cigarettes is that their quality, if anything, increased rather than decreased over time, and mainly from 2009 onwards, which is after the time period studied in this paper (Bai, 2019; Xu et al., 2019). In the next two paragraphs, I collect additional micro-data to verify the constant quality assumption.

Observed cigarette characteristics

First, I collect data on observed cigarette characteristics in the secondary literature. As already mentioned, O'Connor et al. (2010) took repeated samples of cigarette characteristics between 2005 and 2007. Besides the leaf content of cigarettes, they also surveyed other cigarette characteristics that correlate with quality (such as the cigarette ventilation rate) and found very little time series variation in these characteristics. Similar findings can be found in O'Connor et al. (2015), who report that Chinese cigarette characteristics barely changed between 2009 and 2012.

Cigarette quality

Even if other studies point to low cross-sectional and time-series variation in observed cigarette characteristics, there could still be a worry of unobserved cigarette (and hence leaf) quality. There are two ways to infer cigarette quality (and hence leaf quality) from the data. First, the government provides a subsidy to firms who produce cigarettes of the lowest quality grades (Li et al., 2010). I observe these subsidies in 1999-2002 and 2005-2006. The subsidy is on average 1.3% of revenue,

and on average 20% of firms receive it. I use a dummy of whether the firm receives a subsidy and the subsidy as a share of revenue as indicators for low cigarette quality. It has to be noted that the presence of this subsidy affects the size decision of the firms, altering the first order condition of the cost minimization problem. Second, a breakdown of cigarette production into four quality grades is available in the product-level dataset between 1999 and 2003.

In Table A10a, I verify whether cigarette quality changed in response to the consolidation using the same difference-in-differences model as before. The first column uses an indicator of whether the firm produces any cigarettes of the two lowest quality grades as the left-hand side variable. The second column uses an indicator of whether the firm received any subsidies, which are given to low-quality cigarette producers, in a given year. The third column uses the subsidy share of revenue as the left-hand side variable. Although the share of firms that received a subsidy fell, the share of low-quality cigarettes and the subsidy/revenue share remained stable throughout the consolidation, which suggests that cigarette and leaf quality did not change much.

Second, I examine whether latent quality differences across firms affect the estimated markdown treatment effects from the main text in a meaningful way. Controlling for the direct quality measure in the difference-in-differences model is difficult because it is observed only until 2003, when the consolidation policy was still being carried out. The subsidy measure has, however, the benefit of being observed in 2005 and 2006. In Table A10b, I test the robustness of the markdown consolidation treatment effect when controlling for subsidies as a right-hand side variable. I compare the main specification for the difference-in-differences model without control variables, but on the same sample on which subsidies are observed (excluding 2003 and 2004) to the specification that controls for the subsidy dummy and the subsidy share of revenue. The results show that the markdown increase is very similar across these three specifications. Unobserved leaf quality changes hence seem not to be a confounding factor of the effect of the consolidation on leaf price markdowns.

E Additional results and derivations

E.1 Additional results

Measurement error

The model can be extended to allow for measurement error in output. Let the production function be re-defined as follows, with the log of measurement error being ϵ_{ft}^q . I assume that this measurement error is i.i.d. distributed across firms and over time.

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

The production function in logs becomes:

$$q_{ft} = h(l_{ft}, k_{ft}, \boldsymbol{\beta}) + \omega_{ft} + \epsilon_{ft}^q$$

Rewriting the production function to take into account unobserved input quality differences gives:

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

The moment condition using the dynamic panel approach becomes:

$$\mathbb{E} \left[v_{ft} + \epsilon_{ft}^q - \rho \epsilon_{ft-1}^q | (\tilde{l}_{fr-1}, \tilde{k}_{fr}, p_{fr}, w_{fr}^L, \mathbf{Z}_{fr}) \right]_{r \in [2, \dots, t]} = 0$$

