

Is There an Industrial Land Discount in China? A Public Finance Perspective*

Zhiguo He, Scott Nelson, Yang Su, Anthony Lee Zhang, Fudong Zhang

February 2022

Abstract

China's land market features a substantial *industrial discount*: industrial-zoned land is an order of magnitude cheaper than residential land. In contrast to explanations centered on subsidies to industry, we find that a primary determinant of this price gap is local public finance. Under the "land finance" system, land sales are an important source of revenues for Chinese local governments. We show that local governments, who serve as monopolistic land sellers in China, face a trade-off between supplying residential or industrial land that is determined by the different time profiles of revenues from industrial and residential land sales, local governments' financial constraints, and the extent of local governments' tax revenue sharing with other levels of government.

*He: University of Chicago Booth School of Business and NBER, Zhiguo.He@chicagobooth.edu; Nelson: University of Chicago Booth School of Business, Scott.Nelson@chicagobooth.edu; Yang: University of Chicago Booth School of Business, ysu4@chicagobooth.edu; A. Zhang: University of Chicago Booth School of Business, anthony.zhang@chicagobooth.edu; and F. Zhang: Tsinghua University PBC School of Finance, zhangfd@pbcfs.tsinghua.edu.cn. Zhiguo He acknowledges financial support from the John E. Jeuck Endowment at the University of Chicago Booth School of Business. This research was funded in part by the Tsinghua University – University of Chicago Joint Research Center for Economics and Finance. Yiren Ding, Xuechen Hong, Willy Yu-Shiou Lin, and especially Futing Chen provided outstanding research assistance. We are grateful to Ernest Liu and Chad Syverson for comments, and Wei Jiang, Xiao Cen, Zhuo Chen, and Li-an Zhou for sharing data.

1 Introduction

China's land market, which is perhaps one of the key drivers of the extraordinary growth of the Chinese economy in the past forty years, has systemic importance to the world. Because China does not employ a property tax, one of the unique features of the Chinese land market is the practice of "land finance" (Cao et al., 2008; Lin and Yi, 2011; Liu et al., 2014) by which local governments, who serve as monopolistic sellers and who control land supply in their local land markets, heavily rely on land sales for fiscal revenue (see, e.g., Liu and Xiong, 2020). This differs drastically from other developed economies where municipal governments finance a significant part of their public expenditure via property tax revenues (Ahern, 2021).

As in many other countries, there are rigid zoning restrictions in China to classify different land parcels for different uses. As highlighted in Chen et al. (2018), land zoned for residential use sells at roughly a ten-fold higher price than land zoned for industrial use. For example, in 2019 the average price of residential land in China was 3,619 RMB/m², while the average price of industrial land was 304 RMB/m².

We call this price difference between residential and industrial land the *industrial land discount* (or *industrial discount* interchangeably). Our paper aims to offer a comprehensive study of this industrial discount, which has profound implications for the real estate market, public finance, economic growth, and even political economy in China.

The typical view in the literature is that residential land sales are primarily a way for local governments to raise revenues, whereas industrial land is sold primarily to subsidize industry, stimulate economic growth, and support labor demand. Under the "land finance" system, a large share of city governments' operational revenues come directly from land sales, so governments must then sell some high-priced residential land purely for revenue purposes. Governments nevertheless may be willing to sell industrial land at substantially lower prices, in order to stimulate local growth. In support of this view, in discussing the industrial land discount Liu and Xiong (2020, pp. 193) state that "it is common practice for local governments throughout China to offer industrial land at subsidized prices to support local industries."¹

¹There is a broad narrative of "favoring industry and investment over the service sector and domestic

This paper proposes an explanation for the industrial land discount that stems from local public finance rather than from subsidies to industry. We propose that the choice between residential and industrial land sales, from the perspective of city governments as well as the central government, essentially involves an intertemporal revenue tradeoff. Chinese local governments are predominately funded through a combination of corporate tax revenues and land sale revenues; in 2019, these two numbers were roughly 8.7 trillion RMB and 7.3 trillion RMB respectively. Together they count for roughly 60% of local government revenue.² Industrial land generates future tax flows, since industrial firms pay value-added taxes, income taxes along with various fees; residential land does not.³ This simple fact implies that governments face a choice between selling residential land, which pays larger upfront revenues from higher sale prices, and selling industrial land, which pays smaller upfront revenues but comes with a stream of future cash flows from tax revenues over time.

This dynamic perspective implies that the large discount in the upfront sale price of industrial land relative to residential land does not necessarily imply that governments are systematically subsidizing industry through cheap land. Indeed, we show that the flow of tax revenues from industrial land can quantitatively explain the size of the upfront industrial land discount. Rather, our results imply that local governments' financing needs affect land supply to Chinese industry, and hence local public finance plays an underappreciated role in shaping the path of China's economic growth through the land allocation channel.

To reach these conclusions, the first step we take is to measure the cash flows generated over time by industrial and residential land sales. Specifically, we wish to estimate the upfront revenue per square meter of industrial and residential land sales, as well as the

consumption". In more recent years, China has shifted to target subsidies at specific "strategic" industrial sectors (Chen and Naughton, 2016; Liu, 2019). The analysis of the choice between residential and industrial land sales in this paper is more in line with the first broad-based industrial policy, rather than the second subsidies targeted at specific sectors. In Section 4.4, we briefly analyze whether our estimates can be affected by the targeted industrial policies by looking at the non-targeted industries only.

²For detailed calculation, see last paragraph in Section 2.1 on the institutional background of China's land market.

³As of February 2022, there is no nationwide property tax on residential properties in China. However, the formal introduction of a property tax was put on the policy agenda in October 2021 (as it has been several times before). Whether and how to implement a property tax is arguably the most debated public policy in recent China.

annual future tax revenues generated by industrial land sales. This requires two datasets. The first is data on the universe of land parcels sold by the Chinese government, from 2007 to 2019. We observe the price of each parcel and the name of the buyer, whether it is zoned for industrial or residential use, and characteristics of the parcel such as its location and size. The second is data on large Chinese industrial firms, including manufacturing, mining, and utility firms. By merging these two datasets, we are able to identify the industrial firm who acquired each land parcel, as well as the firm's performance before and after purchasing land.

We estimate the industrial land discount based on a potential-outcomes framework. We use observed residential (industrial) sale prices to estimate a hedonic model to predict what the prices of industrial (residential) land parcels would have been if they were, counter-factually, sold as residential (industrial). We then estimate the industrial land discount by taking the difference between the actual (predicted) residential price and the predicted (actual) industrial price. We estimate that the average industrial land discount is 1,176.4 RMB per square meter (hereafter abbreviated RMB/m²).

To estimate the marginal tax revenues from industrial land sales, we first use a differences-in-differences approach to estimate the marginal impact of land purchases on firms' sales. We then estimate marginal tax revenues by multiplying the increase in sales of the land-purchasing firms by an effective tax rate, guided by a simple model of a vertical production network to take into account the spillover effect along the supply chain. We find that the annual future marginal tax revenues from industrial land sale are approximately 113.6 RMB/m² in the first two years, and 214.2 RMB/m² thereafter.

These two sets of estimates for incremental cash flows allow us to back out an internal rate of rate (IRR) implied by the industrial-residential land tradeoff. We find this IRR, or the discount rate that equates the net present value of industrial versus residential land sales, to be 13.94%. Importantly, the estimated IRR is greater than local governments' cost of capital, which when proxied by their bond yields is between 3.5% and 7.5%, by a significant margin. Thus, industrial land sales in China are not subsidized relative to residential land sales, once we take future tax revenues into account. This is the main takeaway of our paper.

What can drive this positive wedge between the IRR and the government cost of

capital? In the conceptual framework developed in Section 3, we discuss the role of tax revenue sharing between local governments and other governments, the wedge between a local government's market power (captured by demand elasticities) in industrial and residential land markets, and other non-pecuniary benefits or strategic considerations such as industrial policies.

The positive wedge we identify between local governments' cost of capital and IRR from land sales implies that future tax revenues more than compensate the current industrial land discount. Our framework shows that this cannot be explained just by the demand for residential land being more inelastic than industrial land. Instead, we find that city governments' sharing of VAT tax revenues with other governments (e.g., the central government) are likely the main driver of the wedge between IRR and cost of capital. Under the current tax system in China, city governments receive almost the entirety of revenue from both residential and industrial land sales – a key feature of the “land finance” system. However, city governments only receive a fraction of future VAT tax revenue generated by firms.⁴ As a result, city governments may only internalize a fraction of future value-added tax revenues, which we find to be the most likely explanation for why we estimate a higher land-sale IRR than local governments' cost of capital.

Based on these findings, we propose that land allocation decisions in China are essentially determined by the interaction of three forces: the “land finance” system, through which land sales are a core source of local governments' operational revenues; the distinct time profiles of revenues from industrial and residential land sales along with the governments' financial constraints; and the way that industrial tax revenues are split between the central government and local governments. Section 5 provides further evidence that industrial discounts are associated with local governments' share of industrial tax revenues and local governments' discount rates. First, we show that industrial land discounts are positively correlated with local governments' shares of value-added taxes in the cross-section. We then exploit a 2016 change in tax sharing between local governments and the central government, and show that cities that experienced a larger increase in their share of value-added taxes subsequently exhibited greater increases in their industrial

⁴This is the part of value-added taxes that are *directly* accrue to city governments. City governments also receive these value-added taxes in an *indirect* way (perhaps from other cities) via “transfers” from higher-level governments.

land discounts. Second, if local governments' choice between industrial and residential land sales represents an intertemporal revenue tradeoff, industrial discounts should be lower when the governments' discount rates are higher, e.g., when the governments are less patient or face greater financial constraints. Consistent with this hypothesis, we show that industrial land discounts are negatively associated with local governments' cost of capital, as measured by local governments' municipal corporate bond yields, in the cross-section of cities. The negative correlation also holds when we instrument for municipal corporate bond yields using a strategy that builds on [Chen et al. \(2020\)](#).

Our paper relates to the following strands of literature. Many papers have argued that local governments in China tend to suppress industrial land prices while inflating residential prices. [Liu and Xiong \(2020\)](#) illustrate the diverging trends between industrial, commercial, and residential land prices, and argue that the industrial price gap is due to local governments' incentives to subsidize industrial land to support local industries. [Lei and Gong \(2014\)](#) argue that, in order to increase fiscal income and city output, it is optimal for local governments to distort the relative prices of industrial and residential lands, due to the agglomeration effects of industrial land sales, as well as future tax revenues from firms.

There are several papers that empirically investigate price distortions in the Chinese land market due to corruption. [Cai et al. \(2013a\)](#) and [Li \(2019\)](#) show that the local governments take advantage of differences between auction formats to influence effective land prices. [Chen and Kung \(2019\)](#) argue that firms with links to Chinese political elites are able to obtain large price discounts in the land market.

Researchers have also related local government incentives to land market subsidies. [Tian et al. \(2019\)](#) show that local government leaders adjust land policies in their jurisdictions at different stages of their term of office, in response to incentives for promotion. [Xie et al. \(2019\)](#) test the effects of VAT sharing and business tax sharing on local governments' land allocation decisions. [Fan et al. \(2015\)](#) model Chinese local governments' incentive to increase the supply of land for industrial use in order to generate more labor inflow and faster urbanization.

Relatively little research directly examines the rate of return on land sales. One exception is [Fu et al. \(2021\)](#), who calculate the average productivity of land using city-

level data and show that a growing share of land conversion (from agriculture to urban) quota is allocated to less productive cities. Several papers relate land supply to industry and firm outcomes in China. [Tian et al. \(2020\)](#) match land transaction data with industry-county-specific characteristics, and show that industries which can generate stronger spillover effects to local incumbents through agglomeration economies were favored in land allocation by governments. [Fei \(2020\)](#) quantitatively examine the impact of firm ownership on the cost of land for Chinese manufacturers.

This paper proceeds as follows. Section 2 describes institutional details of the Chinese land market, as well as our data. Section 3 introduces a conceptual framework illustrating the tradeoff local governments face between industrial and residential land sales. Section 4 shows how we estimate industrial discounts, marginal tax revenues generated by industrial land sales, and government IRRs. In Section 5, we show how industrial discounts are associated with local governments' shares of industrial tax revenues, as well as local government bond yields. We conclude in Section 6.

2 Institutional Details and Data

This section provides a brief summary of the key institutional background for Chinese land markets, followed by the description of data used in this paper.

2.1 Institutional Background

The Chinese land market. There was no formal land market in China before the 1980s. Housing reform and the development of market-based “commodity housing” in the 1990s opened up land leases for the residential market. The industrial land market also developed with the reform of state owned firms and the growth of private firms. The 1994 Tax-Sharing Reform made land lease sales an important source of local government revenue, spurring the take-off of the Chinese land market. However, regulations for the land market were not fully in place yet; without any requirement to release land lease contracts to the public, local governments had much leeway in granting land leases via hidden “negotiations.” Combined with the substantial ambiguities in the scope of

property rights granted by land leases, leases were often granted significantly below market values, leading to corruption and efficiency losses (Cai et al., 2013a).

During the late 1990s and early 2000s, the central government passed several laws and regulations to formalize the land market, with the intention of banning private negotiated deals from the government. Most land leases were required to be sold at market prices, through a variety of different market mechanisms. Auctions are the most popular sales mechanism, during which most of the information about the land, as well as details about the pricing process, are revealed to the public. Governments are also allowed to use “agreements” to set prices, if the requirements for auction-based land sales are not met. Regardless of the sale method used, local governments are now required to make deeds records publicly available.⁵ The permitted uses for each land parcel – analogous to zoning restrictions – are also strictly regulated. There are essentially four land usage categories: residential, industrial, commercial, and public utility (which covers education, health, transportation, and related uses). We provide a detailed explanations of these market mechanisms as well as land categories in Appendix A.1.

Industrial versus residential lands. Our study focuses on residential and industrial lands, which are the two major land categories and generally account for over 60% of total land supply. The quantity and price gaps between those two types of land have important implications for city development and therefore have attracted extensive prior research. There are guidelines routinely published by the Ministry of Housing and Urban-Rural Development on the split of different types of land usage in a city, which are generally loosely stated and not necessarily binding. For example, the 2012 Code states that residential land share range should fall in the range between 25% to 40% and the industrial land share in the range between 15% to 30%, with exceptions permitted for cities with special characteristics.⁶

⁵The Ministry of Natural Resources (then the Ministry of Land and Resources) informally maintained a website which posted the information of all the land transactions at deeds level, which is our data source. Although most of the regulations were in place in early 2000s, it was not until the mid 2000s – which is when our empirical data starts – that local governments started to fully comply with these regulations.

⁶See the 2012 Code ([link here](#)) for classification of urban and rural land use and planning standards of development land. The Ministry of Housing and Urban-Rural Development routinely updates those codes.

Land quota system and its determination. The Chinese government has a land quota system that plans out the maximum amount of newly developed urban land for each city over different horizons. The quota system is first of all “top-down”: it is based on land use master plans (usually covering 15 years) from the national level down to provincial and city level. At higher levels of government, these plans are jointly carved out by several national ministries and the provincial governments. The system is also “bottom-up”. At the lower level, city governments participate in drafting those plans as well as provide feedback to higher level governments. Accompanied by the long-term land use master plan, there are also five-year (more granular) land use plans drafted by city governments and then reviewed and approved by provincial and national governments. In addition, the master plan generally gets modified or amended every five years, reflecting input from all levels of government.⁷

Under the total quotas set by these medium- and long-term plans, local governments decide how to implement these land supply plans in the short-run. Each year, based on its economic development needs, a local government first decides how much quota they would need out of their medium-long-term cap, files a proposal to the Ministry of Natural Resources, and subsequently supplies land according to the quota after approval. It is also important to note that the quota issued by the Ministry of Natural Resources concerns only the total area of land supplied across all uses, so local governments have freedom in allocating the quota to different types of land. Moreover, local governments can occasionally apply for special quota deviations from the regulation of land use plan. Overall, local governments tend to have substantial control over land supply composition in the short run and are an integral part of the medium-long term land-use planning.

The actual supply of land comes out of each city’s “Municipal Land Reserve Center,” which is an executing institution of the municipal government that monopolizes the primary land market. Land reserve centers are responsible for procuring “raw land,” preparing them to “sell-able land,” and reserving them as land reserve assets for the city (to use or to sell later). See Appendix A.1 for a detailed discussion of how the land reserve system works in China.

⁷In addition, sometimes off-schedule modifications are possible, for example when higher-level governments see issues in the land market, or when a local government files an application for a change of land usage plan.

Land allocation and local government financing. Land allocation has important implications for Chinese local government public finance. As highlighted by a team of named Chinese scholars and policy makers (Cai et al., 2013b), “Land finance is a key challenge: most Chinese cities fund their urban infrastructure largely from land sales; 40% of the government debt needs land finance in 2010; land sale revenue accounts for about one third of total local government revenue during 2010-2012.” Another important role played by land in the fiscal equation is that future land reserves can serve as collaterals for local government debt issuing; according to the report (link here) from the Chinese National Audit Office, 37.23% of the debts of local governments explicitly pledged future land sales revenue as collaterals by the end of 2012.

