

The Black Market for Beijing License Plates*

Øystein Daljord[†] Mandy Hu[‡] Guillaume Pouliot[§] Junji Xiao[¶]

December 2, 2019

Abstract

Black markets can reduce the effects of distortionary regulations by reallocating scarce resources towards those who value them the most. The illegal nature of black markets however creates transaction costs that reduce the gains from trade. We take a partial identification approach to infer gains from trade and transaction costs in the black market for Beijing car license plates which emerged following their recent rationing. We use optimal transport methods to non-parametrically estimate a lower bound on the volume of unobserved black market trade under weak assumptions using comprehensive car sales data. We find that at least 11% of the quota of license plates is illegally traded. We next infer gains from trade and transaction costs and tighten the bounds on the volume of trade under further assumptions on black market transactions. The inferred size of the transaction costs suggests severe market frictions: between 61% and 82% of the realized gains from trade are lost to transaction costs, while between 7% and 28% of the potential gains from trade are realized in the black market.

*A warm thanks to Shanjun Li for comments and data. Thanks to John Barrios, Jean-Pierre Dubé, Alfred Galichon, Chris Hansen, Carl Mela, Luxi Shen, Avner Strulov-Shlain, Baohong Sun, Chad Syverson, and Thomas Wollmann for comments. Thanks to Charles Ahlstrom and Xinyao Kong for research assistance.

[†]Chicago Booth School of Business, University of Chicago, 5807 South Woodlawn Avenue, Chicago, IL 60637, USA. E-mail: Oystein.Daljord@chicagobooth.edu. Web: faculty.chicagobooth.edu/oystein.daljord.

[‡]CUHK, Department of Marketing, Room 1105, 11/F, Cheng Yu Tung Building, 12 Chak Cheung Street Shatin, N.T., Hong Kong. E-mail: mandyhu@baf.cuhk.edu.hk. Web: bschool.cuhk.edu.hk/staff/hu-mandy-mantian/.

[§]Harris School of Public Policy, University of Chicago, 1155 E 60th St, Chicago, IL 60637 USA E-mail: guillaumepouliot@uchicago.edu. Web: sites.google.com/site/guillaumeallairepouliot.

[¶]Department of Economics, UTS Business School, University of Technology Sydney 14-28 Ultimo Rd., Ultimo, NSW2007, Australia. E-mail: junji.xiao@uts.edu.au. Web: uts.edu.au/staff/junji.xiao.

Keywords: informal economy/underground economy, optimal transport, partial identification, semiparametric and nonparametric methods

JEL codes: E26, D450, P230, C140

1 Introduction

The informal economy, the trade in goods and services that goes undetected in official statistics, makes up an estimated one-sixth of the GDP in the world economy (Schneider et al. (2010)). A subset of the informal economy is black markets: markets where goods and services are traded illegally. Black markets emerge in response to restrictions on trade, such as prohibition, taxation, and rationing, that create gains from trade. Though black markets are usually considered undesirable *per se*, they may improve welfare by reducing the impact of regulatory distortions (Davidson et al. (2007)). Yet, due to their illegal nature, little is known about how black markets perform. The fact that black markets are illegal creates transaction costs, such as potential legal liabilities and search costs, that reduce the gains from trade. In this paper, we ask how much of the gains from trade are lost to transaction costs in a black market, which is informative about unintended effects of trade restrictions and the efficacy of enforcement.

We address this question in the context of the black market for license plates in Beijing. In 2011, the Beijing government restricted driving within the city limits to all but those who have Beijing license plates in order to regulate increasing pollution and congestion. The license plates were rationed and the quota was allocated by lottery. While Shanghai has used auctions to allocate license plates since the 1990s, the auction format has found limited public support. Li (2018) reports survey evidence which shows that even though a majority of Beijing residents recognized the need for rationing, less than 10% preferred an auction, and about 40% preferred a hybrid lottery and auction mechanism.¹ At the time the Beijing lottery was introduced, there were growing political concerns over the increasing auction prices, which by 2010 had reached the level of the price of a typical new car. Having considered an auction format politically unviable, Beijing chose a lottery out of fairness concerns (Huang and Wen (2019)).