Using the dynamic panel approach, measurement error ϵ_{ft}^q and total factor productivity ω_{ft} are not separately identified. If there would be measurement error in output, all statements in the text about productivity would apply to the sum of log productivity and log measurement error. Furthermore, leaf prices are observed with error if quantities are observed with error: $W_{ft}^M = \frac{W_{ft}^M M_{ft}}{Q_{ft}} \beta^M \exp(\epsilon_{ft}^q)$. Lastly, the productivity decomposition exercise also contains measurement error in output:

$$\hat{\Omega}_{it} \equiv \sum_{f \in \mathcal{F}_{it}} \left(\frac{\Omega_{ft} \exp(\epsilon_{ft}^q)}{F_{it}} \right)$$

$$\bar{\Omega}_{it} \equiv \sum_{f \in \mathcal{F}_{it}} \left(\frac{\Omega_{ft} \exp(\epsilon_{ft}^q) \tilde{L}_{ft}}{\sum_{f \in \mathcal{F}_{it}} (\tilde{L}_{ft})} \right)$$

In the ‘control function approach’ with input inversion, it is possible to separately identify measurement error in output from total factor productivity. As was shown in Appendix C.1, this yields very similar results compared to the dynamic panel approach without measurement error.

Market structure and leaf prices

In Table A12, I regress log leaf prices $\log(W_{ft}^M)$ on dummies that indicate the presence of one or two manufacturing firms at the province-, prefecture-, and county- level.¹⁰ Leaf prices are lower in more concentrated markets at any leaf market definition, and this relationship becomes stronger the narrower leaf markets are defined.

¹⁰In a previous version of the paper, I included dummies up to three firms per market, but this leads to collinearity at the county-level as less than 1% of counties have three or more cigarette manufacturers.

Markdown correlations

Which market and firm characteristics explain markdown variation? In Table A11, I regress the log leaf price markdown ratio on an SOE indicator, log firm revenue, the log value-added tax rate, the log unemployment rate in the county, and an indicator for the county containing only one or two cigarette manufacturers. I block-bootstrap standard errors with 200 iterations. Leaf markdowns are larger in counties with one or two cigarette firms, which is consistent with the classical oligopsony model. Markdown ratios are also higher at larger firms, and at firms with a higher value-added tax rate. Markdowns are lower at state-owned enterprises, and do not correlate significantly with the county’s unemployment rate.

Pre-trends

The difference-in-differences model estimated in the main text relied on Equation (1), which only interacted the treatment group indicator Z_f with one time dummy, $\mathbb{I}[t \geq 2002]$, indicating the post-2002 time period. In order to visually inspect the parallel pre-trends between the treatment and control groups, I estimate an alternative equation that interacts the treatment indicator with all year fixed effects, Equation (8).

$$(8) \quad y_{ft} = \theta_0 + \sum_{\tau=2000}^{2006} \left(\theta_1^\tau \mathbb{I}[t = \tau] + \theta_2^\tau Z_f \mathbb{I}[t = \tau] \right) + \theta_f + \varepsilon_{ft}$$

I estimate the fitted log markdown and log productivity levels for each year, for both the treatment and control group using Equation (8), and plot these in Figure A2 to examine the pre-trends. Whereas log markdowns have parallel trends prior to 2002, this is not the case for log productivity. Productivity was declining for the unconsolidated markets prior to 2002, which is to be expected because the policy objective was exactly to increase low productivity growth by small manufacturers.

E.2 Non-cost minimizing firms

Assumption 3 stated that manufacturers minimize their per-period variable costs. As was discussed earlier, various industry sources confirm that cigarette manufacturers compete against each other on their input markets and have incentives to lower their costs. Non-profit motives of Chinese SOEs are discussed by, among others, Li et al. (2012) and Chen et al. (2021). Chinese firms, and especially those that are state-owned, may incorporate goals such as ‘achieving social stability’ through high and countercyclical employment in their objective function. In this section, I discuss how such deviations from cost minimization affect the markdown estimates.

Suppose manufacturers specifically want to employ a lot of farmers. The objective function becomes Equation (9), with the parameter ς_{ft} indicating the importance given to farmer employment. The higher ς_{ft} becomes, the more the firm values farming employment. If $\varsigma_{ft} = 1$, it does not value farming employment, but simply minimizes costs.

$$(9) \quad \min_{L_{ft}, K_{ft}} \left\{ \left(\frac{W_{ft}^M}{\varsigma_{ft}} M_{ft} + W_{ft}^L L_{ft} - \lambda_{ft} (Q_{ft} - Q(M_{ft}, L_{ft}, K_{ft}, \boldsymbol{\beta})) \right) \right\}$$

The true markdown ψ_{ft}^M is no longer given by Equation (4a), but by the following expression:

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

Given that ς is a latent parameter, the estimated markdown $\hat{\psi}_{ft}^M$ is:

$$\hat{\psi}_{ft}^M = \frac{1}{\alpha_{ft}^M} \left(\frac{1}{\mu} - \frac{\alpha_{ft}^L}{\beta^L} \right)$$

The estimates of β^L are the same in both expressions, because the production function estimation procedure does not rely on cost minimization. The revenue shares $\alpha_{ft}^L, \alpha_{ft}^M$ are observed, and hence also the same in both expressions. It follows that the estimated markdown is an underestimate of the true markdown if and only if $\varsigma_{ft} > 1$.

$$\hat{\psi}_{ft}^M < \psi_{ft}^M \Leftrightarrow \varsigma_{ft} > 1$$

The reason for this is that if $\varsigma_{ft} > 1$, we interpret high leaf usage (in the form of a high cost share of leaf) as evidence for low oligopsony power over farmers, whereas firms in reality have a preference for employing many farmers, and are not fully exerting their oligopsony power to minimize costs.

This bias in the markdown level from non-cost minimizing behavior is not necessarily problematic if we are strictly interested in consolidation treatment effects, because of two reasons. First, it is likely that if firms have non-cost minimizing objectives, they are similar across firms; 98% of sales in the market are under some type of state control, which is the likely source of non-cost minimizing objectives. This makes it likely that the objective parameter ς_{ft} is homogeneous across firms: $\varsigma_{ft} = \varsigma$. In this case, the implied markdown level would be wrong, but the relative treatment effect between the control and treatment group should still be valid. Second, even if manufacturers differ in their objectives, it is likely that these objectives are stable over time, being a preference parameter. In that case, the firm objective ς_f would be absorbed in the manufacturer fixed effects in the difference-in-differences model, and the estimated consolidation effects remain unbiased.

E.3 Derivations

Markup and markdown expressions

In this section, I derive the markup formula of Equation (4a). Taking the first order derivative of variable costs results in the following expression for marginal costs MC_{ft} :

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

Substituting the revenue shares $\alpha_{ft}^V \equiv \frac{V_{ft} W_{ft}^V}{P_{ft} Q_{ft}}$ for $V \in \{L, M\}$ and $\beta_{ft}^L \equiv \frac{\partial Q_{ft}}{\partial L_{ft}} \frac{L_{ft}}{Q_{ft}}$ gives:

$$MC_{ft} = \frac{\alpha_{ft}^L P_{ft}}{\beta_{ft}^L} \psi_{ft}^L + \alpha_{ft}^M P_{ft} \psi_{ft}^M$$

Finally, dividing prices by marginal costs yields Equation (4a).

Input demand function

I now derive the leaf demand function, which is used when identifying the production function using Akerberg et al. (2015) in Appendix C.1. Solving the first order conditions from the cost minimization problem, in Equation (3), with a Cobb-Douglas function for $H(\cdot)$ gives the following optimal output level Q_{ft}^* :

$$Q_{ft}^* = \left[\left(\frac{P_{ft}}{\mu_{ft}} - \frac{W_{ft}^M}{\beta_{ft}^M} \psi_{ft}^M \right) \beta_{ft}^L \omega_{ft}^{\frac{1}{\beta_{ft}^L}} K_{ft}^{\frac{\beta_{ft}^K}{\beta_{ft}^L}} \frac{1}{W_{ft}^L} \right]^{\frac{\beta_{ft}^L}{1-\beta_{ft}^L}}$$

From the production function, it can be seen that the optimal leaf level is equal to $M_{ft}^* = \frac{Q_{ft}^*}{\beta_{ft}^M}$. In short notation, intermediate input demand is a function of cigarette and input prices, capital, the output elasticities of labor and capital, and the leaf input requirement, of total factor productivity, the markup, and the markdown.

$$m_{ft} = m(P_{ft}, W_{ft}^M, W_{ft}^L, K_{ft}, \beta_{ft}^L, \beta_{ft}^K, \beta_{ft}^M, \omega_{ft}, \mu_{ft}, \psi_{ft}^M)$$