Taxes revenue from firms also play an important role in funding local governments in China. According to the recent figure published by the Ministry of Finance, local governments’ total fiscal revenue in 2019 largely comes from three sources: 10 trillion RMB from local government general revenue, 7.5 trillion RMB from the central government transfer payments, and 8 trillion RMB from local government-managed funds (7.3 trillion RMB from land sales).⁸ Since about half of the local government general revenues and central government transfers are from value-added taxes and corporate income taxes, local government revenue from the two tax items is about 8.7 trillion RMB ($\approx \frac{10+7.5}{2}$). The overall picture is that, as the top two sources of revenue, taxes from firms and direct land sales together count for over 60% of local government’s budget in China.

2.2 Data

Land sale data. We use land sale data from the Ministry of Natural Resources. The data cover the universe of land sales by the local governments in China from 2007 to 2019. We focus on residential and industrial land parcels allocated by agreement, tender, auction and listing.⁹ We retrieve data on the geographical coordinates of land parcels using

⁸See the Ministry of Finance’s 2019 Report (link here) on the execution of the central and local budgets.

⁹This excludes a category of allocation called “administrative allotments”, which involves no payment from the land receiver, and is generally used for infrastructures, government offices, military facilities, etc. Throughout the paper, we will use auction to refer to all the three allocation methods of tender, auction and listing.




Figure 1: Average Land Prices Over Time by Land Use: Industrial vs. Residential

Note: This figure reports the average price of residential and industrial land weighted by land size that are sold through auctions for each year during 2007-2019.

the Gaode maps API, a leading Location Based Services (LBS) provider in China.¹⁰ We generally define markets as urban units, which are contiguous urban clusters as identified by satellite images (see Appendix A.3). Figure 1 shows prices of industrial and residential land. Residential land prices exceed industrial land prices by a significant amount, and the price gap increases over time.

Firm data. We use firm data from the NIE database, which is survey data collected by the National Bureau of Statistics on all industrial (manufacturing, mining, and utility) firms in China during 1998-2013 (except 2010). Despite some concerns about the data quality (Nie et al., 2012), the data has been widely used in economic research on China. Some studies use the data until 2005 (Hsieh and Klenow, 2009) or until 2007 (Liu and Lu, 2015; Bai et al., 2019), and others use the data until 2013 (Heinrich et al., 2020; Cen et al., 2021; Tang et al., 2021). For the year 2010, all operating information except sales and employment is missing, and we drop that year due to concerns about data quality. The firm data is subject to some censoring and random dropout concerns, which we analyze in appendices A.2 and C.6.

¹⁰See <https://lbs.amap.com/>.

Since one major exercise of this paper is to estimate the marginal impact of land acquisition on firm sales, we merge firm data to industrial land purchase data using firm names, taking into account firms buying land through their subsidiaries. To simplify our estimation, we exclude firms that purchased land in multiple years during our sample, and we focus on only firms that purchased land in a single year in the sample period. These firms form our treatment group in a matched-pairs difference-in-differences strategy, where the control firms are those who never purchased any new land during our sample period.

In total, we are able to merge 22,636 transactions out of a total of 124,341 industrial land purchases with firm buyers via agreement, tender, auction and listing during 2007-2010. In the NIE sample, around 3% of firm-year observations during 2003-2013 were matched to land purchases during 2007-2010. Appendix Table A.1 compares merged land parcels and firms to the universe of parcels and firms. Merged parcels are slightly more expensive, and indistinguishable in terms of size and distance to the urban unit centers from the population land parcels. Land purchasing firms are also slightly larger than the population of firms in terms of most metrics.

City data. We collect city-level data on GDP and population from Urban Statistic Yearbook published by the National Bureau of Statistics. The data covers all the municipal cities in China during 2007-2018.

3 Conceptual Framework

We start with defining the industrial land discount in this section, followed by a conceptual framework that highlights the intertemporal revenue tradeoff faced by local governments under the Chinese fiscal system.

3.1 Industrial Land Discounts

The core object we will study is the discount of industrial land relative to residential land. Figure 1 shows average land sale prices over time, separately for residential- and industrial-zoned land. The figure shows a striking pattern: industrial land prices per

Table 1: Data Summary

	Obs	Mean	Std Dev	P10	P50	P90
A. Land characteristics						
<i>Residential</i>						
Land Price, RMB/m ²	292,371	2,113.22	2,595.73	264.40	1,200.00	5,101.21
Area, 1000 m ²	292,371	29.80	46.90	0.20	15.30	69.50
Distance to urban unit centers, km	292,371	9.26	11.60	0.91	5.15	22.06
<i>Industrial</i>						
Land Price, RMB/m ²	371,717	252.30	260.43	96.00	199.11	445.48
Area, 1000 m ²	371,717	36.60	73.30	3.30	17.40	80.40
Distance to urban unit centers, km	371,717	10.28	10.93	1.46	7.57	21.77
B. Firm characteristics						
<i>Treated Firms</i>						
profit margin	21,976	0.05	0.08	0.00	0.04	0.13
sale, 1000 RMB	21,752	179,459.00	371,020.60	11,457.00	61,288.00	411,539.00
area, 1000 m ²	4,349	34.11	45.43	4.55	18.83	80.00
<i>Control Firms</i>						
profit margin	22,345	0.05	0.14	0.00	0.04	0.13
sale, 1000 RMB	22,110	174,006.70	353,169.10	10,814.50	60,484.00	402,751.00
C. City characteristics						
IndDisc, RMB/m ²	3,092	1,520.86	1,416.96	223.86	1,082.47	3,472.72
City VAT Share, %	2,839	23.94	10.15	15.00	20.00	40.00
Change of Ctiy VAT Share in 2016, %	216	20.29	6.41	16.25	20.00	27.50
City MCB Coupon rate, %	257	6.96	0.80	5.88	6.98	8.00
Deficit/GDP	257	0.09	0.08	0.01	0.07	0.17
GDP growth rate, %	257	13.39	3.19	10.10	13.20	16.50
GDP per capita, 10,000 RMB	257	2.46	1.71	0.97	1.93	4.98
LateTerm	257	0.14	0.35	0.00	0.00	1.00

Note: This table reports summary statistics at the land, firm-year and city-year level for the sample used in this paper. Panel A is based on the residential and industrial land auction transactions during 2007-2019, which we use to estimate the industrial discounts. Panel B is on the matched sample of firms which we use to estimate the sale effect of land purchase. In Panel C, the first two are time-varying city characteristics during 2007-2019, the fourth-seventh variables are city-level characteristics in 2008 and the last one is a binary indicator for whether the provincial governor had been in offer for more than three years at the end of 2008.

square meter are an order of magnitude lower than residential land prices throughout our sample period. Given the large share of local government revenue that derives from land sales, this raises the question of why a revenue-maximizing local government would not reallocate land sales from industrial to residential uses until the prices in both markets equalize.

Consider a single parcel of land to be sold in period t , and let p_t^{ind} and p_t^{res} respectively denote the parcel's expected price per square meter, if it is sold as industrial or residential land. Section 4.1 gives the details of our estimation method based on a potential-outcomes framework, which addresses the issue that the observed pattern in Figure 1 could be driven by the difference of the underlying characteristics of industrial and residential lands.

Because we are interested in governments' net revenue from selling lands, we need to adjust both p_t^{ind} and p_t^{res} for their respective costs. In general, there are three kinds of costs for land sales, some of which differ between industrial and residential lands; we explain the details in Appendix A.1. First, there is a fixed component, which mainly includes the "standard" compensation to incumbent land users for repossessing and reselling their land, as well as the cost of land development. These fixed costs apply equally to both types of land, and hence cancel out when calculating the industrial discount. The second type of cost is the "non-standard" compensation to local land occupants, which depends on the expected land sale price. This usually involves bargaining with local incumbent land occupants over a "resettlement agreement for demolition", in which incumbents are compensated essentially based on expected future revenues from land sales. Incumbent occupants generally know whether the redeveloped land is intended for residential or industrial uses. Rent-sharing with incumbents is a large part of the costs of residential land sales, but matters relatively little for industrial land sales, since industrial sales generate comparatively negligible upfront revenues, and thus generally involve less negotiation with local occupants. Third, as laid out in the Code for Planning Standards (see footnote 6 Item 4.3.2), selling residential land involves certain auxiliary costs, including extra land and fiscal support for education and other services associated with new residences. The latter two variable costs only apply to residential land. In Appendix C.1 we estimate that the sum of these two variable costs is about 1/3 of residential sale revenues.

Denote the additional cost associated with residential land by $\lambda = 1/3$. Throughout the paper we define the industrial land discount as:

$$\text{IndDisc} \equiv (1 - \lambda)p_t^{\text{res}} - p_t^{\text{ind}}. \quad (1)$$

That is, (1) is the difference in upfront profit per square meter from selling residential land compared to industrial land.

The extra cost by selling residential land could potentially explain part of the (raw) industrial land discount observed in Figure 1. As shown shortly in Section 4, however, under the estimate $\lambda = 1/3$, auxiliary costs alone are too small to explain the entirety of the difference between residential and industrial land prices.

As laid down in the next section, we hypothesize that an intertemporal revenue tradeoff faced by local governments under the Chinese tax system can largely explain the industrial land discount in Chinese land market. Although to our knowledge we are the first to conduct a comprehensive and focused study on this price gap, alternative explanations from prior research could include subsidies to economically important industries (Liu, 2019), subsidies to support labor inflow and urbanization (Lei and Gong, 2014 and Fan et al., 2015), incentives for promotion (Tian et al., 2019), corruption (Wu et al., 2012), or simply lower-quality land being allocated for industrial purposes. While we may refute some explanations, our perspective is largely complementary to these alternative views.

3.2 Tax Revenues and the IRR on Industrial Land Sales

As explained in the introduction, Chinese local governments are predominately funded through a combination of tax revenues and land sale revenues. Given the current land finance system in China, one factor in the choice between allocating a new land parcel for residential versus industrial purposes is the distinct time profiles of revenues from industrial and residential land sales: there are higher revenues upfront from higher-priced residential land sales, and higher revenues in the future from greater industrial tax revenues. The government's decision between industrial and residential land sales thus depends on how it values current versus future cash flows.

Following the terminology in practice in corporate finance (Berk and DeMarzo, 2017), we define the internal rate of return (IRR) on industrial land sales (relative to residential land sales) as follows. Suppose that if a land parcel is sold as industrial land in period t , the government expects a tax flow of tax_s per square meter in future years $s > t$. The government IRR is the discount rate that equates the net present value of industrial versus residential land sales: $(1 - \lambda)p_t^{\text{res}} - p_t^{\text{ind}}$:

$$\text{IRR}^{\text{ind}} \equiv \left\{ \rho : \sum_{s>t} \frac{\text{tax}_s}{(1 + \rho)^{s-t}} = (1 - \lambda)p_t^{\text{res}} - p_t^{\text{ind}} \right\}. \quad (2)$$

Under the premise that land allocation decisions made by Chinese local governments reflect the intertemporal tradeoff outlined above, what should be the relation between IRR^{ind} implied by the industrial land discount and the government's cost of capital, denoted by r^{gov} ? In a frictionless benchmark where the government internalizes all tax revenues fully, the government's indifference condition between selling a marginal land parcel as industrial versus residential indeed implies that $\text{IRR}^{\text{ind}} = r^{\text{gov}}$. There are, however, a number of economic and/or policy factors that could drive a wedge between IRR^{ind} and r^{gov} .

First, in China, city governments keep almost the entirety of their land sales revenue, but only a fraction of taxes paid by local firms. Suppose that the city government only internalizes a share $k \in (0, 1)$ of tax revenues from industrial firms. If we think of city governments as fully determining land allocation decisions and they only care about the revenues they get, k can be thought of as the share of VAT taxes that eventually accrue to city governments. More realistically, city governments negotiate with higher-level governments over land allocation decisions; k can then be thought of as the eventual weight that the bargaining outcome places on tax revenue, in determining land allocation decisions.

In Appendix B, we construct a simple model of partially internalized perpetuity cash flows, and show that we will have:

$$\text{IRR}^{\text{ind}} = \frac{r^{\text{gov}}}{k}. \quad (3)$$

Intuitively, the smaller the share of future tax revenues that a city government receives, the greater the required tax revenue generated from a land sale in order for the government's indifference condition to hold, and hence the higher IRR^{ind} , as defined in Eq. (2), will be.

Second, it is well known that Chinese local governments are monopolistic sellers in their local land markets. Thus, their sales can move land prices, and marginal revenue may differ from price. Taking the simple model in Appendix B, but further incorporating different demand elasticities in industrial and residential land markets, we show that

$$IRR^{ind} = \frac{r^{gov}}{k} \times \left[1 - \frac{\overbrace{\sigma_{res}^{-1} - \sigma_{ind}^{-1}}^{\text{wedge in inverse semi-elasticities}}}{IndDisc} \right] \quad (4)$$

where σ_{res} (σ_{ind}) is the negative demand semi-elasticity of residential (industrial) land, and $IndDisc = (1 - \lambda)p_t^{res} - p_t^{ind}$ is the industrial discount. In China, the demand elasticity for industrial land is likely to be greater than that for residential land, as firms typically shop around across different cities for the most favorable land price, while most households do not move across cities.¹¹ We therefore would expect the term in brackets in Eq. (4) to be less than 1. Accordingly, we would expect a smaller IRR^{ind} than the term r^{gov}/k . Intuitively, the monopolistic local government takes into account the price impact it has in the residential housing market and so will tend to maintain a higher industrial land discount, generating a lower implied IRR^{ind} .

Finally, IRR^{ind} in Eq. (2) incorporates the industrial firms' marginal tax revenues only; it ignores potential non-pecuniary benefits or costs the government derives from choosing industrial rather than residential zoning. One advantage of our approach, which provides a gauge of the magnitude of IRR^{ind} , is that it gives a clear guidance on whether one particular economic force can explain the empirical pattern. For instance, comparing our estimate of IRR^{ind} to values of r^{gov} will shed light on the importance of governments' market power relative to any such non-pecuniary benefits of industrial land sales.

¹¹One reason for households' immobility across cities is China's "hukou" residence restrictions (Li et al., 2017).

4 Estimation

In the framework laid out in Section 3, measuring industrial land discounts and estimating government IRRs requires three key quantities: p_t^{ind} and p_t^{res} , the representative prices per square meter of industrial and residential land; λ , the auxiliary costs from selling residential land; and Tax_s , the stream of expected future industrial tax revenues generated by industrial land sales. We provide these estimates in this section. After estimating industrial discounts and auxiliary residential sale costs in Section 4.1, we use a differences-in-differences approach in Section 4.2 to estimate the marginal effect of land purchases on firms' sales. Section 4.3 further estimates incremental tax revenue from land sales, and we then combine these estimates together to calculate IRRs on industrial land sales in Section 4.4.

4.1 Industrial Land Discount Estimation

For each parcel of land indexed by i , we first estimate the price of the land if it were sold for the alternative use (industrial or residential). Let p_{it}^{res} (p_{it}^{ind}) denote the price per square meter of the parcel assuming it is sold as residential (industrial) land. Let $\mathbf{1}_{it}^{\text{res}}$ be a dummy representing whether parcel i is actually sold as a residential parcel. The sale price for parcel i that we observe is:

$$p_{it} = \begin{cases} p_{it}^{\text{res}} & \mathbf{1}_{it}^{\text{res}} = 1, \\ p_{it}^{\text{ind}} & \mathbf{1}_{it}^{\text{res}} = 0. \end{cases} \quad (5)$$

Our goal is to estimate both outcomes p_{it}^{res} and p_{it}^{ind} , only one of which is observed.

The main challenge is that land parcels are not randomly zoned as residential or industrial use. For example, parcels closer to the city center are more likely to be used as residential, and hence it is likely that $\mathbb{E}[p_{it}^{\text{res}} | \mathbf{1}_{it}^{\text{res}} = 1] \neq \mathbb{E}[p_{it}^{\text{res}} | \mathbf{1}_{it}^{\text{res}} = 0]$. Therefore, one cannot directly take the average observed prices of residential land parcels as the predicted price of the industrial land parcels, if they were instead zoned for residential use. We must control for the differences in land characteristics between the two types of land parcels.

We proceed by using the sample of observed residential (industrial) sale prices to estimate a hedonic model to predict what the prices of industrial (residential) land parcels would have been if they were, counter-factually, sold as residential (industrial). Formally, let \mathcal{J}_{res} and \mathcal{J}_{ind} represent the sets of residential and industrial parcels, respectively:

$$\mathcal{J}_{\text{res}} \equiv \{i : \mathbf{1}_{it}^{\text{res}} = 1\}, \mathcal{J}_{\text{ind}} \equiv \{i : \mathbf{1}_{it}^{\text{res}} = 0\}.$$

For all residential parcels $i \in \mathcal{J}_{\text{res}}$, we will estimate the following regression specification:

$$p_{it} = X_{it} \cdot \beta^{\text{res}} + \gamma_{\text{ut}}^{\text{res}} + \epsilon_{it}, \quad \forall i \in \mathcal{J}_{\text{res}}. \quad (6)$$


Eq. (6) is a hedonic model that predicts p_{it} , the price per square meter of each residential land parcel. To control for geographical variation in prices within cities, we construct “urban units,” which are geographical units smaller than cities, by grouping contiguous pieces of urban land into blocks. We describe details of this procedure in Appendix A.3.¹² Parcel characteristics X_{it} consist of the following control variables: second-order polynomials in the log of the area of the land parcel, the distance to the center of the closest urban unit, and the year-quarter in which the land is sold. We also include urban-unit-by-year fixed effects γ_{ut} .