Allocating license plates by lottery creates gains from trade: some lottery winners may prefer to trade their license plates rather than use the plates themselves. In Section 2, we report anecdotal evidence of a black market for license plates that emerged soon after the introduction of the lottery. News reports note that the market largely took

¹See its footnote 13.

place online, while car dealerships seem to have acted as occasional intermediaries. The market was supplied by both corrupt officials and lottery winners. Enforcement of the non-transferability was seemingly lax. Both rentals, without formal ownership, and purchases, with or without formal ownership, were offered in the market. The reported transaction prices are of a similar magnitude to the Shanghai auction prices.

We find evidence of black market trade in comprehensive car sales data. We argue that if the license plates were allocated to a random selection of the population of car buyers, and there was no black market trade, then the sales distribution would not shift materially immediately following the lottery. If instead black market trade allocated license plates towards wealthier households, whom [Li and Xiao et al. \(2017\)](#) shows empirically purchase more expensive models, the average prices would shift upwards. We use a standard difference-in-differences approach to document such an upward shift in the average Beijing car prices following the rationing. We find no such price shifts in comparable nearby cities that did not ration license plates.

We do find evidence of a similar shift in the average prices in Tianjin after it introduced a hybrid auction/lottery rationing mechanism, which presumably allocates the share of the license plates that are auctioned towards wealthier households that buy more expensive cars. The price jump in Tianjin suggests that a hybrid lottery/black market mechanism may similarly have car buyers from a selected subset of wealthier households. We do not find that the shift was caused by car dealerships adapting their pricing to the rationing. Instead, we find evidence that vertical price restraints precluded supply side pricing responses to the market contraction.

To assess the gains from trade that are realized in the black market, we need to at least know the volume of trade. Though the difference-in-differences estimates are consistent with the existence of black market trade, they carry little information about the volume of trade. Our main empirical challenge is that the black market transactions are unobserved. We develop a novel, transparent, and intuitive empirical approach based on optimal transport methods to identify a lower bound for the gains from trade. The empirical strategy is to exploit the fact that a lottery which randomly samples car buyers from the population will generate a different sales distribution than a black market that

allocates licenses to a selected subset of the population. We show how a distance between the two sales distributions is informative about the volume of black market trade. We develop an analogue to difference-in-differences for distributions which uses shifts in the sales distribution in the nearby city of Tianjin, where there was no rationing at the time, to control for common trends. We find that at least 11% of the quota is traded on the black market.

In the final step, we combine information from a variety of sources, including anecdotal evidence (news reports), to infer transaction costs in a market equilibrium model. We use [Li's](#) recent estimates of the willingness-to-pay for Beijing license plates combined with anecdotal evidence to further bound the volume of trade. We consider transaction costs to be a tax that restricts the volume of trade. Transaction costs are inferred as the wedge between demand and supply that is necessary to rationalize the estimated volume of trade.

Though the anecdotal evidence points to lenient enforcement of the non-transferability of the license plates, our estimates suggest otherwise. Firstly, we find that sizable gains from trade are left unrealized. While the market would realize RMB 18.8 billion in gains from trade in the absence of transaction costs, we find that the realized net gains from trade lie in a plausible range from RMB 1.3 billion to RMB 8.2 billion.² Secondly, we find that a plausible range from 61% to 82% of the realized gross gains from trade in license plates are lost to transaction costs.

Our paper ties into different literatures. One is an older, and mostly theoretical, literature on the economic impact of rationing on welfare and incentives for illegal trade, e.g. [Tobin \(1952\)](#), [Drèze \(1975\)](#), [Stahl and Alexeev \(1985\)](#), [Dye and Antle \(1986\)](#). Our results also complement a literature within the field of public finance that estimates the size of the informal economy. Various indirect measures have been proposed, from monitoring excessive electricity consumption to currency velocity, see e.g. [Schneider et al. \(2010\)](#). The interest in this literature however tends to be in the scale of tax evasion in the macro economy. Our interest is instead in measuring the performance of a particular black market, which calls for a more tailored empirical approach.

²Dividing by six gives a rough estimate of the dollar equivalents.

Our novel non-parametric estimators contribute to a recent literature on optimal transport. Though optimal transport methods have been applied in the econometric theory literature, see e.g. Galichon (2016) for an overview, these methods have rarely been used in the applied literature outside of matching models (Chiappori and Salanie (2016)). We believe our application shows that optimal transport methods can complement standard event study approaches. In particular, we demonstrate that optimal transport methods can recover information that standard event study approaches, e.g. difference-in-differences, cannot.