Table A1: Full model with markup variation

<i>panel A: Leaf supply estimates</i>	Outside option = all industries		Outside option = non-agricultural	
	Est.	S.E.	Est.	S.E.
Leaf price	0.674	0.262	0.475	0.283
Nesting parameter	0.288	0.422	-0.403	0.392
1(Exporter)	1.330	0.336	0.775	0.366
Export share of revenue	-1.411	6.140	-4.749	5.527
Wage	1.428	0.238	1.085	0.229
1(SOE)	1.365	0.777	1.197	0.783
1(Collective firm)	-1.002	1.216	-1.450	1.006
Cigarette price	-1.956	0.550	-1.683	0.682
Year	0.021	0.049	0.024	0.076
1st stage F-stat (leaf price)		29.11		15.48
1st stage F-stat (market share)		7.93		27.83
R-squared		0.30		0.25
Observations		625		591
<i>panel B: Markdown, markup</i>	Markdown		Markup	
	Average	Median	Average	Median
Moment	2.842	3.388	0.507	0.644
5%-95% C.I.	[-1.46 , 5.233]	[-1.27 , 5.294]	[-1.22 , 1.921]	[-.152 , 1.502]
<i>panel C: Consolidation effects</i>	log(Markdown)		log(Markup)	
	Est.	S.E.	Est.	S.E.
1(Year \geq 2002)*1(treatment)	0.391	0.210	-0.144	0.199
R-squared		0.78		0.72
Observations		919		918

Notes: Panel (a) reports the nested logit leaf supply estimates. In column (I), the outside option of tobacco farmers is to work in any other industry, in column (II) the outside option is restricted to non-agricultural industries. Panel (b) reports the corresponding markdown and markup moments. Average markups and markdowns are weighted by input usage. Panel (c) re-estimates the consolidation treatment effects for both markdowns and markups All standard errors are block-bootstrapped using 200 iterations.

Table A2: Full model with markup variation, not using own TFP as an instrument

<i>panel A: Leaf supply estimates</i>	Outside option = all industries		Outside option = non-agricultural	
	Est.	S.E.	Est.	S.E.
Leaf price	0.565	0.330	0.157	0.404
Nesting parameter	0.375	0.409	-0.134	0.389
1(Exporter)	1.252	0.352	0.538	0.393
Export share of revenue	-1.558	5.961	-5.093	6.752
Wage	1.363	0.266	0.894	0.306
1(SOE)	1.140	0.818	0.531	0.981
1(Collective firm)	-0.934	1.109	-1.263	1.040
Cigarette price	-1.699	0.696	-0.890	0.945
Year	0.012	0.049	-0.021	0.084
1st stage F-stat (leaf price)		28.25		27.36
1st stage F-stat (market share)		8.11		20.77
R-squared		0.40		0.55
Observations		625		591
<i>panel B: Markdown, markup</i>	Markdown		Markup	
	Average	Median	Average	Median
Moment	3.267	3.842	0.463	0.576
5%-95% C.I.	[-6.40 , 6.724]	[-2.23 , 7.093]	[-.941 , 1.630]	[-.210 , 1.529]
<i>panel C: Consolidation effects</i>	log(Markdown)		log(Markup)	
	Est.	S.E.	Est.	S.E.
1(Year \geq 2002)*1(treatment)	0.411	0.234	-0.159	0.229
R-squared		0.79		0.73
Observations		919		918

Notes: Panel (a) reports the nested logit leaf supply estimates. In column (I), the outside option of tobacco farmers is to work in any other industry, in column (II) the outside option is restricted to non-agricultural industries. Panel (b) reports the corresponding markdown and markup moments. Average markups and markdowns are weighted by input usage. Panel (c) re-estimates the consolidation treatment effects for both markdowns and markups. All standard errors are block-bootstrapped using 200 iterations. Own TFP is not used as an instrument.

Table A3: Summary statistics

	Mean	S.D.	Obs.
Revenue (million USD)	106.86	200.28	1132
Quantity (thousand cases)	354.96	446.46	1132
Price per case (USD)	1588.52	13925.49	1132
Profit (million USD)	12.53	45.55	1132
Wage bill (million USD)	3.51	6.12	1132
Material expenditure (million USD)	36.23	52.75	1132
Capital stock (million USD)	48.55	73.46	1132
Employees	1197.70	1080.60	1132
Export dummy	0.23	0.42	1132
Export share of revenue	0.01	0.05	1132
County population (millions)	0.57	0.36	932
Leaf content per cigarette (mg)	686.37	30.49	110
Filter density (mg/ml)	112.71	3.79	110

Notes: A case contains 50,000 cigarette sticks. Prices are factory-gate prices. Revenue, prices, profits, and input expenditure are denoted in 1998 US dollars.