We estimate Eq. (6) by restricting the sample to the set of land parcels sold by auction. To account for the possibility that the coefficients may vary over time and across cities, we estimate (6) separately for each prefecture city, and separately for three time periods: 2007-2011, 2012-2015, and 2016-2019. Since specification (6) requires enough data to precisely estimate, we restrict to cities and periods in which we observe at least 100 industrial land sales as well as 100 residential land sales in the city in that period. This leaves us with 272 out of 341 cities, which collectively constitute 90.67% of all industrial and residential land sales through auction during 2007-2019.¹³

Using our estimates from specification (6), we can then predict residential prices for

¹²Appendix Figure A.1 shows some examples of the urban units in large and small cities.

¹³When constructing the industrial discount estimates without controlling for land characteristics X_{it} , we restrict to cities and years in which they are at least 10 industrial and 10 residential land sales. This leaves us with 308 cities, which collectively constitute 94.28% of all industrial and residential land sales through auction during 2007-2019.



(a) Prov. Industrial Discount vs GDP Per Capita

(b) Industrial Discounts Over Time

Figure 2: Industrial Discount Estimates Summary.

Note: Panel (a) plots the province-level industrial land discount against the GDP per capita, both taken as simple average across cities and years (during 2007-2018) for each province. Panel (b) shows the average quality-adjusted industrial discount estimates across cities for each year from 2007 to 2019.

industrial parcels by plugging characteristics of these parcels into Eq. (6):

$$\hat{p}_{it}^{\text{res}} = X_{it}\hat{\beta}^{\text{res}} + \hat{\gamma}_{ut}^{\text{res}}, \forall i \in J_{\text{ind}}. \quad (7)$$

That is, $\hat{p}_{it}^{\text{res}}$ is the predicted price of parcel i if it were sold as residential land. Analogously, we fit a hedonic model to industrial land parcels, with the same control variables as in (6):

$$p_{it} = X_{it}\beta^{\text{ind}} + \gamma_{ut}^{\text{ind}} + \epsilon_{it}, \forall i \in J_{\text{ind}}. \quad (8)$$

We then predict the counter-factual industrial prices for residential parcels as:

$$\hat{p}_{it}^{\text{ind}} = X_{it}\hat{\beta}^{\text{ind}} + \hat{\gamma}_{ut}^{\text{ind}}, \forall i \in J_{\text{res}}. \quad (9)$$

Using our estimates of $\{p_{it}^{\text{res}}, p_{it}^{\text{ind}}, \hat{p}_{it}^{\text{res}}, \hat{p}_{it}^{\text{ind}}, \lambda\}$, we can estimate industrial land dis-

counts for each parcel using equation (1) as follows:

$$\text{IndDisc}_{it} = \begin{cases} (1 - \lambda)p_{it}^{\text{res}} - \hat{p}_{it}^{\text{ind}}, & i \in \mathcal{J}_{\text{res}}; \\ (1 - \lambda)\hat{p}_{it}^{\text{res}} - p_{it}^{\text{ind}}, & i \in \mathcal{J}_{\text{ind}}. \end{cases}$$

In words, IndDisc_{it} is the actual (predicted) residential sale price minus the predicted (actual) industrial price for residential (industrial) parcels, where the residential prices are adjusted by $1 - \lambda = \frac{2}{3}$ (recall Section 3.1).

The estimation delivers IndDisc_{it} at the land parcel level. We then aggregate to form city-year level estimates, IndDisc_{ct} , by taking averages of IndDisc_{it} weighted by the size of each land parcel. Figure 2 shows the distribution of IndDisc_{ct} in the cross section and over time. Panel (a) plots the industrial land discounts against GDP per capita across provinces, showing a positive correlation between the level of industrial discount and economic development. Panel (b) plots the time series of the estimated industrial land discount. It was about 400-500 RMB/m² during 2007-2009 and increased to 750 RMB/m² around 2010. It remained stable during 2010-2015, but increased significantly during 2015-2018. In 2019, the average industrial discount reached about 2,000 RMB/m², five times the level in 2007.¹⁴ Section 5.1.2 provides one explanation for the significant increase of industrial land discounts from 2015 to 2018.

To show how controlling for characteristics affects our estimates of industrial discounts, we compare estimates with and without controls for characteristics (X_{it}) in Appendix Figure A.2. Overall, the estimates of industrial discounts decrease once we adjust for parcel characteristics; this is because industrial land parcels are in general located further from city centers (Table 1). However, a large fraction of the price difference remains after controlling for land characteristics.

¹⁴From 2007 to 2015, the simple average of residential (industrial) land price, where we use predicted value if not observed, across cities increased by a factor of 2.23 (1.41) in our data. Liu and Xiong (2020) controls for changing land characteristics and shows that residential land price increased by a factor of about 3.12 and the industrial land price barely changed during the same period.

4.2 Effects of Land Purchases on Firm Sales

4.2.1 Estimation Methodology

In this section, we discuss how we estimate the effects of land purchases on firms' sales using a differences-in-differences approach.

To bring the conceptual framework to data, we impose a few econometric assumptions. Suppose that in period τ_j , firm j purchases a land parcel of size Δ_j . Define $\Delta_{j,t} \equiv \Delta_j \cdot \mathbf{1}_{t \geq \tau_j}$; then firm j 's sales in period t take the following form:

$$S_{j,t} = \alpha_j + \eta_t + \theta_{t-\tau_j+1} \cdot \Delta_{j,t} + \varepsilon_{j,t}. \quad (10)$$

In words, Eq. (10) states that firms' sales are determined by time-varying factors η_t , time-invariant firm-specific factors α_j , and land purchases $\Delta_{j,t}$, whose effect depends on a parameter $\theta_{t-\tau_j+1}$. The time-varying factors η_t may represent, in reduced form, factors such as growth, demand, and input prices, while α_j represents persistent firm-specific productivity differences.

In this framework, treated firms are those that ever acquired some new industrial land during their presence in the sample, i.e., firms with $\Delta_j > 0$ for some $\tau_j < \infty$. For cleaner identification, we focus on the sample of firms who have purchased land only in one year in our sample period.¹⁵ In contrast, control firms are those that never acquired any new industrial land during their presence in the sample (so that $\tau_j = \infty$).

A natural concern with estimating the parameters $\theta_{t-\tau_j+1}$ in Eq. (10) is that land purchase decisions may be endogenous with firm-time-specific shocks $\varepsilon_{j,t}$. To address this concern, we decompose these shocks as

$$\varepsilon_{j,t} = f(p(x_{j,t})) + e_{j,t} \quad (11)$$

In this decomposition, f can be any function, and $p(x_{j,t})$ is a firm's probability of purchas-

¹⁵If a firm purchased multiple land parcels in one year, then we aggregate these purchases together as one firm-year observation.

ing land given observables $x_{j,t}$. We make the identifying assumptions:

$$\mathbb{E}[e_{j,t}\Delta_{j,t} \mid \alpha_j, \eta_t, \tau_j \in \{\tau, \infty\}] = 0 \quad \text{and} \quad \mathbb{E}[e_{j,t}\mathbf{1}_{\Delta_{j,t}>0} \mid \alpha_j, \eta_t, \tau_j \in \{\tau, \infty\}] = 0, \quad \forall \tau \quad (12)$$

In words, the two requirements for $e_{j,t}$ are that these shocks to firm sales be uncorrelated with i) the amount of land purchased, and ii) with the decision of whether to buy land, among firms that either purchase land in a particular year ($\tau_j = \tau$) or do not purchase land at all ($\tau_j = \infty$). The conditioning on α_j and η_t reflects that these assumptions only need to hold after we control for firm and time fixed effects. Because $e_{j,t}$ is the component of firm-time-specific shocks that are *unrelated* to the probability of land purchase predicted by $x_{j,t}$ (see equation (11)), we view this as a plausible identifying assumption, and moreover an assumption that we can partially test by examining pre-trends in sales among treatment and control firms.

Motivated by this framework, we match treated firms with control firms using propensity scores $\hat{p}(x_{j,t})$ for land purchase using firm characteristics in year $t = \tau_j - 1$. Recall that the control firms are those that did not acquire any new industrial land during their presence in the sample. After stratifying by event year, province, and two-digit National Industries Classification code, we estimate $\hat{p}(x_{j,t})$ based on three following observables at the firm level:

$$x_{j,t} = \left\{ \log S_{j,t-1}, \log S_{j,t-2}, \frac{\text{Profit}_{j,t-1}}{S_{j,t-1}} \right\}.$$

Here, $S_{j,t}$ is the firm j 's sales in period t and $\text{Profit}_{j,t}/S_{j,t}$ is firm j 's profit margin in period t . In our data, we find these three variables are predictive of land purchase decisions; other observables do not provide additional explanatory power for whether the firm purchases land in $t = \tau_j$.

After matching, one test of our assumption on the residuals $e_{j,t}$ will be whether treated firms and control firms exhibit parallel trends in sales prior to τ_j . We conduct this test as part of our differences-in-differences strategy below and confirm (fail to reject) parallel trends for all purchase cohorts τ .

We estimate the effects of land purchase, $\theta_{t-\tau_j+1}$, using difference-in-differences on the matched sample. To do so, we define the average land size in a given land-purchase

year τ as,

$$\bar{\Delta}_\tau \equiv \mathbb{E} [\Delta_j \mid \Delta_j > 0, \tau_j = \tau]. \quad (13)$$

In words, we estimate the average land size at a particular year τ by averaging over all land transactions in that year. Using Eq. (10), firm sales can therefore be equivalently written as

$$S_{j,t} = \alpha_j + \eta_t + \theta_{t-\tau_j+1} \cdot \mathbf{1}_{\Delta_j > 0} \cdot \bar{\Delta}_{\tau_j} + \varepsilon'_{j,t}, \quad (14)$$

where we define

$$\varepsilon'_{i,t} \equiv \begin{cases} e_{j,t}, & \Delta_{j,t} = 0; \\ e_{j,t} + \theta_{t-\tau_j+1} \cdot (\Delta_{j,t} - \bar{\Delta}_{\tau_j}), & \Delta_{j,t} > 0. \end{cases} \quad (15)$$

Note that,

$$\mathbb{E}[\varepsilon'_{jt} \bar{\Delta}_{\tau_j} \mid \alpha_j, \eta_t] = 0, \quad (16)$$

where we use conditioning on α_j and η_t to reflect controlling for firm and time fixed effects. This follows from (12) thanks to the definition of $\bar{\Delta}_{\tau_j}$ in equation 13. In light of Eq. (16), we can consistently estimate $\theta_{t-\tau_i+1}$ with difference-in-differences estimation using Model (14).

4.2.2 Empirical results

Table 2 reports the estimates of specification (14). We allow the time fixed effects η_t to vary at the province-year level in all the regressions to absorb differences in time trends across provinces. For each purchase year $\tau \in \{2007, 2008, 2009, 2010\}$, we use data from years $\tau - 4$ through the year 2013 (there are very few firms with data before $\tau - 4$).

Table 2 shows three important patterns. First, estimated treatment effects are positive and are both economically and statistically significant. Each square meter of land generates, for example, 428.2 RMB in additional sales in the first year after land purchase in 2007. Second, overall, the estimated treatment effects grow over time.

Third, and importantly for validating our matched difference-in-differences identification assumptions, treated and control firms are not significantly distinguishable in years prior to the event. Note, our matching procedure guarantees that the parallel trend holds between the treated and control from $t = \tau - 2$ to $t = \tau - 1$. The fact that the

Table 2: Dynamic Treatment Effect of Land Purchase on Sales

Event Year	2007	2008	2009	2010
Dep Var: Sale	(1)	(2)	(3)	(4)
$\bar{\Delta} \cdot \text{Treat} \cdot (t - \tau = -4)$	192.9 (0.562)	-77.30 (-0.110)	-17.82 (-0.0717)	141.4 (0.444)
$\bar{\Delta} \cdot \text{Treat} \cdot (t - \tau = -3)$	2.781 (0.0143)	185.0 (0.337)	-105.7 (-0.534)	339.4 (1.568)
$\bar{\Delta} \cdot \text{Treat} \cdot (t - \tau = -2)$	10.21 (0.0936)	-107.3 (-0.363)	69.11 (0.554)	191.5 (1.558)
$\bar{\Delta} \cdot \text{Treat} \cdot (t - \tau = 0)$	428.2*** (2.869)	938.1** (2.257)	287.3** (2.133)	
$\bar{\Delta} \cdot \text{Treat} \cdot (t - \tau = 1)$	783.3*** (3.235)	1,097** (2.486)		772.3*** (2.695)
$\bar{\Delta} \cdot \text{Treat} \cdot (t - \tau = 2)$	687.0** (2.207)		655.6** (2.129)	1,048*** (2.985)
$\bar{\Delta} \cdot \text{Treat} \cdot (t - \tau = 3)$		1,691** (2.070)	602.9* (1.678)	1,497*** (3.299)
$\bar{\Delta} \cdot \text{Treat} \cdot (t - \tau = 4)$	814.2* (1.847)	2,222* (1.725)	965.8** (2.333)	
$\bar{\Delta} \cdot \text{Treat} \cdot (t - \tau = 5)$	1,293* (1.757)	1,600 (1.059)		
$\bar{\Delta} \cdot \text{Treat} \cdot (t - \tau = 6)$	2,081** (1.968)			
Firm FE	Yes	Yes	Yes	Yes
Province-Year FE	Yes	Yes	Yes	Yes
Observations	9,189	4,046	13,132	16,510
R ²	0.475	0.522	0.521	0.497

Note: This table reports estimation results of Model (14) with the matched sample. We drop the matched pairs whenever the treated or the control firm exits the sample. For each treatment year $\tau \in \{2007, 2008, \dots, 2010\}$, the sample ranges from $\tau - 4$ to 2013 (but the data for 2010 is missing). The variable sale is in 1,000 RMB and $\bar{\Delta}$ is in 1,000 m². The year of $t = \tau - 1$ is used as the base year. Standard errors are clustered by firms. Robust t-statistics in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

parallel trend holds from $t = \tau - 4$ to $t = \tau - 1$ lends certain support to our identification assumption.

Motivated by these patterns, Table 3 summarizes the estimated treatment effects more

Table 3: Baseline Estimation of Marginal Output of Land

Event Year	2007-2010	2007	2008	2009	2010
Dep Var: Sale	(1)	(2)	(3)	(4)	(5)
$\bar{\Delta} \cdot \text{Treat} \cdot (t - \tau \in \{0, 1, 2\})$	636.2*** (4.367)	561.7** (2.562)	1,003** (2.205)	393.6** (2.327)	751.3** (2.352)
$\bar{\Delta} \cdot \text{Treat} \cdot (t - \tau > 2)$	1,199*** (4.453)	1,283** (2.050)	1,836* (1.745)	736.3** (2.073)	1,342*** (2.887)
Firm FE	Yes	Yes	Yes	Yes	Yes
Province-Year FE		Yes	Yes	Yes	Yes
EventYear-Province-Year FE	Yes				
Observations	43,671	9,425	4,196	13,171	16,879
R ²	0.505	0.439	0.548	0.561	0.488

Note: This table reports estimation results of Model (17) with the matched sample. We drop the matched pairs whenever the treated or the control firm exits the sample. The first estimate is for $\theta_{\text{short-run}}$ and the second is for $\theta_{\text{long-run}}$. For each treatment year $\tau \in \{2007, 2008, \dots, 2010\}$, the sample ranges from $\tau - 4$ to 2013 (but the data for 2010 is missing). The variable sale is in 1,000 RMB and $\bar{\Delta}$ is in 1,000 m². Standard errors are clustered by firms. Robust t-statistics in parentheses. *** p<0.01, ** p<0.05, * p<0.1

concisely. In words, we separately estimate a treatment effect for the first three years after purchase, which captures the more modest effects on sales that we observe as firms presumably are making other fixed investments (e.g., new plants) that complement the land purchase, and another treatment effect for the third and subsequent years after purchase, which captures the long-run effects of new land. Formally, we estimate,

$$S_{j,t} = \alpha_j + \eta_{t,\tau_j} + \theta_{\text{short}} \cdot \mathbf{1}_{\Delta_j > 0, t - \tau_j \in \{0, 1, 2\}} \cdot \bar{\Delta}_{\tau_j} + \theta_{\text{long}} \cdot \mathbf{1}_{\Delta_j > 0, t - \tau_j > 2} \cdot \bar{\Delta}_{\tau_j} + \varepsilon'_{j,t}. \quad (17)$$

In the first column we report these estimates using land sales that occur in 2007-2010, for which we have sufficient sample to estimate long-run treatment effects. In the next four columns we report the estimated effects year-by-year.

Overall, we observe that in the first three years after land purchase, land sales generate an additional 636.2 RMB/m² in sales on average per year, and in subsequent years after land purchase, land sales generate a long-run effect of 1,199 RMB/m² in sales on average per year.