Finally, the market we study is of interest in itself. Our paper complements a small literature on welfare effects of the recent rationing of car sales in larger Chinese cities, e.g. Li, Xiao et al. (2017), Tan et al. (2019), and Huang and Wen. Importantly, this literature has assumed away the existence of a black market. Our results show that black market trade plausibly ranges from 11% to 37%, which affects welfare calculations.

We give an overview of the lottery rationing mechanism and the black market in Section 2. We describe the data in Section 3 and report a standard event study in Section 4. The optimal transport methods are laid out in Section 5 along with the estimation results. We develop and evaluate the market equilibrium model in Section 6, including the estimated gains from trade and transaction costs. We conclude with a brief discussion in Section 7.

2 The lottery and the black market for license plates

In late December 2010, the Beijing government announced that, effective January 2011, car license plates would be required to drive freely within the city limits.³ The license plates were rationed. A quota of non-transferable license plates, which was set to about 40% of the previous year's sales, was allocated by lottery. The lottery application process was simple: the pecuniary costs were low and the application could be completed quickly online. Applicants had a uniform probability of winning a license plate and

³See Section B for the exact restrictions and eligibility criteria for the lottery.

there were limits on the number of applications each household can submit. Immediately following the introduction of the lottery, the number of newly registered cars in Beijing was reduced to the level of the quota. The quotas have since been tightened further.

Allocating the license plates by lottery creates gains from trade: rationed prospective car buyers may find lottery winners, or corrupt officials, that are willing to sell. Despite the license plates being non-transferable by law, there are numerous news reports of a black market for license plates that emerged soon after the introduction of the lottery. Table 1 gives an overview of results from a search in Chinese news reports that include mentions of black market transaction prices. The most frequent mentions of yearly rental prices for a license plate is between RMB 6 000 and RMB 12 000 annually. To provide context, the average monthly wage in Beijing in 2011 was around RMB 5 000. Given an interest rate of about 6% at the time, the twenty year net present value (NPV) of the rental income stream of a license plate ranges from RMB 73 000 to RMB 145 000.

Chinese news outlets report of license plates purchase prices ranging from RMB 16 000 to RMB 650 000. The latter price was for a so-called Jing-A plate, which carries particular prestige, and is predominantly used by party officials and certain institutions. These seem to be rare transactions of exceptionally high value, see [New York Times](#). The majority of reported purchase prices, without ownership, are between RMB 70 000 to RMB 120 000. The upper end of that range is on par with the best selling car at the time (Ford Focus), see [South China Morning Post](#).

[Reuters](#) and [New York Times](#) report of a number of online sites that match license owners with buyers. Car dealerships seem to act as intermediaries. In 2014, [Reuters](#) quotes a sales representative, with full name at an identifiable dealership, who claims that he can provide prospective buyers with license plates.⁴ Such reports suggest that the non-transferability of license plates is leniently enforced.

The market is reportedly supplied by both lottery winners and corrupt officials. According to [South China Morning Post](#), officials involved with drafting the lottery rules

⁴“Wang Shaoyong, sales manager at a Peugeot dealer in Beijing, said his shop provides car buyers with license permits from a partner firm that has many car plates registered in its name.”

Table 1: Chinese news reports that mention black market transaction prices in 2011 and 2012

Article Date	Newspaper (EN)	Rental Price (RMB)	Purchase Price (RMB)
July 23, 2011	China Times	0.5k-2k per mo.	20k(N); 40k-90k(N)
Sep 18, 2011	Beijing Times	1st yr free, then 1k per mo	
Oct 8, 2011	China National Radio	0.5k-1k per mo	
Oct 11, 2011	China Economic Weekly	1000 per mo.	80,000-100,000(U)
Dec 19, 2011	The Beijing News		85k(U); 130k(U); 75k(U)
Feb 18, 2012	The Rule of Law Weekend		650k(Y), Jing A. 260k(Y). 50k(Y).
Aug 29, 2012	Beijing Youth Daily	0.5k-1k per mo.	150k(Y); 30k-40k(N)
Sep 10, 2012	Beijing Times	800 per mo.	40,000 (N)
Dec 15, 2012	Workers Daily	500 per mo.	200,000+ (Y), 170k, 16k (N), 30k-50k (N)
Dec 19, 2012	Guangming Daily		30k, (U)
Dec 26, 2012	ChinaNews.com		80k,(U) 160k, (U) 200k, (U)

The letters in the bracket in the column of purchase prices indicate the ownership status.