Table A4: Control function approach

<i>panel A: Production function</i>	Exogenous exit		Endogenous exit	
	Est.	S.E.	Est.	S.E.
Output elasticity of labor	0.720	0.460	0.859	0.256
Output elasticity of capital	0.486	0.265	0.267	0.203
Scale parameter	1.206	0.214	1.127	0.184
R-squared	0.96		0.96	
Observations	923		923	
<i>panel B: Markdown moments</i>	Exogenous exit		Endogenous exit	
	Est.	S.E.	Est.	S.E.
Average markdown	3.941	0.494	4.090	0.992
Median markdown	2.078	0.081	2.110	0.186
<i>panel C: Treatment effects on markdowns</i>	Exogenous exit		Endogenous exit	
	Est.	S.E.	Est.	S.E.
Treatment*1(Year \geq 2002)	0.303	0.105	0.317	0.106
R-squared	0.70		0.73	
Observations	1127		1130	
<i>panel D: Treatment effects on TFP</i>	Exogenous exit		Endogenous exit	
	Est.	S.E.	Est.	S.E.
Treatment*1(Year \geq 2002)	-0.059	0.081	-0.061	0.082
R-squared	0.89		0.91	
Observations	1132		1132	

Notes: Panel (a) reports the production function estimates using the control function approach of Akerberg et al. (2015). The first column assumes exit is exogenous to firm productivity, the second column allows for endogenous exit. Panels (b)-(d) report the corresponding markdown moments and treatment effects. All standard errors are block-bootstrapped with 200 iterations.

Table A5: Input substitutability

<i>panel A: Elasticity of substitution</i>	Labor and leaf		Labor and capital	
	Est.	S.E.	Est.	S.E.
Elasticity of substitution	0.010	0.288	0.915	0.209
1st stage F-stat	38.52		38.52	
R-squared	0.24		0.48	
Observations	1129		1129	
<i>panel B: Capital intensity changes</i>	log(Capital/labor)			
	Est.	S.E.		
1(Treatment)*1(year>2002)	0.029	0.082		
R-squared	0.89			
Observations	1915			

Notes: Panel (a) reports the estimated elasticity of substitution between labor and leaf and between labor and capital in the CES model. Panel (b) reports how the log capital-to-labor ratio changed in response to the consolidation.

Table A6: Substitutable leaf model

<i>panel A: Production function</i>	Substitutable leaf			
	Est.	S.E.		
Output elasticity of labor	0.174	0.126		
Output elasticity of capital	0.381	0.122		
Output elasticity of materials	0.529	0.095		
Scale parameter	1.084	0.260		
R-squared	0.96			
Observations	849			
<i>panel B: Markdowns and markups</i>	Markdown		Markup	
	Est.	S.E.	Est.	S.E.
Average markdown	0.892	2.868	5.256	3.425
Median markdown	0.337	0.457	3.813	2.356
<i>panel C: Treatment effect</i>	log(Markdown)		log(TFP)	
	Est.	S.E.	Est.	S.E.
Treatment*1(Year \geq 2002)	0.643	0.263	0.341	0.129
R-squared	0.73		0.89	
Observations	1130		1132	

Notes: This table reports the estimated production function, markdowns, and treatment effects for the gross-output production function in which tobacco leaf is substitutable. Standard errors are block-bootstrapped with 200 iterations.

Table A7: Alternative production functions

<i>panel A: Production function</i>	Translog		Cost shares	
	Est.	S.E.	Est.	S.E.
Output elasticity of labor	0.501	0.219	0.617	0.059
Output elasticity of capital	0.605	0.124	0.383	0.059
Returns to scale	1.107	0.128	1.000	0.000
R-squared	0.92		0.27	
Observations	849		659	
<i>panel B: Markdowns</i>	Translog		Cost shares	
	Est.	S.E.	Est.	S.E.
Avg. markdown	3.497	2.201	4.008	0.425
Med. markdown	1.980	0.092	2.103	0.055
<i>panel C: Consolidation treatment effects</i>	Translog		Cost shares	
	Est.	S.E.	Est.	S.E.
Markdown effect	0.310	0.111	0.374	0.112
TFP effect	-0.050	0.085	-0.147	0.124
Observations	1132		932	

Notes: This table reports the model estimates when using a translog production function (first column), and a model in which labor and capital are both variable and in which the production function is estimated using a cost shares approach (second column). All standard errors are block-bootstrapped with 200 iterations.