Before leaving this section we discuss one econometric issue regarding panel balance. Firms enter and exit our panel due to data linkage issues and due to firm births and deaths. For example, we use firm names as firm identifiers, so name changes or inconsistencies in name reporting can lead to panel imbalance if we fail to track a firm over time. Panel imbalance can also arise due to censoring when firm sales fall below a threshold for inclusion in our data. We address imbalance by excluding matched treated-control pairs from our analysis whenever either firm's data are imbalanced. In Appendix Section C.6 we study the causes of panel imbalance and conclude the majority of imbalance is due to idiosyncratic reasons such as imperfect matches on firm name over time. We also find evidence that a modest amount of imbalance is due to censoring, which we argue in Appendix Section C.6 makes our estimates of land-purchase treatment effects conservative.

4.3 Firm Tax Estimation

We now calculate the marginal tax revenues generated by increases in firms' sales. We first provide a simple model of vertical production network that guides us in estimating the total impact of the land purchase of one firm to the entire production chain. Essentially, the increase of taxes paid by all firms on the production chain is the sales increase of the land-purchasing firm, as estimated in Section 4.2, multiplied by an effective tax rate.

4.3.1 Marginal Tax Revenue Estimation

The main tax on industrial firms in China is the value-added tax; for a detailed description of the system of firm taxation in China, see Appendix C.3. Appendix C.4 provides a simple model of a vertical production network, demonstrating the fact that, under homogeneous tax rates across firms, the sum of value-added taxes across firms is equal to the value-added tax rate multiplied by the sum of output across final goods producing firms. Thus, if the land is sold to a final good producer, and if the output of all other final good producers remains constant, then the increase in the total tax revenues is equal to the increase in the land buying firm's sale multiplied by the marginal tax rate.

Guided by the above framework, we estimate the marginal tax income of land with

the following steps. First, we estimate firm tax rates in China in Appendix C.3. Appendix Figure A.4 shows that while tax rates differ across firms in China, the average value-added tax rate is approximately 12.10%, and this is relatively stable for firms of different sizes. We also estimate that income taxes and other administrative fees amount to approximately 5.77% of firms' value-added. Combining these estimates, firms face an average tax rate of approximately 17.87%.

In the second step, we calculate the marginal effects of land sales on the total tax revenues by multiplying the tax rate with the estimated effect of land sales on output from column (1) of Table 3. The tax effect in the first three years after the land purchase is estimated to be:

$$\text{Tax}_{t=0,1,2} = 636.2 \times 17.87\% = 113.6\text{RMB}. \quad (18)$$

The permanent increase in taxes and fees paid by the land acquiring firm as well as the upstream firms for one square meter of industrial land is estimated to be:

$$\text{Tax}_{t>2} = 1,199 \times 17.87\% = 214.2\text{RMB}. \quad (19)$$

There are two potential issues with our approach in estimating marginal tax income of land. First, our approach only applies to final goods producers. If land is sold to intermediate-goods producers, their sales may increase output and value-added further downstream in production networks, so we are likely to underestimate the tax effect of land sales for intermediate goods producers. Second, we assume that changes in the output of land buying firms do not affect sales, or input purchases, for other firms not in the same production chain. In Appendix C.5, we study a simple model building on Hulten (1978) and Baqaee and Farhi (2019) to show how changes in individual firm productivity affects aggregate output. We show that the increase in land buyers' output will tend to over-estimate changes in total output, if land purchases lead land-buying firms to increase their input purchases. Essentially, this is because land buyers cannibalize some input purchases, and hence decrease output of other firms.

4.3.2 Complementary Evidence

Our estimates of marginal tax income of industrial land sales from the government's perspective square nicely with the following two pieces of complementary evidence.


Average VAT income from industrial land. As a first benchmark, we compare our estimated marginal effect of land sales on tax receipts to the average VAT per square meter of land. For each province during our sample period, we calculate the average VAT per square meter of land as total VAT revenue from China Tax Yearbook,¹⁶ divided by total industrial land size (from China City Construction Yearbook). Figure 3a shows the average VAT income per square meter of land for each province in 2011, a year that is right after the sample period of 2007-2010 that we use to estimate the marginal taxes on land. Across all provinces, the simple average VAT income per square meter of land is 332 RMB/m². This has the same magnitude as, though is slightly larger than, the long-run tax revenues land, $Tax_{t>2} = 214.2$ RMB/m², estimated in Eq. (19).

Official guidance on minimum required tax on industrial land. As the second source of evidence on tax income, we use the government's direct guidance on the "required minimum" tax paid by firms operating on industrial land.


Although the land allocation decision is at the discretion of the city government, the central government does propose some guideline on industrial land supply. In 2008, the Ministry of Land Resources initiated the Guidelines on Land Supply to Industrial Projects, which required the local land bureaus to impose restrictions on the industrial land supply along certain dimensions (for example, a green land ratio).¹⁷ Some provincial land bureaus modified the guideline further by adding additional requirement on the tax payment by firms, with Jiangsu province being the first local government to explicitly impose an industry-specific minimum requirement on tax payments by firms on industrial land in 2018. Some provinces, such as Hunan, followed and imposed the same minimum requirement in 2020.

¹⁶We calculate the total VAT paid by firms in each province as the summation of both the local governments' and the central government's VAT revenues.

¹⁷Other restrictions are on the amount of fixed investment, floor ratio, the fraction of land for buildings and construction, and the fraction of land for offices and utilities.



(a) Total VAT Over Industrial Land for Each Province, 2011, RMB/m²



(b) Requirement on Minimum Tax Payment by Firms on Industrial Land, RMB/m²

Figure 3: Supplementary Evidence on Tax Income of Land

Notes: Panel (a) plots the total VAT paid by firms in each province divided by the stock of industrial land in that province in 2011. Panel (b) plots the industry-specific requirement on minimum tax payment by firms on the industrial land set by Jiangsu Province in 2018 and Hunan Province in 2020. Values are in RMB/m².

Figure 3b plots the industry-specific minimum requirement on annual tax payments set by Jiangsu and Hunan province for all non-tobacco-related manufacturing industries. The minimum tax requirement for most industries is around 100 RMB/m², and if we average across industries using the industrial composition of land sales in our data during 2007-2010, we find an average minimum tax requirement of 113.5 RMB/m². We thus conclude our estimate of the marginal tax revenues of 214.2 RMB/m² is comparable to these minimum requirements, further validating our estimates.

4.4 IRR Estimates and Comparison to r^{gov}

We now can use the framework of Section 3 to translate our empirical estimates into an IRR estimate on industrial land sales.

Baseline estimate of IRR^{ind} . In Section 4.1, we estimate industrial land discounts at the city-year level, $IndDisc_{ct}$. To aggregate these estimates to the national level in a way comparable with the tax estimates, we calculate the average $IndDisc_{ct}$ weighted by the number of treated firms in city c that purchased land in year t , using the estimation sample of Table 3 Column (1). The weighted-average industrial land discount during 2007-2010 is 1,176.4 RMB/m², as shown in the first row and first column of Table 4. Combining with the estimates of tax revenues, which are 113.6 RMB/m² in the first two years given by Eq. (18) and 214.2 RMB/m² thereafter given Eq. by (19), we calculate the government IRR in Eq. (2) to be $IRR^{ind} = 13.94\%$.

Different aggregation methods. To show how the aggregation method matters, we also report the simple average value of $IndDisc_{ct}$ across cities during 2007-2010, which is given in the second row and the first column in Table 4. This alternative aggregation method gives an estimate of 529.2 RMB/m², smaller than the weighted average value of 1,176.4 RMB/m². This difference is driven by the fact that in the sample of firms based on which we estimate the sale effect of land purchase, there are more firms located in cities with higher industrial discounts, which are positively correlated with GDP per capita, as shown in Figure 2. With this lower estimate of industrial discount, we obtain a higher implied government IRR of $IRR^{ind} = 26.64\%$.

Table 4: Industrial Land Discounts and IRR^{ind}

Aggregation Weight	Estimation Method			
	Quality Adj.		Simple Ave.	
	Discount	IRR^{ind}	Discount	IRR^{ind}
Number of firms in sale estimation	1176.4	13.94%	1246.2	13.28%
Not weighted	529.2	26.64%	641.4	22.83%

Note: This table shows the industrial land discount estimates during 2007-2010 and the corresponding IRR^{ind} , calculated with Equation (2), with different aggregation methods. We aggregate the land-level industrial discount estimates with two approaches: one using simple average and the other weighted by the number of treated firms in that city treated in that year in the estimation of Table 3 Column (1).

For robustness, in the second column of Table 4 we also show the aggregated industrial discount estimates using the simple average version of the industrial discount estimates (which does not control for land characteristics X_{it}). We observe that one tends to overestimate industrial discounts if one ignores the quality differences between residential and industrial lands. The implied government IRRs are slightly lower than the quality-adjusted estimates, with $IRR^{ind} = 13.28\%$ or 22.83% .

Targeted industrial policy. We conduct another robustness check to address the concern that our results are confounded by targeted industrial policy in China. The Chinese government has implemented various policies, such as lower land price, tax rebates, and subsidies on sales and investment, to support industries that are regarded as having special national importance. The implementation of these policies will incur costs to the governments, likely beyond lower industrial land sale prices.

Conceptually, focusing on the pecuniary side of this trade-off, one can estimate the IRR^{ind} for any industrial land supply independent of supporting industrial policies. Although we observe an industrial land discount specific to these policy-targeted industries, we do not observe other forms of costs that should have been deducted from future tax revenues. As a result, ignoring those costs will lead to overestimation of the “pecuniary” benefits of industrial land supply to these targeted industries, and ultimately an overestimation of the IRR^{ind} . On the other hand, estimates of IRR^{ind} that are based on

the non-targeted industries should not suffer from this upward bias (caused by industry supporting policies).

With this important conceptual issue in mind, we first identify industries that are ever targeted by the Eleventh and Twelfth China Five-year Plan, which highlights the key sectors the government plans to support during the period 2006-2015 (Cen et al. (2021)).¹⁸ In our sample, among all the treated firms, 57.0% (43.0%) are from targeted (non-targeted) industries and they account for 62.7% (37.3%) of the size of the matched industrial lands. Table 5 reports the estimated industrial land discounts, tax effect, and IRR^{ind} of targeted and non-targeted industries, respectively. The targeted industries generate much more taxes out of land than the non-targeted industries, which might partially be attributed to government support. Regarding industrial land discount, offering a greater industrial land discount could be the direct consequence of the governments' supporting policies, but it is also possible that the higher profitability of the targeted industries could lead firms to bid more aggressively for the land, and hence indirectly lowers the industrial discount. Table 5 reports that our industrial discount estimate for the targeted industries is marginally smaller than that of the non-targeted industries, suggesting the second indirect force dominates.

Taking these cash-flows estimates together, we find that the IRR^{ind} on targeted industries is 15.84%, which, as we discussed, tends to overestimate the true IRR due to overestimated future net revenues. Importantly, based on the firms in non-targeted industries that are free from Chinese industrial policies, we find the resulting IRR^{ind} to be 11.36%, which is only slightly smaller than the baseline value of 13.94% in Table 4.

Comparison to r^{gov} One of the main takeaways from our paper is that our estimate of IRR^{ind} is higher than most estimates of government discount rates in the literature, which we will call r^{gov} .

We proxy the city government's cost of capital using the issuance yield of municipal corporate bonds (MCBs, or Chengtuo Bonds in Chinese). MCBs are bonds issued by city government financial vehicles, which are state-owned enterprises, to support infrastructure investment at both the provincial and the city level.¹⁹ Since the four-trillion

¹⁸See Appendix C.7 for the list of targeted industries.

¹⁹As explained in Chen et al. (2020), MCBs have the implicit backing of the corresponding city government

Table 5: Targeted vs Non-targeted Industries

Targeted Industries?	Land Discount	Cash-flows		IRR ^{ind}
		Tax _{t=0,1,2}	Tax _{t>2}	
Yes	1,113.9	121.5	239.6	15.84%
No	1,223.1	98.9	171.8	11.36%

Note: This table reports the IRR^{ind} for the two groups of industries based on whether they are ever targeted by the Eleventh and Twelfth China Five-year Plan that covers the period 2006-2015. For each group, we estimate Model (17) separately to get the tax estimates, calculate the city-year level industrial discount using residential land and only industrial land sold to firms in the targeted (or non-targeted) industry, and calculate the average discount weighted by number of firms from the targeted (or non-targeted) industries in the sale estimation.

stimulus plan in 2009, MCBs have become the major financing source for Chinese city governments besides selling land directly (Bai et al., 2016; Chen et al., 2020), and their market-determined yields reflect the city governments' fiscal conditions.²⁰

We find that our estimate of IRR^{ind} is larger than city governments' cost of capital r^{gov} , which ranges from 3.5% to 7.5%, by a significant margin. This suggests that city governments' sharing of tax revenues with other governments plays an important role in the equilibrium land allocation decisions. If the city government received the entirety of VAT revenues, then the discussion of governments' market power over residential and industrial land in Section 3 (see Eq. (4)) indicates that IRR^{ind} should be smaller than city governments' cost of capital. In contrast, the fact that city governments only retain a fraction of tax revenues can explain why IRR^{ind} is higher than r^{gov} .

Thus, we posit that the high IRR on industrial land sales is driven by the intersection of three forces: the "land finance" system, in which the revenues from land sales accrue

(hence the name municipal), but in a strict legal sense they are issued by LGFV entities just like other regular corporations (hence corporate).

²⁰We do not use the yields of municipal bonds for three reasons. First, in China, the official municipal bonds (i.e., issued by the municipal directly) were with rather limited supply before central government launched the second major tax reform in 2014. Second, municipal bonds issued by local municipals have a experienced dramatic increase since 2015, but they are explicitly guaranteed by the central government, which removes any risk premia associated with fiscal conditions of municipals. Finally, municipal bonds are subject to strict issuance quotas, and hence do not serve as the marginal financing method for city governments.

entirely to city governments and are an important source of governments' operational funds; the distinct time profiles of revenues from industrial and residential land sales along with the governments' discount rates; and the asymmetric treatment of industrial tax revenues, which are shared between city governments and upper-level governments. We empirically analyze the association between tax sharing schemes, city governments' discount rates, and industrial discounts in the following section.

5 City-level Evidence: Tax Shares and Discount Rates

In this section we exploit cross-sectional heterogeneity at the city level and an event study across cities to present two pieces of empirical evidence that are consistent with our perspective of public finance in explaining the industrial land discount. Consistent with the theoretical framework in Section 3.2, we show that a city has a greater industrial land discount if it has a larger VAT share, or a lower discount rate as measured by city governments' municipal corporate bond yields.

5.1 City Tax Shares

As we discussed in the previous section, IRRs on industrial land sales are higher than city government discount rates, suggesting that city governments do not fully internalize the tax revenues generated by industrial land sales. This is consistent with the institutional details of value-added taxes in China: city governments directly receive almost the entirety of revenue from land sales, but only directly receive a small fraction of revenues from firm taxes. Value-added tax revenues largely accrue to central governments: prior to 2016, 75% of the VAT was allocated to the central government, and the remaining 25% was shared between province-level governments and city-level governments. As a result, during 2007-2010, the average share of VAT that accrues to city-level governments was 18.9%. Combining this rate with similar tax-sharing rates for corporate income taxes, we find that the effective share of all firm taxes that accrue to city-level governments is

22.8%.²¹ That is, city governments only receive 22.8% of the tax revenues generated by industrial land sales.

In this subsection, we investigate the relationship between industrial discounts and the share of value-added taxes that accrue to city governments. In each province, the province-level government has the discretion in setting how to split the taxes between itself and the city-level governments, and there is variation in the share of VAT accruing to the city governments in different provinces.²² Although the actual share of VAT that accrues to the city governments may underestimate the extent to which the city governments internalize tax revenues from industrial land sales,²³ we assume the city VAT share is at least positively correlated with the extent to which governments internalize future tax revenues in their land allocation decisions. We first present cross-sectional evidence that industrial discounts are sensitive to the city VAT shares, and we then conduct an event study exploiting a change in tax sharing between the central government and local governments in 2016.

5.1.1 Cross-Sectional Evidence

In Figure 4, we show a binned scatterplot of industrial land discounts against city VAT shares. The relationship is positive: industrial discounts are higher in cities where city governments receive a larger share of VAT taxes, consistent with our theoretical predictions.

We then test this hypothesis by regressing industrial discounts on the share of value-added taxes which accrues to city governments, using the following specification:

$$\text{IndDisc}_{ct} = \beta \text{VATShare}_{ct} + \gamma \text{GDPPercapita}_{c,t-1} + \eta_t + \epsilon_{ct}, \quad (20)$$

²¹Corporate income taxes are similar: 60% accrues to the central government, and within each province, the provincial and city governments split the remaining 40%. During 2007-2010, the average share of the corporate income tax that accrues to city-level governments was 30.9%.

²²Wu and Zhou (2015) shows that the city government VAT share tends to be higher if there is less variation in economic development across cities in the province, if the industrial sector in the city is more developed, and if there are less state-owned firms under the control of the province governments.

²³For example, city leaders may partially internalize the welfare of provincial or central governments, either via promotion incentives or bargaining process in revising land allocation plans.