Abbreviations: Y = with ownership, N = without ownership, U = ownership unclear. Jing A is a license plate that is used predominantly by governments and institutions. These license plates command a premium in the market.

have experienced unusual luck in the lottery and there have been accusations of the lottery being rigged.⁵ Some reports however suggest strong enforcement. The [New York Times](#) reported that the head of the Beijing department of transportation was sentenced to life in prison in 2015 for selling Jing-A license plates. This sentence however appears to be part of general secretary Xi Jinping’s campaign against corruption amongst party officials. The evidence of the level of enforcement is mixed, leaning towards lenient.

In sum, we find anecdotal evidence in support of the existence of a material market for license plates, some indications of the level of enforcement, and a plausible range of transaction prices. The news reports however give little indication pertaining to either the volume of trade in license plates or the transaction costs. We next turn to car sales data in order to infer the size of the market.

⁵“The Beijing News reported on Thursday that a record 1.26 million residents competed for fewer than 20,000 plates this month. “Liu Xuemei” was so lucky that this person - or perhaps people - won two plates in May. Cynical car-plate hunters wondered if they should change their name to stand a better chance in the next contest. Eagle-eyed internet users soon discovered Liu Xuemei was the name of the director of the vehicle and driver management department of the Ministry of Public Security. Liu Xuemei, in her 30s, is in charge of drafting rules for vehicle permits.”

3 Data

We obtained vehicle registration data from a Chinese marketing research company, the Webinsight Technology and Information Corporation⁶. We also collected public information on vehicle quotas and city characteristics from other channels specified below. The vehicle registration data have observations at the city-month level on aggregate registration of all vehicle models available from January 2010 to December 2015 in 35 China cities. The vehicle models are identified by their unique codes as catalogued by the Motor Vehicles’ Type and Model Designation, published by the National Standard of People’s Republic of China, and various car characteristics are provided. Table 3 presents summary statistics.

The vehicle quota data is collected from each city’s official publications.⁷ The data include the number of applicants, the quota, the average value of winning bids (in the auction rationing mechanism), and the share of each type of rationing mechanism if multiple mechanisms are applicable, see Table 2. Table 3 aggregates the variables across cities and years.

Table 2: Rationing mechanisms

City	2015 Per Capita GDP (USD)	Rationing Mechanism	Announcement Date	Implementation Date	Average Quota per Month	Quota Allocation (electric:lottery:auction)	Restriction on vehicles without local plate
Beijing	18731	Lottery	12/23/2010	1/1/2011	10,000	0:1:0	Yes
Tianjin	13355	Auction and Lottery	12/15/2013	12/16/2013	8,000	1:5:4	Yes
Shijiazhuang	4576	No rationing	n/a	n/a	n/a	n/a	n/a

Table 3: Summary statistics

Variable	count	mean	std dev	min	max
MSRP (RMB)	5 065 520	150 356	100 919	20 800	1 305 000

⁶<http://www.webinsight.cn>

⁷The official publications of information on allocation mechanisms are from Shanghai International Commodity Auction Co. Ltd., Beijing Passenger Vehicle Quota Administration Office, Tianjin Information System of Passenger Vehicle Quota Administration

4 Difference-in-differences analysis

Figure 2 displays average car prices in Beijing, Tianjin, and Shijiazhuang before and after the introduction of the Beijing lottery. Figure 2 shows that the average price of a Beijing car jumped 23% following the introduction of the lottery. To control for unrelated economic trends, we compare the Beijing prices to the prices in the nearby cities of Tianjin and Shijiazhuang. Tianjin is the fourth largest city in China and is about 130km away from Beijing. It experienced the fastest economic growth of the major Chinese cities in the period covered by the data, with a GDP of about $\frac{2}{3}$'s that of Beijing. Shijiazhuang is the capital and the largest city of North China's Hebei Province and is about 270 km away from Beijing. It has about 23% of the GDP of Beijing, but a similar population size to Beijing. Shijiazhuang had not introduced rationing by 2016, while Tianjin introduced a hybrid auction/lottery format in January 2014. The average price in Tianjin and Shijiazhuang increased by about 7% in the same period. Assuming common trends in Beijing, Tianjin, and Shijiazhuang, about 17% of the Beijing price jump remains, see Section A for the corresponding difference-in-differences regressions.