Table A8: Other robustness checks

<i>panel A: Different market definitions</i>	log(Markdown)		log(TFP)	
	Est.	S.E.	Est.	S.E.
Province-level treatment	0.086	0.057	-0.068	0.118
Prefecture-level treatment	0.214	0.061	-0.044	0.080
County-level treatment	0.315	0.103	-0.055	0.083
<i>panel B: Dropping small firms</i>	log(Markdown)		log(TFP)	
	Est.	S.E.	Est.	S.E.
1(Treatment) *1(year>2002)	0.315	0.103	-0.055	0.083
<i>panel C: Leaf content variation</i>	log(Leaf content per cigarette)			
	Est.	S.E.		
1(One firm in the county)	0.010	0.015		
1(Two firms in the county)	-0.005	0.020		

Notes: Panel (a) compares the markdown and productivity effects from the consolidation when re-defining leaf markets at the provincial and prefectural levels. Panel (b) drops SOEs that produced below the exit threshold of 100K cases from the sample. Panel (c) regresses the log leaf content per cigarettes at the county-level on indicators of the presence of one or two cigarette manufacturers in that county.

Table A9: Different markup calibrations

	Markdown level		Treatment effect	
	Est.	S.E.	Est.	S.E.
$\mu = 0.644$	4.813	0.415	0.314	0.107
$\mu = 1.000$	2.904	0.262	0.315	0.108
$\mu = 1.502$	1.750	0.182	0.257	0.102

Notes: The first column reports the average markdown for three calibrated markup values: $\mu = 0.644$, which is the median markup found in Appendix A. Second, $\mu = 1.000$, which is the markup value used in the main text. Third, $\mu = 1.502$, which is the upper bound of the 95% confidence interval for the median markup in Appendix A. The second column reports the corresponding markdown treatment effects of the consolidation.

Table A10: Quality

<i>panel A: Quality and consolidation</i>	1(Low quality)		1(Subsidy)		Subsidy/Revenue	
	Est.	S.E.	Est.	S.E.	Est.	S.E.
1(Treatment) *1(year>2002)	-0.014	0.043	-0.093	0.047	0.014	0.015
R-squared	0.74		0.47		0.64	
Observations	530		1575		1482	
<i>panel B: Markdowns and consolidation</i>	log(Markdown)		log(Markdown)		log(Markdown)	
	Est.	S.E.	Est.	S.E.	Est.	S.E.
1(Treatment) *1(year>2002)	0.241	0.110	0.244	0.110	0.243	0.110
Control for subsidy dummy:	No		Yes		Yes	
Control for subsidy/revenue:	No		No		Yes	
R-squared	0.76		0.76		0.76	
Observations	881		881		881	

Notes: Panel (a) estimates the difference-in-differences model with three quality indicators as the left-hand variable: an indicator for the firm producing the two lowest quality grades, an indicator for the firm receiving subsidies, and the subsidy-to-revenue ratio. Panel (b) adds the subsidy variables as controls to the treatment effects estimation for markdowns, on the sample on which subsidies are observed (1999-2002 and 2005-2006). Standard errors are block-bootstrapped with 200 iterations.

Table A11: Markdown covariates

	log(Markdown)	
	Est.	S.E.
1(SOE)	-0.194	0.087
log(Revenue)	0.141	0.029
log(Tax rate)	0.135	0.042
log(Unemp. rate)	-0.002	0.037
1(# Firms = 1)	0.286	0.101
1(# Firms = 2)	0.128	0.102


Notes: This table reports a regression of log average leaf markdowns on firm and county characteristics. Standard errors are block-bootstrapped with 200 iterations.

Table A12: Market structure and leaf prices

	log(Leaf price)					
	Province		Prefecture		County	
	Est.	S.E.	Est.	S.E.	Est.	S.E.
1 firm	-0.788	0.138	-0.917	0.108	-1.168	0.283
2 firms	-0.644	0.201	-0.330	0.119	-0.837	0.293
R-squared	0.08		0.11		0.04	
Observations	221		787		1034	

Notes: I regress the logarithm of the leaf price on dummies that indicate whether each market contains one or two cigarette manufacturers. Each column uses a different leaf market definition.