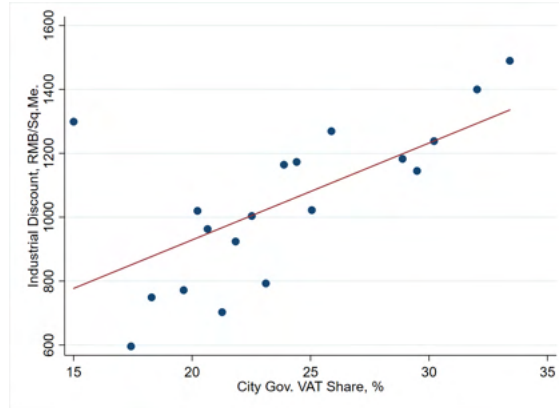


Figure 4: Industrial Land Discount and City Government’s VAT Share

Note: This figure plots bin scatter of city-level industrial discounts, $IndDisc_{ct}$, against the city government’s VAT share for the sample period 2007-2019.

for city c in year t . That is, we regress industrial land discounts on city VAT shares, with time fixed effects η_t and controlling for the level of GDP per capita of city c in year t . Because we initially focus on the variation across cities, this regression does not include a city fixed effect; we relax this later in Section 5.1.2.

Table 6 reports the result, where we divide the sample into two periods. In the first two columns, we use sample period 2007-2010 to match the baseline estimation of the sale effect in Table 3, while in the last two columns we use all the years after 2010. The coefficient on VAT taxes is positive and significant across both sample periods, and the magnitude of the coefficients only changes slightly when we control for the level of economic development using GDP per capita.

5.1.2 Event Study: the 2016 VAT Tax Change

Next, we present evidence on the relationship between VAT shares and industrial discounts using an event study. We analyze a change in tax-sharing schemes that occurred in 2016, which increased the share of VAT revenues accruing to city governments heterogeneously across provinces. We show that provinces which had a larger increase in the city government’s share of tax revenues experienced larger increases in the industrial land discount.

Table 6: Industrial Land Discount and City Government's VAT Share

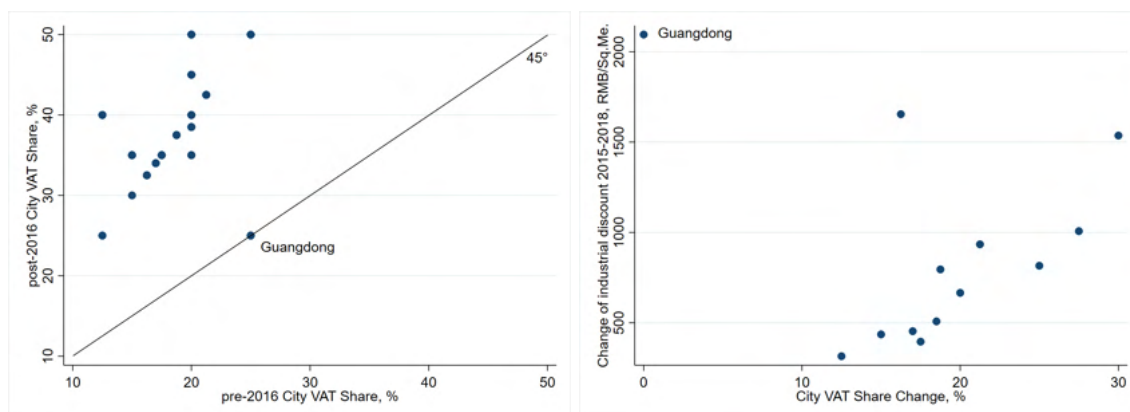
Sample Period	2007-2010		2011-2019	
	(1)	(2)	(3)	(4)
Dep Var: Industrial Discount				
Vat Share	30.53*** (2.829)	32.33*** (3.559)	30.28*** (2.963)	31.71*** (3.732)
GDP per capita		160.0*** (4.246)		138.1*** (4.729)
Year FE	Yes	Yes	Yes	Yes
Observations	849	849	1,942	1,907
R ²	0.083	0.232	0.178	0.336

Note: This table shows the correlation between the city government's VAT share and the quality-adjusted industrial land discount. City VAT share is in percent and GDP per capita is in 10,000 RMB/person. Standard errors are clustered by cities. Robust t-statistics in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

We mentioned in the beginning of Section 5.1 that prior to 2016, roughly 25% of value-added tax revenues was shared between province- and city-level governments, and there existed a significant heterogeneity in city VAT shares due to various economic and political factors. On May 1, 2016, the central government launched a major tax code change – the so-called "Business to Value-added" program – which enlarged the coverage of value-added taxes. More importantly, this reform modified the tax-sharing scheme between central and local governments, such that the share of value-add taxes retained by the local governments increased from 25% to 50%.²⁴ The province-level government would then decide how to split the incremental 25% of the value-added taxes between itself and the city governments.

City VAT Share Changes and Industrial Land Discounts: Raw Data Panel A in Figure 5 shows the pre-2016 city VAT share on the x-axis, and the post-2016 city VAT share on the y-axis. Most cities experienced a rise in their share, except for cities in Guangdong whose share remained at 25%; we will explain the special circumstance of Guangdong shortly. There is also substantial heterogeneity in the magnitude of the tax share increase

²⁴Local governments previously received the entirety of business taxes. After the business taxes were replaced with value-added taxes and shared by the central government, the central government increased the VAT share of the local governments to keep their fiscal revenue stable.



(a) City VAT Share Pre- and Post-2016 (b) Change of Ind. Discount vs VAT Share

Figure 5: Change of City VAT Share and Industrial Land Discount

Notes: Panel (a) plots the city government’s share of VAT before and after 2016. Most cities within the same province receives the same share with very few exceptions. Panel (b) plots a binscatter of the change of city-level industrial land discount from 2015 to 2018 against the change of city VAT share.

across cities, allowing us to investigate how industrial discounts respond to their VAT shares. Indeed, Panel B in Figure 5 shows a binned scatterplot of the change in the key variable of interest, the industrial land discount, from 2015 to 2018 relative to the city VAT share change in 2016. There is a strong positive correlation between the two variables (without counting cities in Guangdong).

In both panels of Figure 5 We also observe that cities in Guangdong province appear to be outliers. Although they experienced zero increase in their share of VAT, the industrial land discount increased substantially from 2015 to 2018 in these cities. One possible explanation is confounding policies that also encouraged industrial land supply in Guangdong. On August 20, 2017, the provincial government of Guangdong initiated a list of actions to secure the industrial land supply by the city government. All these actions are taken by Guangdong only, and are not in place before 2017. Appendix D.1 provides more details on the land-related policies for Guangdong province. Due to these factors, we remove Guangdong from our analysis in the rest of this section.

Dynamic treatment effect. We apply a straightforward difference-in-differences estimation strategy to study how local governments' land allocation decisions respond to these changes in city VAT shares:

$$y_{ct} = \alpha_c + \gamma_t + \sum_{\tau \neq 2015} \beta_{\tau} \times \mathbf{1}_{t=\tau} \times \Delta \text{VATShare}_c + \varepsilon_{ct}, \quad (21)$$

for city c in year t , with city and year fixed effects. In Eq. (21), we use the year before the taxation change, 2015, as the base year. We also include interactions with years before 2015 to test the assumption of parallel trends between cities that experiences different changes in VAT shares.

If city governments' land allocation decisions are indeed sensitive to tax revenues, then as the share of tax revenues accruing to the city government increases, they should be willing to offer a higher industrial land discount.²⁵ The estimation results reported in Table 7 support this hypothesis. We observe a significant and positive treatment on the industrial land discount in all the years since 2016. Moreover, there was no significant difference between cities with differential treatment prior to 2016 (with only a few occasional exceptions), which lends support to the parallel trends assumption underlying this difference-in-differences strategy.

In Column (2) and (3), we investigate the industrial and residential land price separately. As we can see, the effect is mostly driven by an increase of the residential land price. This could be due to the fact that the demand elasticity for industrial land is higher than that for residential land. Industrial firms typically shop around different cities for the most favorable land price, but most households do not move across cities and demand for residential land is more fixed; therefore, residential land price shall adjust in a significant way when the supply changes.

²⁵There are two possible mechanisms through which the industrial discount adjusts. First, the city government may allocate more industrial land relative to residential land in the future, with an immediate adjustment in prices (and hence industrial discounts). The quantity adjustment may not occur in the short run given the planning constraint; see Section 2. Second, if the government and the potential buyers can negotiate on the land transaction, the buyer who knows that more future taxes go to the local government may ask for a greater industrial discount.

Table 7: City VAT Share and Industrial Land Discount - DID Estimation

Dep Var:		Industrial Discount	$(1 - \lambda)p^{res}$	p^{ind}
		(1)	(2)	(3)
$\Delta VATShare \times$				
Year=	2010	13.52 (0.845)	10.49 (0.652)	-3.033* (-1.945)
	2011	7.833 (0.606)	5.854 (0.437)	-1.979 (-1.368)
	2012	-6.958 (-0.621)	-7.122 (-0.623)	-0.164 (-0.158)
	2013	-9.484 (-0.864)	-9.185 (-0.814)	0.299 (0.366)
	2014	-13.06 (-1.501)	-12.21 (-1.344)	0.852 (1.006)
	2016	28.08** (2.225)	27.49** (2.172)	-0.592 (-0.787)
	2017	59.18** (2.548)	58.27** (2.430)	-0.903 (-0.647)
	2018	54.66*** (2.623)	54.91*** (2.608)	0.259 (0.164)
	2019	45.93** (2.401)	48.23** (2.481)	2.305 (1.270)
Observations		1,885	1,885	1,885
R ²		0.858	0.871	0.857
Year FE		Yes	Yes	Yes
City FE		Yes	Yes	Yes
#City		204	204	204

Note: This table shows how the change of the city VAT share affects the industrial land discounts. The sample includes all the municipal cities for which we have the industrial discount estimates from 2010-2019, and the year 2015 is used as the baseline. We drop few city-year observations before year $t < 2016$ if the city's share of VAT changed in year t . The treatment variable, $\Delta VATShare$, is in percentage. Standard errors are clustered by cities. Robust t-statistics in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Discussion on economic magnitudes Table 7 allows us to gauge the economic importance of city VAT share in explaining the observed industrial land discounts. During 2010-2015, the standard deviation of industrial discounts and city VAT share across city-year is 768.6 RMB/m² and 6.36%, respectively. Thereby, the city VAT share can explain approximately 44% ($=\frac{6.36 \times 53.55}{768.6}$) variation of industrial discounts.

We can also link the magnitude of the estimated effect in (21) to the marginal tax revenue of industrial land estimated in Section 4.3, which are 113.6 RMB/m² in the first three years in Eq. (18) and 214.2 RMB/m² permanently afterwards in Eq. (19). If we take the average government borrowing rate of 5.25% as a proxy for the government discount rate,²⁶ then the present value of total tax revenues is 3,824 RMB/m². The estimation in (21) then implies that an increase of 1% of city VAT share allows the city government to get 38.24 RMB/m² more of taxes from the industrial land. This is fairly close to our estimate in Table 7, where we find a 1% increase of city VAT share is associated with an increase of 53.55 RMB/m² (if we combine the three post-treatment year 2007-2019 together) in industrial land discount.²⁷

Together, the results in this section suggest that governments' land allocation decisions are sensitive to the share of tax revenues they receive: increasing the share of value-added tax revenues accruing to local governments tends to increase the industrial land discount.

5.2 Government Discount Rates

If governments' land sale decisions reflect intertemporal revenue trade-offs, then changes in city governments' discount rates should affect industrial land sale decisions: less constrained city governments should sell less residential land and more industrial land, increasing industrial discounts. We test this hypothesis by analyzing the relationship between city governments' discount rates, proxied by the corresponding MCBs' issuing yields, and their industrial discounts. Consistent with our hypothesis, Figure 6 shows that MCB yields and industrial discounts are strongly negatively correlated in the cross-

²⁶By 2021, the average coupon rates of all municipal corporate bonds, which are the bonds issued by local government financing vehicles, is 5.25%.

²⁷The difference between the two numbers, 38.24 vs. 53.55, can be attributed to the government market power in the land market, which we discuss in details in Appendix B.




Figure 6: Industrial Land Discount and Average Municipal Corporate Bond Yields

Note: This figure plots bin scatter of industrial discounts and the average yields of municipal corporate bonds issued by the governments in that city for the sample period 2007-2019.

section.

A concern for interpreting Figure 6 is that MCB yields may be endogeneous: certain forces may affect both MCB yields and industrial discounts, so the cross-sectional correlation between MCB yields and industrial discounts may not reflect the causal effect of government discount rates on industrial discounts. To address this concern, we build on [Chen et al. \(2020\)](#) and use an instrumental variable for MCB yields related to China’s four-trillion stimulus plan in 2009. The instrument uses local political officials’ job tenure at the launch of the stimulus; [Chen et al. \(2020\)](#) show that cities in provinces with governors who were late in their term engage in more local infrastructure investment in 2009, which has long-lasting effects on the local government’s fiscal position in the future and hence on future bond yields. Local government officials’ tenure is plausibly related to investment choices in 2009 because the incentive to comply with the central government in general increases with the governor’s term.²⁸

Following [Chen et al. \(2020\)](#), we construct an instrument, $LateTerm_c$, which takes a

²⁸More broadly, this instrumental variable is motivated by the existing literature on China’s political economy that links local government officials’ promotion to their incentives of pursuing local economic growth during different years of their official terms ([Ru, 2018](#); [Liu et al., 2018a](#)). What is more, the city government, which is our unit of observation, has a strong incentive to comply with his or her provincial governor’s political agenda, because of China’s “one-level-up” policy: the promotion of the local officials is largely determined by their immediate superior officials ([Chen and Kung, 2019](#)).

value of one if the city c 's provincial governor had been in office for at least three years in the beginning of 2009, and zero otherwise. The first stage is strong and statistically significant: LateTerm_c is negatively correlated with the MCB yield in subsequent years, and in particular during the period 2012-2019, when we are able to calculate the industrial discount; the first-stage F-statistic for this time period is 28.1. This negative sign is consistent with greater infrastructure investment in 2009 leading to a stronger future fiscal position, for example in the form of greater land reserve values (both for sale, and for the government to use as collateral).²⁹ The exclusion restriction for the instrument then is that these fiscal changes only are correlated with the future industrial discounts through changes in MCB bond yields. In particular, the exclusion restriction requires that the size of a city's land reserve, which can be developed for both industrial and residential purposes, does not directly affect the choice of what mix of residential or industrial land to sell.³⁰

We then use the tenure instrument to estimate the causal effect of MCB yield shifts on industrial discounts, using the following specification:

$$\text{IndDisc}_{ct} = \beta \times \text{MCBYield}_{ct} + \sum_{\tau} \gamma'_{\tau} \cdot \mathbf{1}_{t=\tau} \cdot X_{c,2008} + \epsilon_{ct}, \quad (22)$$

where MCBYield_{ct} is the average yields of MCB bonds issued by city c year t and weighted by issuance amounts,³¹ and where we instrument MCB yields using LateTerm_c . To separate our estimation sample from the potential direct effect of the governor term in 2009, we estimate Eq. (22) based on the sample period 2012-2019.³² We control for time-varying effects of initial city-level economic conditions using $X_{c,2008}$, which includes

²⁹See Section 2 and A.1 for the discussion of land finance in China and the role of land reserves in shaping local government's fiscal position.

³⁰Chen et al. (2020) show that provinces with greater stimulus bank loans in 2009, due to the future refinancing needs, experience faster MCB growth and more shadow banking activities during 2012-2015. Chen et al. (2020) concern a pure quantity implication, while the price implication of 2009 stimulus bank loans on future MCB yields is ambiguous, exactly because of the expanded land reserves mentioned here.

³¹Our sample includes all bonds identified as MCBs by either China Banking and Insurance Regulatory Committee or Wind.

³²One concern on the exclusion restriction is that the governor term in 2009 would predict the terms of future provincial governors, which could have direct effect on the land supply. Chen et al. (2020) shows that thanks to the anti-corruption campaign launched in 2012 by Xi Jinping, there is negligible correlation between governor term in 2009 and governor term in future years after 2012.

Table 8: Industrial Land Discount and Municipal Corporate Bond Yield

Specification	OLS	OLS	IV	IV
Dep Var: Industrial Discount	(1)	(2)	(3)	(4)
MCBYield, %	-561.5*** (-8.870)	-362.3*** (-6.127)	-2,162*** (-6.859)	-3,259*** (-3.278)
Controls	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	1,386	1,386	1,386	1,386
R-squared	0.348	0.450	-1.477	-3.994
#City	257	257	257	257
F statistic			28.08	8.194

Note: This table shows the regression of industrial land discounts on City MCB yields, i.e., the average yields of MCBs weighted by the bond size. The first two columns report the OLS estimation results and the last two columns report the 2SLS estimation results where the City MCB yield is instrumented by $LateTerm_c$, i.e., an indicator of whether the provincial governor had been in office for at least three years in the beginning of 2009. The sample period is from 2012-2019. Standard errors are clustered by cities. Robust t-statistics in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

the GDP per capita, the growth rate of GDP from the previous year, and the fiscal deficit over GDP, all measured in the year 2008.