We first rule out that the price jump follows from a supply side price response to the rationing. When we compare the prices of cars that were sold in 2010 with the prices of the same car models in 2011, we find virtually no changes in prices. Figure 3 shows the distribution of the change in MSRP for car models that were registered in both 2010 and 2011. The lack of supply side price responses is explained by the vertical restraints that the industry used at the time. Li reports that car manufacturers used resale price maintenance at the national level at the time, which precludes car dealerships from adapting their pricing to changes in local demand. Li also notes that discounts on the MSRP of luxury goods more generally are rare in China, which suggests that measurement errors in prices (discrepancies between purchase prices and Manufacturer Suggested Retail Prices) is a minor issue.

Though car dealerships were not allowed to change their prices, Feenstra

(1988) suggests that car dealerships may upgrade the quality of their product lines in response to the quota. The gain from an increase in the intensive margin can partially offset the loss from a loss of extensive margin following the rationing. The increase in the average price could therefore be the consequence of manufacturers' strategic responses to the rationing. If so, the average prices for the models which are sold following the rationing, but not before, would be higher than the models that were sold both before and after the introduction of the lottery. Figure 1 shows barely noticeable changes in the distribution of the prices of the models offered before the lottery (pre) and the models that were only offered after the lottery (post). The post-lottery distribution seems shift slightly to the left, if anything. In sum, we find little evidence that the price jump was caused by supply side responses.

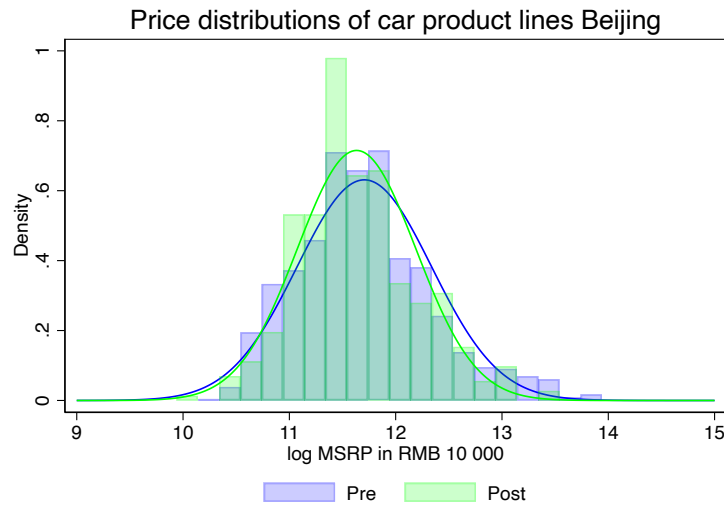


Figure 1: Changes in the distributions of prices of the product lines offered before and after the lottery (pre) and only after the lottery (post).

The price jump may instead be explained by black market trade. If the Beijing lottery randomly selected car buyers from the population, if there was no trade in license plates, and if income and preferences did not change before and after the lottery, then we would expect the sales distribution before the lottery to be the same as the sales distribution after the lottery. The observed shift in the sales distribution is instead consistent with a black market that reallocates the license

plates to wealthier households that buy different cars.

Similar price jumps are observed in other Chinese cities that use market based allocation mechanisms. One example is Tianjin, which introduced a hybrid allocation mechanism in January 2014. The mechanism allocates 50% of the license plates by lottery, 40% by auction, and reserves the remainder for electric cars. Figure 2 shows a similar, if smaller, price jump soon after the introduction of the hybrid mechanism in Tianjin. The auction presumably allocates its share of the license plates towards wealthier households who buy more expensive car models. The jump in the Beijing prices may be explained by a combination of a lottery, which randomly selects car buyers from the population, and a black market, which selects wealthier households, resulting in a price jump similar to that observed following the introduction of the lottery-auction hybrid mechanism in Tianjin.

It is however not obvious that a black market would shift sales towards more expensive cars. We can think of the price of a license plate as a tax on car purchases that substitutes lottery winners towards less expensive cars. Yet, similar shifts towards more expensive cars are documented empirically by [Tan et al. \(2019\)](#) and [Xiao et al. \(2017\)](#) for other Chinese cities that later introduced hybrid allocation mechanisms. Positive correlations between auction prices, that similarly tax car purchases, and car expenditures are also observed in Shanghai, which has used auctions to allocate license plates since 1996.⁸

We conclude that the shift in the average prices in Beijing is consistent with a black market that reallocates license plates towards wealthier households. Though a 17% increase in the average price may seem substantial, it is however hard to infer the volume of trade, which is one of our quantities of interest, from the first moments of the sales distributions. In the next section, we show that shifts in the sales distributions are informative about the volume of trade in a way that shifts in

⁸ The average Shanghai auction price in 2010 was RMB 39 000 and increased to RMB 81 000 in 2015, while the average car price increased from RMB 173 000 to RMB 215 000 over the same period.