Figure A1: Annual firm size distributions, pre-consolidation.




Notes: This graph plots the distribution of the number of cigarette cases produced per firm in 1999, 2000, and 2001 for firms producing less than 400,000 cases. There is no evidence for ‘bunching’ just above the exit threshold of 100,000 cases or the merger threshold of 300,000 cases, so self-selection into the group just above these thresholds is unlikely.

Figure A2: Pre-trends

(a) Markdown



(b) Productivity



Notes: This graph plots the predicted average log markdown and productivity level for the treatment and control groups, which are obtained by estimating Equation (8).

References

- Akerberg, D., K. Caves, and G. Frazer (2015). Identification properties of recent production function estimators. *Econometrica* 83(6), 2411–2451.
- Azar, J. A., S. T. Berry, and I. Marinescu (2022). Estimating labor market power. Technical Report 30356, National Bureau of Economic Research.
- Bai, X. (2019). *Essays on China's Cigarette Industry*. Ph. D. thesis, Clemson University.
- Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica* 64(4), 841–890.
- Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, 242–262.
- 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 (2019). Wto accession and performance of chinese manufacturing firms: Corrigendum. *American Economic Review* 109(4), 1616–21.
- 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.
- Brown, J. D., J. S. Earle, and A. Telegdy (2006). The productivity effects of privatization: Longitudinal estimates from Hungary, Romania, Russia, and Ukraine. *Journal of Political Economy* 114(1), 61–99.
- Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics* 36(1), 13–70.
- 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.
- Collard-Wexler, A. and J. De Loecker (2015). Reallocation and technology: Evidence from the us steel industry. *American Economic Review* 105(1), 131–71.
- 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.
- Doraszelski, U. and J. Jaumandreu (2018). Measuring the bias of technological change. *Journal of Political Economy* 126(3), 1027–1084.
- Doraszelski, U. and J. Jaumandreu (2021). Reexamining the De Loecker and Warzynski (2012) method for estimating markups. *Unpublished*.
- 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.
- Foster, L., J. Haltiwanger, and C. Syverson (2008). Reallocation, firm turnover, and efficiency: selection on productivity or profitability? *American Economic Review* 98(1), 394–425.
- 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.
- Hijmans, R. J., L. Guarino, C. Bussink, P. Mathur, M. Cruz, I. Barrentes, and E. Rojas (2004). Diva-gis.
- 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.
- Levinsohn, J. and A. Petrin (2003). Estimating production functions using inputs to control for unobservables. *Review of Economic Studies* 70(2), 317–341.
- Li, Q., A. Hyland, G. T. Fong, Y. Jiang, and T. Elton-Marshall (2010). Use of less expensive cigarettes in six cities in China: findings from the International Tobacco Control (ITC) China Survey. *Tobacco Control* 19(2), 63–68.
- 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.
- Li, Y., K. Ren, M. Hu, X. He, K. Gu, B. Hu, J. Su, Y. Jin, W. Gao, D. Yang, et al. (2021). Cold stress in the harvest period: effects on tobacco leaf quality and curing characteristics. *BMC Plant Biology* 21(1), 1–15.
- Macchiavello, R. and A. Morjaria (2015). The value of relationships: evidence from a supply shock to Kenyan rose exports. *American Economic Review* 105(9), 2911–45.

- 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.
- 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.
- 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.
- O’Connor, R. J., L. M. Schneller, R. V. Caruso, W. E. Stephens, Q. Li, J. Yuan, and G. T. Fong (2015). Toxic metal and nicotine content of cigarettes sold in China, 2009 and 2012. *Tobacco Control* 24(4), 55–59.
- 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.
- Sumner, D. A. and J. M. Alston (1987). Substitutability for farm commodities: the demand for us tobacco in cigarette manufacturing. *American Journal of Agricultural Economics* 69(2), 258–265.
- United Nations (2019). Un comtrade database. <http://comtrade.un.org/>.
- 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.
- Xu, S. S., S. Gravely, G. Meng, T. Elton-Marshall, R. J. O’Connor, A. C. Quah, G. Feng, Y. Jiang, G. J. Hu, and G. T. Fong (2019). Impact of China National Tobacco Company’s ‘Premiumization strategy: longitudinal findings from the ITC China surveys (2006–2015). *Tobacco Control* 28(1), 68–76.