The results are shown in Table 8. The first two columns report OLS estimation results, confirming the negative correlation in Figure 6. In Column (3) and (4), we instrument $MCBYield_{ct}$ with $LateTerm_c$ and find a significantly negative causal effect of bond yield on the industrial land discounts.

The association between bond yields and industrial discounts highlights how city governments' land allocation decisions can be entangled with their liquidity management needs. City governments with more of a liquidity shortfall or greater financial constraints may reallocate land sales from industrial to residential purposes, trading future cash flows from industrial taxes to more immediate cash flows from residential sales under the "land finance" system. While these patterns are evident in the cross-section across geography, they also may matter over time too, for example suggesting if distress were to emerge in Chinese municipal debt, reductions in industrial land supply (absent any tax reform to correct for this) may be an important knock-on effect.

6 Conclusion

In this paper, we analyze the industrial land discount in the Chinese land market. Counter to conventional wisdom, the return of supplying industrial land instead of residential land, accounting for all the future tax revenues the industrial land generates, is higher than the usual range of government discount rates proxied by the city governments' MCB yields. The higher total return on industrial land sales is potentially driven by the fact that city governments directly receive all upfront revenue from land sales, but only a fraction of the revenues from taxes paid by industrial firms. We show evidence for this hypothesis from the 2016 increase in cities' shares of industrial tax revenues. Our results have implications for understanding the drivers of land prices in China, and how they are linked to the tax sharing scheme with the central government, as well as local governments' intertemporal revenue tradeoffs. From the central government's perspective, the tax sharing scheme between the central and local governments can be carefully designed to counteract the effect of the local governments' differential market powers in the local land market to achieve desired land allocation outcomes.

References

- Ahern, Kenneth R, 2021, The business of city hall, Technical report, National Bureau of Economic Research.
- Bai, Chongen, Chang-Tai Hsieh, and Zheng (Michael) Song, 2016, The long shadow of china's fiscal expansion, *Brookings Papers on Economic Activity* 129–165.
- Bai, Yan, Keyu Jin, and Dan Lu, 2019, Misallocation under trade liberalization, Technical report, National Bureau of Economic Research.
- Baqae, David Rezza, and Emmanuel Farhi, 2019, The macroeconomic impact of microeconomic shocks: beyond hulten's theorem, *Econometrica* 87, 1155–1203.
- Berk, Jonathan, and Peter DeMarzo, 2017, *Corporate Finance: Global Edition* (Pearson).

- Cai, Hongbin, J. Vernon Henderson, and Qinghua Zhang, 2013a, China's land market auctions: evidence of corruption?, *The RAND Journal of Economics* 44, 488–521.
- Cai, Hongbin, Bo Li, Zan Lin, Ge Wu, Lin Xu, Yu Xi, Li Yuan, and Chengjun Zhou, 2013b, On the land reform and new urbanization, *Journal of Financial Research (in Chinese)* 3, 114–125.
- Cao, Guangzhong, Changchun Feng, and Ran Tao, 2008, Local “land finance” in china's urban expansion: challenges and solutions, *China & World Economy* 16, 19–30.
- Cen, Xiao, Vyacheslav Fos, and Wei Jiang, 2021, A race to lead: How chinese government interventions shape the us-china production competition, Working paper, Columbia University.
- Chen, Ling, and Barry Naughton, 2016, An institutionalized policy-making mechanism: China's return to techno-industrial policy, *Research Policy* 45, 2138–2152.
- Chen, Ting, and James Kai-sing Kung, 2019, Busting the “princelings”: The campaign against corruption in china's primary land market, *The Quarterly Journal of Economics* 134, 185–226.
- Chen, Ting, Laura Liu, Wei Xiong, and Li-An Zhou, 2018, Real estate boom and misallocation of capital in china, *Working Paper, Princeton University, Princeton, NJ* .
- Chen, Zhuo, Zhiguo He, and Chun Liu, 2020, The financing of local government in china: Stimulus loan wanes and shadow banking waxes, *Journal of Financial Economics* 137, 42–71.
- Fan, Jianyong, Jiawei Mo, and Jipeng Zhang, 2015, Housing models and urbanization in china-empirical research from the perspective of land supply, *Social Sciences in China (in Chinese)* 44–63.
- Fei, Xuan, 2020, The misallocation in the chinese land market, SSRN Scholarly Paper ID 3716540, Social Science Research Network.
- Fu, Shihe, Xiaocong Xu, and Junfu Zhang, 2021, Land conversion across cities in china, *Regional Science and Urban Economics* 87, 103643.

- Heinrich, Torsten, Jangho Yang, and Shuanping Dai, 2020, Levels of structural change: An analysis of china's development push, Technical report.
- Hsieh, Chang-Tai, and Peter J Klenow, 2009, Misallocation and manufacturing tfp in china and india, *The Quarterly Journal of Economics* 124, 1403–1448.
- Hulten, Charles R, 1978, Growth accounting with intermediate inputs, *The Review of Economic Studies* 45, 511–518.
- Lei, Xiaoyu, and Liutang Gong, 2014, Industrialization and urbanization based on land supply, *Management World (in Chinese)* 29–41, ZSCC: 0000042.
- Li, Hongbin, Prashant Loyalka, Scott Rozelle, and Binzhen Wu, 2017, Human capital and china's future growth, *Journal of Economic Perspectives* 31, 25–48.
- Li, Ming, 2019, Information and corruption: Evidence from china's land auctions, Working paper.
- Lin, George CS, and Fangxin Yi, 2011, Urbanization of capital or capitalization on urban land? land development and local public finance in urbanizing china, *Urban Geography* 32, 50–79.
- Liu, Chang, and Wei Xiong, 2020, 7. china's real estate market, in *The Handbook of China's Financial System*, 183–207 (Princeton University Press).
- Liu, Ernest, 2019, Industrial policies in production networks, *The Quarterly Journal of Economics* 134, 1883–1948.
- Liu, Laura, Haibing Shu, Shujing Wang, and KJ Wei, 2018a, The political cycle of corporate investments: new evidence from chinese manufacturing firms, in *Working Paper* (Peking University).
- Liu, Qing, and Yi Lu, 2015, Firm investment and exporting: Evidence from china's value-added tax reform, *Journal of International Economics* 97, 392–403.
- Liu, Xiaoping, Guohua Hu, Yimin Chen, Xia Li, Xiaocong Xu, Shaoying Li, Fengsong Pei, and Shaojian Wang, 2018b, High-resolution multi-temporal mapping of global urban

- land using landsat images based on the google earth engine platform, *Remote sensing of environment* 209, 227–239.
- Liu, Yansui, Fang Fang, and Yuheng Li, 2014, Key issues of land use in china and implications for policy making, *Land Use Policy* 40, 6–12.
- Nie, Huihua, Ting Jiang, and Rudai Yang, 2012, A review and reflection on the use and abuse of chinese industrial enterprises database, *World Economy (in Chinese)* 142–158.
- Qu, Maohui, and Zhifang Zhou, 2009, A study on compensation standard for expropriation in china, *Chinese Journal of Law (in Chinese)* 3, 163–177.
- Ru, Hong, 2018, Government credit, a double-edged sword: Evidence from the china development bank, *The Journal of Finance* 73, 275–316.
- Tang, Maogang, Ruihan Zhang, Zhen Li, and Baijun Wu, 2021, Assessing the impact of tradable discharge permit on pollution reduction and innovation: micro-evidence from chinese industrial enterprises, *Environment, Development and Sustainability* 1–23.
- Tian, Wenjia, Zhi Wang, and Qinghua Zhang, 2020, The visible hand and the invisible hand in china's industrial land market post-2007, SSRN Scholarly Paper ID 3573226, Social Science Research Network.
- Tian, Wenjia, Jingwen Yu, and Liutang Gong, 2019, Promotion incentives and industrial land leasing prices: A regression discontinuity design, *Economic Research Journal (in Chinese)* 54, 89–105.
- Wu, Min, and Li-An Zhou, 2015, Tax sharing rates among sub-provincial governments in china: facts and explanation, *Journal of Financial Research (in Chinese)* 10, 64–80.
- Wu, Wenfeng, Chongfeng Wu, and Oliver M Rui, 2012, Ownership and the value of political connections: Evidence from china, *European Financial Management* 18, 695–729.
- Xie, Zhenfa, Kairong Zhu, and Pei Li, 2019, Tax sharing, fiscal incentives and urban land allocation, *Economic Research Journal (in Chinese)* 54, 57–73.

Online Appendix

A Supplementary material for Section 2

A.1 Further Institutional Details

Pricing mechanisms. There are three major means to allocate lands: administrative allotment; agreement; and auctions, which are further subdivided into tenders, auctions, and listings. Auction-based pricing is required if the land parcel is to be used for commerce, tourism, entertainment, or commercial residences. Auctions must also be used for industrial land sales, when two or more land users show interest upon the publication of the land supply plan. The government can choose one of the three distinct auction formats (tender, auction, and listing), which have different degrees of pricing competitiveness, as stated in the national regulation guide, titled [Provisions on the Assignment of State-owned Construction Land Use Right through Bid Invitation, Auction and Quotation](#). Li (2019) argues that the choice of auction format is used by local governments in the process of corrupt distortions to land allocation. Agreement-based pricing is much less competitive than auctions, but must meet stricter requirements (see [Provisions on the Agreement-based Assignment of the Right to Use State-Owned Land](#)). The third allocation method, “administrative allotment”, is primarily used for non-profit public utility usage; examples of such uses include military use, municipal infrastructure, energy and power industries, schools, hospitals, and other public facilities.

Land types. All land contracts in China are long-term leases for land use rights, which grant the holder secured tenure and control over the piece of land for a limited period. The length of a land lease typically depends on the types of the land use. For example, the leasehold for residential land is usually 70 years, while for industrial and commercial lands the leaseholds are typically 50 or 40 years. Detailed regulations on different types of land are laid out in the [Interim Regulations of the People’s Republic of China Concerning the Assignment and Transfer of the Right to the Use of the State-owned Land in the Urban Areas](#). Our study excludes commercial lands, which are much smaller (less than 10%) in total size compared to residential and industrial lands. Moreover, a large fraction

of commercial lands are sold with designated purposes to accompany certain residential needs.

Land reserve system and the cost to land supply. The land reserve system emerged in the late 1990s in a few cities in China, and was formalized by the central government in 2001. After the central government issued the “[Measures for Land Reserve Administration](#)”, municipal land reserve centers were granted an effective monopoly on supplying land to the market.³³ It is stated that the purpose of the land reserve system is “to enhance the control of lands and the regulation of land market, and ultimately to allocate land resources efficiently”. Municipal land reserve centers are the only executing government institutions that are responsible for procuring “raw land”, preparing them to “sell-able land”, and reserving them as land reserve assets for the city (to use or to sell later).³⁴ As explained in 2.1, land finance is critical in local government public finance, the amount of land reserves grows ever bigger over time. According to the report ([link here](#)) from the Chinese National Audit Office, as of June 2013, the land reserve asset held by 34 major Chinese cities totaled to the amount of 16 trillion m².

Land reserve centers face some cost for acquiring land parcels for redevelopment and resale. The potential reserve cost of each land can be complicated to calculate, varying with the prior status and usage of the land. However, generally the redevelopment cost has a “fixed” component, which mainly includes the “standard” compensation for land, as well as the cost of land exploitation and an “non-standard” compensation component. The official rules for the “standard” compensation price are based on the prior usage value of the land: this price usually is set equal to a certain multiple of the the land’s annual output (which mostly consists of agricultural goods). This tends to be fairly standardized, though there is some variation at the city level ([Qu and Zhou, 2009](#)).

³³The formal “Measures” was first issued in 2007 (with the most recent update in 2018), preceded by two related regulations in 2001 and 2006, respectively. Before the establishment of the land reserve system, various de facto land occupiers (for example, state-owned enterprises) could effectively supply lands for private leasing purposes.

³⁴In China, the ownership of lands are all public (either state-owned or collective-owned), but the occupants of the land might have various types of usage rights. By regulation, only state-owned construction land are “marketable”. Therefore, depending on the prior status of each piece of “raw land”, the government needs to procure it through consolidation, expropriation, acquisition, repossession, and replacement (most of those processes involve both monetary and non-pecuniary compensations) before it can be developed into “sell-able land”. See more details in “[Measures for Land Reserve Administration](#)”.

Unsurprisingly, compensation based on prior use values created discontent to prior owners of land parcels when the sale price of redeveloped land parcels greatly exceeded the compensation paid to previous owners. Thus, local governments gradually started to introduce compensation elements based on post-redevelopment usage values in their agreements with incumbent landowners. While there is wide variation in how these schemes are implemented across cities, these schemes are essentially proportional revenue sharing schemes: local incumbent land users are given either a fraction of redeveloped land, houses on the land, a share of net land profits. In other cases, occupants and developers negotiate directly. We collectively refer to these as “non-standard” costs of land procurement, and we explain how we estimate them from the data in Appendix C.1.

A.2 Data Cleaning

Land data. Our land sale data is from the Ministry of Natural Resources. We adopt the following procedures to remove outliers. First, the recorded size of a number of land parcels is above 10 million square meters, which are probably errors. We correct it by dividing the size by 10,000, which is the standard multiplier in Chinese unit systems. Second, the recorded price of a number of land parcels is over 100,000 yuan per square meter, which are also errors. Similarly, we scale the price down by 10,000.

We retrieve geographical coordinates of each land parcel by inputting their street addresses into the Gaode maps API. To verify the accuracy of the retrieved coordinates, we collect the Gaode address corresponding to the retrieved coordinates and compare them with the raw address in the land-sale data. We keep lands for which the Gaode address and the raw address are in the same town.

Firm data. Our firm data is from the from the NIE database, collected by the Chinese National Bureau of Statistics. There is no consistent firm identifier in the NIE database that is non-missing in all years. We thus rely on firm names to match firms across years. The database is censored from below, in the sense that one industrial firm will enter the database only in years when its annual sale exceeds certain threshold, and if in the next year its annual sale falls below the threshold, it will not be in the database for that year.

We can lose track of a firm in the NIE database not only due to censoring, but also to

other reasons such as the data collecting process or changing firm names. In appendix [C.6](#), we show that most of the firms dropping out of the sample can be attributed to random dropping, rather than censoring. This suggests that our DID estimates of the sales effect of land purchases should not be biased substantially by dropped firms.

Merging. We merge the land-sale data with firm data by the name of land buyers. We merge not only land parcels directly bought by the firm, but also those bought by the firm's immediate controlling subsidiaries (ICS), and the ICSs of the firm's ICSs, and so forth. We define firm A as firm B's ICS if firm B has at least a 50% equity share in firm A. The ownership data come from firm registry information which covers the population of firms in China. Appendix Table [A.1](#) shows how the merged sample compares to the full samples of land parcels and firms.

A.3 Constructing Urban Units

Cities are a relatively large unit of geography, and cities may have multiple clusters of developed land with different prices. To account for this possibility, we divide cities into "urban units". To do this, we use geographic data from [Liu et al. \(2018b\)](#), who use Google Earth images to classify 30m×30m cells as urban or non-urban land, where urban land refers to impervious surface such as pavement, concrete, brick, stone and other man-made impenetrable cover types. We then cluster urban land into contiguous blocks, using the ArcGIS function `arcpy.AggregatePolygons_cartography`. Essentially, this function produces blocks of land, iteratively connecting blocks to form larger blocks, as long as they are within a specified distance of each other. The function has two parameter settings: the maximum permitted separation distance between units, which we set as 1 mil, and the maximum area of holes to fill, which we set as 1 square mile. We keep urban units of size bigger than one square mile, extract their centroids, and map each land parcel to the closest urban area centroid.

In Appendix Figure [A.1](#), we first show the distribution of urban units throughout the country. A larger fraction of land is covered by these urban units in the more developed coastal areas, especially the Circum-Bohai Sea Region, the Yangtze River delta, and the Pearl River Delta. In Panel B we use Shanghai as an example to show the urban units

Table A.1: Summary Statistics of Industrial Lands and Land Buying Firms

	Obs	Mean	Std Dev	Obs	Mean	Std Dev
A. Industrial Lands Characteristics	Sample, 2007-2010			Population, 2007-2010		
Land price per square meter (yuan)	22,566	207.74	217.96	122,901	180.77	284.72
Area (1,000 m ²)	22,636	38.16	50.23	124,340	39.04	103.88
Distance to urban unit centers (km)	22,636	10.69	9.9	124,341	10.92	11.29
B. Firm Characteristics	Merged Firms, 2003-2013			All Firms, 2003-2013		
Sales revenue	70,466	260.4	1,206.591	2,151,097	178.87	1,626.65
Sales cost	70,464	222.75	1,085.41	2,150,925	151.88	2,141.12
Total assets	70,462	210.89	1,221.79	2,151,003	151.74	2,113.02
Gross value of industrial output	70,326	264.85	1,126.2	2,148,079	179.89	1,543.06
Enterprise income tax	60,334	2.75	36.92	1,969,737	1.9	38.38
Value-added tax	68,429	7.58	53.34	2,115,965	5.9	89.94
Sales tax and surtax	68,603	1.84	32.06	2,122,431	2.72	146.12
Total profit	70,345	16.71	91.4	2,149,174	11.91	269.09
Sales value	70,320	258.65	1,118.07	2,147,941	178.28	3,469.36
Average annual number of employees	69,288	363.05	1,437.91	2,124,366	287.57	7,841.25

Panel A is summary statistics of the sample and population of industrial land parcels sold during 2007-2010. Panel B is summary statistics of firm-year (2003-2013) observations in our sample of merged firms that purchased lands during 2007-2010 and in the population of all NIE firms. In total, there are 19,602 unique merged firms that purchased lands during 2007-2010 and 711,023 unique NIE firms between 2003-2013. All variables except the last one in Panel 2 are measured in one million yuan.

in big cities. Each blue polygon with a black outline represents one urban unit. Panel C shows the urban units in a small city, Taizhou, where the urban units are mostly disconnected with each other.

There are 21,048 different urban units across the country. The median and mean of the total number of urban units in each prefecture city is 44 and 57, respectively. This large number is because, as Appendix Figure A.1 shows, there are many very small urban units. The median size of the urban units is 0.51 square kilometers and the mean is 8.39 square kilometers.

We match each land parcel to the nearest urban unit. In our estimation of industrial land discount, we use all the residential and industrial land parcels sold through auction during 2007-2019, and we impose additional restriction to the sample size in terms of the number of land sales in each prefecture city. This leaves us with 3,837 different urban units, and the mean and median number of land parcels matched to these 3,837 urban units are 173 and 92, respectively. The mean and median size of these 3,837 urban units are 11.57 square kilometers and 0.52 square kilometers.

B Supplementary Material for Section 3


In this model, we decompose the IRR into government discount rate and the differential market power of the government between the residential and industrial land market. We will also show how the equilibrium industrial discount responds to the change of the city VAT share.

Assume the total land supply is fixed at \bar{L} , which the government can allocate between residential land L_R and industrial land L_I . Denote the value-added tax rate as τ and the city VAT share as k . Assume the production function is $Y = f(L_I)$, then the city's VAT revenue as a function of L_I is $k\tau f(L_I)$. Denote the semi-elasticity of demand for residential (industrial) land as $-\sigma_R$ ($-\sigma_I$). Denote the government discount rate as r^{gov} . The city government chooses the land supply policy to maximize the land sale revenues plus the

Urban Unit - China



Urban Unit - Shanghai



Urban Unit - Taizhou




Figure A.1: Examples of Urban Units.

Note: Panel A is the allocation of all urban units in China. Panel B and C show the urban units in the city of Shanghai and Taizhou. Each blue polygon with black outline represents one urban unit.

present value of its own tax revenues:

$$\max_{L_I, L_R} \frac{1}{r^{gov}} k\tau \cdot f(L_I) + L_I P_I + L_R P_R, \text{ s.t. } L_I + L_R = \bar{L}$$

Replace $L_R = \bar{L} - L_I$, and then the FOC with respect to L_I is:

$$0 = \frac{1}{r^{gov}} k\tau f'(L_I) + P_I + L_I \frac{dP_I}{dL_I} - P_R - L_R \cdot \frac{dP_R}{dL_R} = \frac{1}{r^{gov}} k\tau f'(L_I) + P_I - \sigma_I^{-1} - P_R + \sigma_R^{-1} \quad (23)$$

Decomposition of IRR^{ind} . Equation (23) implies that in equilibrium, the marginal effect of land on tax revenues is:

$$\frac{k}{r^{gov}} \tau f'(L_I) = P_R - P_I - (\sigma_R^{-1} - \sigma_I^{-1}) \quad (24)$$

The IRR can then be written as

$$IRR^{ind} \equiv \frac{\tau f'(L_I)}{P_R - P_I} = \frac{r^{gov}}{k} \left(1 - \frac{\sigma_R^{-1} - \sigma_I^{-1}}{P_R - P_I} \right) \quad (25)$$

Equation (25) decomposes the IRR^{ind} into three components. The first is the government discount rate r^{gov} . The second is the government differential market power in the residential and industrial market, which is captured by the difference of the inverse semi-elasticity scaled by the industrial discount. The last term is k , i.e., the city government share of taxes.

If the government has no monopoly power in both the two land markets, i.e., $\sigma_R = \sigma_I = \infty$, then

$$IRR^{ind} = \frac{r^{gov}}{k}.$$

Effect of Tax share k . Consider how the industrial discount changes when the government share of taxes, k , increases. Denote the price elasticity of demand for residential (industrial) land as $-\epsilon_R$ ($-\epsilon_I$) and assume they are constant. Then we can rewrite Equation (24) as:

$$\frac{k}{r^{gov}} \tau f'(L_I) = P_R - P_I - \left(\frac{P_R}{\epsilon_R} - \frac{P_I}{\epsilon_I} \right) \quad (26)$$

Taking derivatives with respect to k on both sides of Equation (26), we get:

$$\begin{aligned} \frac{1}{r^{gov}} \tau f'(L_I) + \frac{k}{r^{gov}} \tau f''(L_I) \frac{dL_I}{dk} &= \frac{d(P_R - P_I)}{dk} - \frac{dP_R}{dk} \frac{1}{\epsilon_R} + \frac{dP_I}{dk} \frac{1}{\epsilon_I} \\ \frac{d(P_R - P_I)}{dk} &= \frac{1}{r^{gov}} \tau f'(L_I) + \frac{k}{r^{gov}} \tau f''(L_I) \frac{dL_I}{dk} + \frac{dP_R}{dk} \frac{1}{\epsilon_R} - \frac{dP_I}{dk} \frac{1}{\epsilon_I} \end{aligned} \quad (27)$$

Equation (27) states that the effect of tax share on the industrial discount equals the marginal tax revenues of the industrial land, plus the adjustment of the land allocation and the price impact on both the residential and industrial land market.

Consider an example where $f(L_I) = A \times L_I$. Assume the market for industrial land is competitive. The industrial land price would be $P_I = (1 - \tau)A$. Equation (27) simplifies to

$$\frac{d(P_R - P_I)}{dk} = \frac{1}{r^{gov}} \tau f'(L_I) + \frac{dP_R}{dk} \frac{1}{\epsilon_R} \quad (28)$$

On the LHS of Equation (28) is the effect of city VAT share on the industrial land discounts, which we estimate to be 53.55 RMB/m². The first term on the RHS of Equation (28) is the present value of the marginal tax revenues of additional industrial land supply, which we estimate to be 38.24 RMB/m². Using these two numbers we can back out the demand elasticity for residential land, which is:

$$53.55 = 38.24 + 53.55 \times \frac{1}{\epsilon_R} \rightarrow \epsilon_R = 3.5$$

C Supplementary material for section 4

C.1 Estimating λ

We first estimate the “non-standard” compensation to local land occupants (such as “resettlement cost for demolition”). As the “non-standard” nature of this type of cost implies, the data on it is not available at the land parcel level. Therefore, we choose to infer it as a proportional cost from the aggregate data of budget accounts of local

government-managed funds. In particular, we calculate a fraction λ_1 of the land sale must be shared with local land occupants, as the division between the budgeting total expenditure on “Compensation for Using Land and Removing” and the budgeting total revenue on “Sale Receipt of State-owned Land-use Rights”.³⁵ Since we only have data on those numbers between 2010–2014 and we need to use lagged budget revenue to adjust for the time lag between land reserving and land sales, in the end we get $\lambda_1 = 0.28$ using the averages between years 2010–2012, which is in the middle of our data sample.³⁶

For the auxiliary cost associated with providing public services to new residences. We also impose a linear cost structure: if the parcel is sold as residential land, an additional fraction λ_2 of the land must be allocated to build schools to support the residences. We estimate λ_2 by regressing the total area of educational lands on total area of residential lands across different cities, both sold during 2007-2010 and scaled by city population in 2010, after controlling for province fixed effects. The time window 2007-2010 is chosen because we estimate the marginal output of land input based on land sold in 2007-2010. We also conduct the same regression for time interval 2011-2019. To explore potential heterogeneity of λ_2 , we divide the cities into three groups based on the average price of land sold during 2007-2019.

Table A.2 shows estimates of λ_2 . In 2007-2010, for every 100 square meters of residential land, the city government will supply about 8 square meters of land for schools. There is not much heterogeneity across cities with different land price levels. In 2012-2019, the supply of education land seems to have doubled for cities with high and medium price levels, but remains mostly unchanged for the cities with low price levels. We use the λ_2 estimates for 2007-2010 to match our estimates of the taxes.

Our above estimates provide us with the additional cost factor associated with residential land $\lambda = 1 - \frac{1-\lambda_1}{1+\lambda_2} = 1/3$.

³⁵Note that those items do not distinguish between industrial and residential, but they’re generally dominated by residential land, so this is as good an approximate as we can get.

³⁶There is also no data on the time lag between land reserving and land sales, so we chose to take the averages of lagging one to three years. Specifically, we take the total budget compensation between 2010 and 2012, divide it by the total budget revenue between 2011-2013, 2012-2014, and 2013-2015, respectively, and finally take the average of these three ratios. Note that, we see an increasing time trend in λ_1 within our limited sample, unfortunately, we don’t have enough data to track the whole time trajectory of λ_1 .

Table A.2: Lambda Estimates

Price Tier	Sample Period	
	2007-2010	2011-2019
High	0.073*** (4.231)	0.171*** (7.404)
Medium	0.079*** (6.666)	0.146*** (7.231)
Low	0.094*** (6.386)	0.077** (2.798)
Total	0.087*** (10.737)	0.114*** (8.695)

Note: Price tiers are divided based on the 1/3 and 2/3 quantile of the distribution of city-level average land price between 2007-2019. Robust t statistics clustered at province level are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

C.2 Industrial Discount Estimates

To check how the quality control affects the estimates, we exclude X_{it} but keep γ_{ut}^{res} and γ_{ut}^{ind} from Equation (6)-(9) and then construct the sample average version of the industrial land discount using Equation (1). We restrict the sample to city-year for which there are at least 10 auction transactions for both residential land and industrial land in the city-year.

Figure A.2 plots the two estimates. The linear projection of the quality-adjusted estimates on the simple average estimates indicates that controlling for quality control reduces the industrial discount estimates on average.

To visualize the variation of the industrial land discounts across the country, Figure A.3 shows the average industrial land discount across provinces. In general, the industrial discount is higher in more developed areas. It is highest in Beijing and Shanghai across the country, and highest in Chongqing among all the inland provinces.




Figure A.2: Industrial Discount Estimates: Simple Ave. vs Quality Adj.

Note: The figure compares the two city-year level industrial discount estimates: one without controlling for land quality and the other with controlling for land quality.

C.3 Estimating Marginal Tax Revenues

In this section, we explain how the value-added tax system in China works, and explain how we estimate marginal tax rates.

Value-added taxation and accumulated taxes. The main tax collected by the Chinese government from industrial firms is the value-added tax. The VAT rate may differ across firms in different industries, and foreign exports are taxed at a lower rate than domestic sales.

The VAT works as follows. Suppose firm A faces a tax rate of 10%, whereas firm B faces a tax rate of 20%. These tax rates are not marginal tax rates: rather, they imply that the total tax associated with each firm's output must always equal the tax rate multiplied by output. For example, suppose firm A produces 1 unit of output, which is used by firm B as an input, to produce 2 units of output. Total taxes paid by firm A on A's output must equal $10\% \times 1 = 0.1$. Total taxes paid by firm A and B, associated with firm B's output, must equal $20\% \times 2 = 0.4$.

The government implements this scheme recursively. It first calculates the tax payments owed by upstream firms: in the example, A owes $10\% \times 1$, which is 0.1. Firm B is then directly responsible for paying the shortfall between the total tax burden associated with its output, which is 0.4, and the taxes already paid by A, which is 0.1: thus, B is



Figure A.3: Industrial Discount Estimates across the Country.

Note: This figure shows the average quality-adjusted industrial discount estimates across cities during 2007-2019 for each province.

responsible for 0.3 in taxes.³⁷ Logistically, firm B is first responsible for the entirety of its tax burden of 0.4: From an accounting standpoint, this quantity is referred to as B’s “accumulated tax”. B can deduct from its accumulated tax by presenting the government with receipts showing that it has purchased inputs from upstream firms, such as A. For each receipt, the government deducts from B’s tax bill amounts equal to the inputs quantity purchased, multiplied by the tax rate that the producing firm faces. In this example, B would present the government with a receipt showing a purchase of 1 unit from A, and would be entitled to a deduction of $10 \times 1\% = 0.1$.

The accumulated tax that B is responsible for also depends on whether B’s output is supplied domestically or internationally, and also on the export tax rate, which differs across industries. Output which is exported faces a lower tax rate. Hypothetically,

³⁷Note that this implies that the effective marginal tax rate that B pays on value-added can differ from 20%: in this example, B pays 0.3 of taxes, when it has produced only 1 in value-added, so its effective tax rate on value-added is 30%.

suppose that the government's tax rate on exports for B is 5%. If B exports 1 units of output, and sells the other 1 units domestically, B's total accumulated tax is:

$$1 \times 5\% + 1 \times 20\% = 0.25$$

B can still subtract taxes paid by A as deductions from this accumulated tax quantity. However, since all deductions correspond to taxes paid by firm A, the total tax paid by B, including taxes paid by upstream firms on inputs supplied to B will still be the quantity 0.25.

Measurement of marginal tax rates. Conceptually, marginal tax rates are the marginal revenue that the government collects if, for example, firm B increases output by 1 unit. We will attempt to estimate a single representative marginal tax rate associated with all firms in the economy. We will do this simply by regressing accumulated taxes on firms' sale. The regression prediction can be interpreted as follows: if a given firm produces 100 units of output, we infer that the amount of taxes it would pay if it produced 110 units of output is simply the average accumulated tax of all firms currently producing 110 units of output.

To show that this method produces reasonable results, in Figure A.4, we show a scatterplot and a binned scatterplot of firms' accumulated tax against firms' output. The scatterplot shows that the ratio of accumulated tax to output differs nontrivially across firms: some firms pay a smaller share of output as accumulated tax than others. However, the binscatter shows that the relationship between accumulated tax and output across firms is well described by a straight line passing through 0, with slope 12.10%. This means that, on average, firms pay roughly 12.10% of output as taxes, and this does not vary substantially across firms of different sizes.

Besides value-added taxes, firms also pay income taxes and a variety of administrative fees, which we will collectively call $ITF_{j,t}$. Income taxes and fees are charged base on the firm's profit; we will assume these are homogeneous across industries. If we ignore wages and predict the firm's profit with value-added ($S_{j,t} - COGS_{j,t}$), we can write:

$$ITF_{j,t} = (S_{j,t} - COGS_{j,t}) \cdot \psi_t$$

Following a similar logic to our calculations for value-added taxes, the accumulated income taxes and fees associated with firm j 's output, paid by j and its upstream suppliers, is $S_{j,t} \cdot \psi_t$. To account for these taxes and fees, we simply add ψ_t to the marginal tax rate associated with firms' output. Since we do not observe accumulated income taxes and fees in the firm data, we instead estimate the rate ψ_t by regressing income taxes and fees, ITF_{jt} , on firms' value-added, $S_{jt} - COGS_{jt}$. The estimate for the marginal rate is 5.77%, with a tight 95% confidence interval of [5.72%,5.83%].³⁸

Combining these estimates, our final estimate of the effective tax rate facing firms is $(12.10\% + 5.77\%) = 17.87\%$

C.4 Value Added Taxes in Production Networks

In this appendix, we build a simple production-network model to illustrate the assumptions under which the government's incremental tax revenue from land sales can be calculated by multiplying the marginal effect of land sales on output by the value-added tax rate, as we do in expressions (18) and (19).


We consider a finite-layered production network in a single market. The network has M layers, indexed by m ; higher values of m denote more downstream firms. Let \mathcal{J}_m denote the set of firms in layer m of the network. Firm j in layer m produces output $S_{(m,j)}$. Output in the final layer M is sold directly to consumers, whereas output from firms in layer $m < M$ is sold to downstream firms as inputs. Total output of firm j in layer $m < M$ is the sum of its sales to downstream firms in layer $m + 1$:

$$S_{(m,j)} = \sum_{j \in \mathcal{J}_{m+1}} S_{(m,j) \rightarrow (m+1,j)}$$


Cost-of-goods-sold for firm j is the sum of its inputs from upstream firms:

$$COGS_{(m,j)} = \sum_{j \in \mathcal{J}_{m-1}} S_{(m-1,j) \rightarrow (m,j)}$$


³⁸Note that our estimate of income tax as a fraction of value-added is much lower than the official corporate income tax rate, which is 25%. This is because income taxes are applied to firm profits, which are a small fraction of value-added.




(a) Scatter Plot: Accu. VAT vs Sale



(b) Bin Scatter Plot: Accu. VAT vs Sale



(c) Scatter Plot: Income Tax + Fees vs VA.



(d) Bin Scatter Plot: Income Tax + Fees vs VA.

Figure A.4: Marginal VAT Rate and Income Tax and Fees Rate

Note: Panel (a) is the scatter of the "accumulated VAT" vs. sale based on a randomly chosen 1% of the sample and Panel (b) is the bin scatter of the two based on the full sample. Panel (c) is the scatter of corporate income tax plus fees vs. value-added (i.e., output minus input) based on a randomly chosen 1% of the sample and Panel (d) is the bin scatter of the two variables based on the full sample.

If the VAT rate is ψ , value-added taxes for firm j are thus:

$$\psi [S_{(m,j)} - \text{COGS}_{(m,j)}] \quad (29)$$

For firms in layer 1, who do not purchase inputs, $\text{COGS}_{(m,j)} = 0$.

The total tax revenue collected by the government is the sum of (29) across all firms, that is:

$$\psi \sum_{m=1}^M \sum_{j \in \mathcal{J}_m} [S_{(m,j)} - \text{COGS}_{(m,j)}] \quad (30)$$

Now, the sum of $\text{COGS}_{(m,j)}$ for firms in layer m is simply the output of layer $m - 1$. To see this, note that:

$$\begin{aligned} \sum_{j \in \mathcal{J}_m} \text{COGS}_{(m,j)} &= \sum_{j \in \mathcal{J}_m} \sum_{j \in \mathcal{J}_{m-1}} S_{(m-1,j) \rightarrow (m,j)} \\ &= \sum_{j \in \mathcal{J}_{m-1}} \sum_{j \in \mathcal{J}_m} S_{(m-1,j) \rightarrow (m,j)} = \sum_{j \in \mathcal{J}_{m-1}} S_{(m,j)} \end{aligned}$$

Hence, (30) is equal to:

$$\psi \sum_{m=1}^M \left(\sum_{j \in \mathcal{J}_m} S_{(m,j)} - \sum_{j \in \mathcal{J}_{m-1}} S_{(m,j)} \right) \quad (31)$$

Expression (31) is a telescoping sum, which is simply equal to total output in the lowest layer, of final goods. Hence, the government's total tax revenue is simply:

$$\psi \sum_{m=1}^M \sum_{j \in \mathcal{J}_m} [S_{(m,j)} - \text{COGS}_{(m,j)}] = \psi \sum_{j \in \mathcal{J}_M} S_{(M,j)} \quad (32)$$

Expression (32) captures the familiar idea that value-added taxation produces equivalent revenue to taxation of final outputs for the government.

Now, suppose firm j in layer M purchases land, and increases output from $S_{(M,j)}$ to $S'_{(M,j)} > S_{(M,j)}$. Expression (32) implies that, as long as output of other final goods

producers in layer M stays constant, so:

$$S'_{(M,\tilde{j})} = S_{(M,\tilde{j})}, \forall \tilde{j} \in \mathcal{J}_M, \tilde{i} \neq j$$

Then the incremental tax revenue collected by the government is simply:

$$\psi \sum_{j \in \mathcal{J}_M} S'_{(M,j)} - \psi \sum_{j \in \mathcal{J}_M} S_{(M,j)} = \psi \left(S'_{(M,j)} - S_{(M,j)} \right) \quad (33)$$

that is, the increase in output of firm j . Note that while we require that output of final-goods firms to be held fixed, (33) holds regardless of what happens to the output of upstream firms. For example, firm j may increase purchases from input suppliers in layer $M - 1$ to produce more output, who then increase purchases from suppliers in $M - 2$, and so on. Conversely, j may increase efficiency, allowing it to produce more output using less inputs. (33) holds in either case. Intuitively, (33) implies that the government's total tax revenue can always be calculated as if final goods are taxed directly, so changes in upper layers of the production network can are not needed for calculating incremental tax revenues.

This discussion also highlights a few settings in which our assumptions are violated. First, our calculations for marginal tax revenue require that output of other final-goods producers firms in layer M are fixed. In practice, firm j could produce output which is a substitute or complement for other final-goods producers. If firm j 's sales cannibalize other final good producers' sales, then our approach will overestimate incremental tax revenue. If firm j 's output is complementary to other final goods producers, so higher output from j increases other firms' output, then our approach will underestimate incremental tax revenue.

Second, our approach only holds for final goods producers. For intermediate goods producers, expansion in output will tend to expand output of downstream firms; our approach applied to intermediate producers will thus tend to underestimate the marginal effect of land sales on taxes.

Third, the analysis assumes that all firms in the production network operate within the same market, and thus are taxed by the same local government. If a firm purchases inputs from a firm in a different market, subject to taxation by a different government,

our approach will tend to overestimate tax revenue from land sales, because some of the revenue in the production chain does not accrue to the local government selling the land.

C.5 The Impact of Land Purchases on Total Output in a Domar Aggregation Model

The foundational theorem of [Hulten \(1978\)](#) states that in a competitive market with a representative consumer, the impact on aggregate TFP of a microeconomic TFP shock is equal to the Domar weight, i.e., the shocked producer's sales as a share of GDP. Hulten's theorem is significant in the sense that sales summarize the macroeconomic impact of microeconomic shocks and we do not need to concern ourselves with the details of the underlying production network structures. If we think of the land-purchase as a shock to the producer's TFP, using the same framework as [Baqae and Farhi \(2019\)](#), we can show that when a firm purchases additional land, the impact on total output in the economy is smaller than the effect on the sale of the land-purchasing firm.

Following [Baqae and Farhi \(2019\)](#), the representative consumer maximizes a constant-returns aggregator of final demand for individual goods

$$Y = \max_{\{c_1, \dots, c_N\}} \mathcal{D}(c_1, \dots, c_N),$$

subject to the budget constraint

$$\sum_{i=1}^N p_i c_i = \sum_{f=1}^F w_f \bar{\ell}_f + \sum_{i=1}^N \pi_i,$$

where c_i is the representative consumer's consumption of good i , p_i is the price and π_i is the profit of producer i , w_f is the wage of factor f which is in fixed supply $\bar{\ell}_f$.

Producer i maximizes its profit

$$\pi_i = p_i y_i - \sum_{f=1}^F w_f \ell_{if} - \sum_{j=1}^N p_j x_{ij},$$

subject to the following production technology

$$y_i = A_i F_i(\ell_{i1}, \dots, \ell_{iF}, x_{i1}, \dots, x_{iN}),$$

where A_i is technology, x_{ij} are intermediate inputs of good j used in the production of good i , and ℓ_{if} is factor f used by i .

The market clearing conditions are

$$y_i = \sum_{j=1}^N x_{ji} + c_i \text{ and } \bar{\ell}_f = \sum_{i=1}^N \ell_{if}$$

Define the equilibrium output Y as a function of the technology to be $Y(A_1, \dots, A_N)$. Since the first welfare theorem holds, the equilibrium allocation solves

$$Y(A_1, \dots, A_N) = \max_{c_i, x_{ij}, \ell_{if}} \mathcal{D}(c_1, \dots, c_n) + \sum_i \mu_i (A_i F_i((\ell_{if})_f, (x_{ij})_j) - \sum_j x_{ji} - c_i) + \sum_f \lambda_f (\bar{\ell}_f - \sum_i \ell_{if})$$

where μ_i and λ_f are Lagrange multipliers.

The envelope theorem implies that the increase in total output satisfies:

$$\frac{dY}{dA_i} = \mu_i F_i$$

Let the price of Y to be p_c . We can then show that $\mu_i = \frac{p_i}{p_c}$. For each product j , it is either used as an input for another product i or consumed by the households. If it is used as an input for i , then the profit-maximization of firm i and the optimization of the social planner imply:

$$p_i \frac{A_i \partial F_i}{\partial x_{ij}} = p_j \text{ and } \mu_i \frac{A_i \partial F_i}{\partial x_{ij}} = \mu_j$$

If it is consumed by the households, then

$$p_c \frac{\partial \mathcal{D}}{\partial c_j} = p_j \text{ and } \frac{\partial \mathcal{D}}{\partial c_j} = \mu_j$$

Therefore, $\mu_i = \frac{p_i}{p_c}$ for any i , and hence we have that the increase in total output induced by the productivity change is:

$$p_c \frac{dY}{dA_i} = p_i F_i \quad (34)$$

The discussion so far follows exactly [Baqae and Farhi \(2019\)](#). Building on this framework, we want to compare the total output increase with the sale increase of the firm when the firm's productivity improves. Suppose in each sector of i , there are infinite number of firms indexed by k , and each has its own productivity A_i^k . Following the same steps as above, the impact on total output is

$$p_c \frac{dY}{dA_i^k} = p_i F_i^k$$

Now consider the profit-maximization problem of this firm (to simplify the notation we will drop k):

$$\max_{l_{if}, x_{ij}} p_i A_i F_i(l_{i1}, \dots, l_{iF}, x_{i1}, \dots, x_{iN}) - \sum_{f=1}^F w_f l_{if} - \sum_{j=1}^N p_j x_{ij}$$

The profit maximization conditions are

$$p_i A_i \frac{\partial F_i}{\partial l_{if}} = w_f \text{ and } p_i A_i \frac{\partial F_i}{\partial x_{ij}} = p_j$$

The effect on the firm's sale, $p_i y_i$, is

$$p_i \frac{dy_i}{dA_i} = p_i F_i + p_i A_i \left(\sum_f \frac{\partial F_i}{\partial l_{if}} \frac{\partial l_{if}}{\partial A_i} + \sum_j \frac{\partial F_i}{\partial x_{ij}} \frac{\partial x_{ij}}{\partial A_i} \right) = p_i F_i + \left(\sum_f w_f \frac{\partial l_{if}}{\partial A_i} + \sum_j p_j \frac{\partial x_{ij}}{\partial A_i} \right) \quad (35)$$

Thus, the increase in firms' sales, (35), will exceed the increase in total output, (34), as long as the difference term is positive:

$$\left(\sum_f w_f \frac{\partial l_{if}}{\partial A_i} + \sum_j p_j \frac{\partial x_{ij}}{\partial A_i} \right) > 0 \quad (36)$$

The LHS of expression (36) involves the derivatives $\frac{\partial \ell_{if}}{\partial A_i}$ and $\frac{\partial x_{ij}}{\partial A_i}$, which are the changes in inputs induced by the increase in productivity. These will generally be positive: more productive firms will expand inputs. Thus, the increase in sales of the affected firm will be larger than the increase in total output.³⁹

The intuition for this result is as follows. When a firm's productivity increases, there is a direct effect on sales from higher productivity, and an indirect reallocation effect as the firm changes its purchases of inputs, in response to increased productivity. When the first welfare theorem holds, the reallocation effects do not have a first-order effect on total output, since inputs are equally productive in all industries on the margin. Hence, the sales increase of the affected firm overestimates the increase in total output, whenever the affected firm tends to increase inputs in response to increased productivity.

C.6 Panel Imbalance

In this subsection we further study the causes and consequences of panel imbalance in our difference-in-differences design. As noted in Section 4.2.2, firms enter and exit our panel due to data linkage issues, firm births and deaths, and sales falling below a threshold for inclusion in our data. While panel imbalance arising from data linkage issues is likely to be idiosyncratic, there is a concern that left-censoring due to sales falling below the threshold for inclusion could affect our estimate of the effect of land purchase. To assess the importance of censoring, we test whether panel imbalance (firm attrition) is more likely for firms close to the censoring boundary: to the extent panel imbalance is idiosyncratic, we should not see differences in the distribution of sales for firms that do and do not attrite.

Figure A.5 shows the results of this test in the form of kernel densities of past-year sales, separately for firms that do and do not attrite in a given year. We see the two distributions are strikingly similar. If anything, firms near the 2011 censoring boundary (denoted by the second of the two vertical dashed lines in the figure) are disproportionately likely to be *not* censored. We are reassured that the role of censoring is likely modest in generating panel imbalance.

³⁹Note that it is possible for the LHS of (36) to be negative; for example, if demand for firm output is sufficiently inelastic, increasing productivity will cause the firm to tend to scale down input purchases.




Figure A.5: Distribution of Log(Sale) in the Past Year Conditional on Exiting or Not

Note: This figure reports the kernel densities of the past year log(sale) for firms that do and do not exit in a given year separately. For 2011, the past year is 2009 as we do not have data for 2010. The two vertical dashed lines represent the censoring boundaries, which is 5 million RMB before 2011 and 20 million RMB after 2011.

We also examine whether panel imbalance varies by treatment status (land purchase). Appendix Table A.3 shows the survival rates of the treated and matched control firms for each event year, i.e., the percentage of firms remaining in the sample. By construction, all the firms are observed in the two years before treatment. There is not much difference between the treated and control firms in $t = \tau - 3$ and $t = \tau - 4$ in terms of the survival rates, confirming that the matching generates a comparable control group for the treatment group. However, after the treatment year, the survival rate of the treatment group is higher than the control group. This is consistent with the firm's expansion on the newly acquired land increasing sales and making the firm more likely to stay above the censoring threshold. While our evidence in Figure A.5 suggests the consequences of such censoring is likely to be modest, this does imply that our estimate of the treatment effect of land purchase is conservative: when the control firms exit the sample, dropping this observation removes a higher difference between the treated and control firms in

Table A.3: Survival Rates of the Matched Sample

Event Year	2007		2008		2009		2010	
Treat	0	1	0	1	0	1	0	1
t=-4	37%	39%	59%	57%	64%	57%	56%	49%
t=-3	81%	78%	78%	73%	79%	75%	68%	64%
t=-2	100%	100%	100%	100%	100%	100%	100%	100%
t=-1	100%	100%	100%	100%	100%	100%	100%	100%
t=0	87%	100%	75%	100%	84%	100%		
t=1	68%	87%	59%	78%			63%	100%
t=2	55%	71%			50%	71%	59%	95%
t=3			36%	52%	46%	68%	54%	89%
t=4	32%	48%	32%	46%	40%	63%		
t=5	29%	44%	29%	42%				
t=6	26%	41%						

Note: This table reports the survival rates, i.e., the percentage of firms remaining in the sample in each year, for the treated and matched control firms for each event year. Note the rates are 100% for the two years with data before the event year by construction.

sales, so dropping out these pairs will make the treatment effect estimates downward biased. This ultimately translates into a corresponding downward bias in our estimates of the effects of land sales on tax revenues, and hence a downward bias in our estimate of local governments' IRR from land sales.

C.7 Classification of Targeted Industries

Table A.4 shows the list of industries that were ever targeted by one of both of the Five-year Plans initiated in 2006 and 2011.

Table A.4: Targeted Industries of Five-Year Plan 2006 & 2011

Targeted Industries
Mfg. of agricultural and non-staple foodstuff
Chemical feedstock and chemical mfg.
Medicine mfg.
Non-ferrous smelting and extrusion
Specialized facility mfg.
Transport and comms facilities mfg.
Automobile mfg.
Electric machinery and equip mfg.
Mfg. of comms equip, computers and other electronic equip
Production and supply of electric power and heat power
Gas generation and supply
General-purpose equip mfg.
Exploitation of petroleum and natural gas
Chemical fiber mfg.
Coal mining and washing
Ferrous metal smelting and extrusion

Note: This table lists the industries that were ever targeted by one or both of the two Five-year Plans initiated in 2006 and 2011, which cover the period 2006-2015.

D Supplementary material for section 5

D.1 Guangdong Policy Changes in 2016

The main policy change that most provinces made in 2016 was to change the share of VAT taxes accruing to city governments. However, Guangdong is an outlier: it did not change the city government VAT tax share, but implemented a number of other policies to encourage city governments to allocate more industrial land. These policies thus confound our analysis of the effect of VAT tax changes on industrial land sales. When we include cities in Guangdong when estimating Equation (21), there is no significant results from dynamic treatment effect analysis (the results are available upon request).

Guangdong made the following policy changes in 2016. All cities within the province were required to specify a region within which all land had to be sold as industrial, not residential land. Cities were also broadly required to guarantee “sufficient” industrial land supply to advanced manufacturing industries. Incentives to do so included, for example, policies stating that industrial land allocated to major investment projects would not count towards land quotas, that is, the maximal amount of land that cities could sell within a certain period of time.

These policies were imposed upon city governments and supervised by the provincial government. Cities which experienced higher growth in manufacturing were to be rewarded with larger quotas for future land sales. To verify in the data that this policy encouraged more industrial land sales, we regress an indicator for whether a city received a reward of higher land quotas in 2019, on the share of land sold as industrial in the year 2017. The results are shown in Table A.5: as predicted, cities allocating more industrial land were substantially more likely to be rewarded.

The majority of these policies were applied only in the Guangdong province.⁴⁰ These policies are likely to have contributed to Guangdong increasing industrial land supply, despite the fact that the city government VAT share in Guangdong stayed constant in 2016. Thus, we drop Guangdong from our event study analysis.

Table A.5: Industrial Land Supply and Reward in Guangdong

Dep Var: Reward	(1)	(2)	(3)
Share of industrial land supply	0.380*	1.963*	1.215*
	(1.689)	(1.843)	(1.931)
Observations	121	118	118
R ²	0.138	0.0776	0.0788
Spec	OLS	Logit	Probit
City FE	Yes	Yes	Yes

Note: This table reports the correlation between the share of industrial land supply in 2017 with whether the district/county received reward in 2019 across the 118 districts/counties in Guangdong. Robust t-statistics in parentheses. *** p<0.01, ** p<0.05, * p<0.1

⁴⁰Two exceptions are that the policy of rewarding cities with higher manufacturing growth with larger land quotas was also implemented in Guangxi in 2017, and Sichuan in 2019.