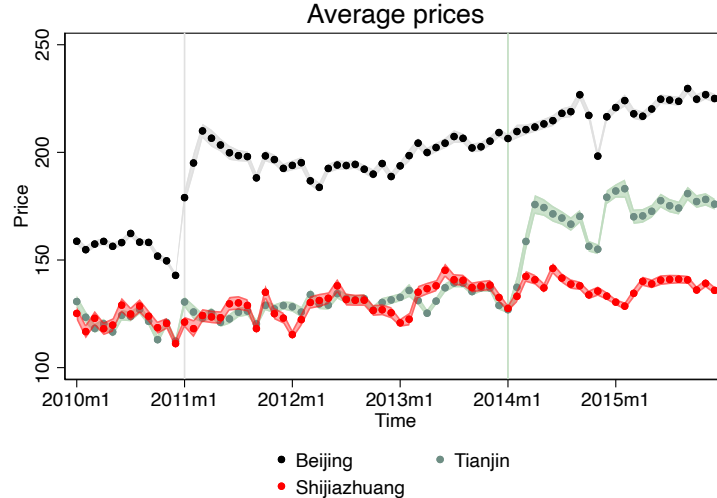


Figure 2: Tianjin introduced a hybrid lottery and auction mechanism January 2014. Confidence bands at 0.1 percent level.

the first moments are not.

5 Inferring bounds on the volume of trade using optimal transport

Our main empirical challenge is that we do not observe transactions in the black market directly. We instead infer the volume of trade from observed changes in car sales. Figure 4 shows overlaid, smoothed empirical distributions of car prices in Beijing in 2010 and 2011-2012. The sales distribution clearly shifts to the right after the introduction of the lottery. This shift is consistent with a black market that reallocates license plates to households that buy more expensive cars, but is inconsistent with a lottery that randomly selects car buyers from the population. Our strategy is to infer a lower bound for the volume of black market trade by quantifying the displacement of the car sales distributions in Figure 4. We show that this strategy can be cast as an optimal transport problem.

Optimal transport has a long history in economics and operations research, see [Kantorovitch \(1958\)](#) for an early treatment. Optimal transport has recently

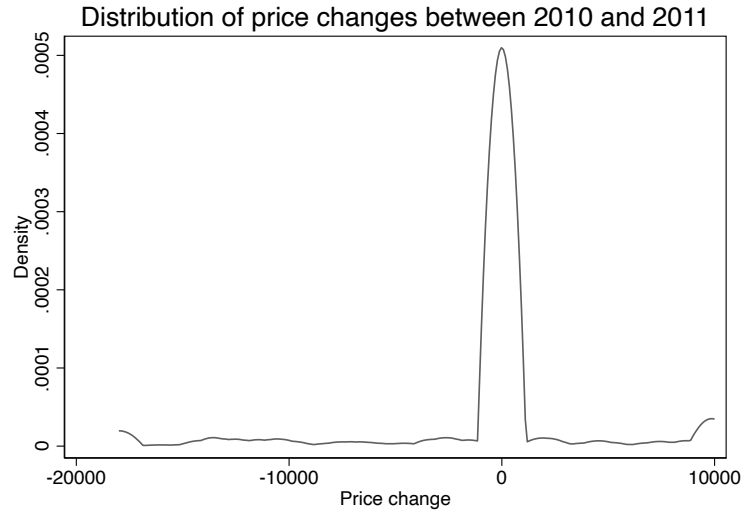


Figure 3: Distribution of price changes

witnessed renewed interest in economics and applied econometrics. It has found sophisticated uses in both economic and econometric theory, for instance in the analysis of identification of dynamic discrete-choice models (Chiong et al. (2016)), in vector quantile regression (Carlier et al. (2016)), and in empirical matching models (Galichon et al. (2018)).

A point we press here is that optimal transport is also a natural and easily implementable method for transparent, applied empirical analysis under weak and economically motivated assumptions. As in standard regression analysis, we must choose a cost specification, which in the case of optimal transport is the distance between points in the support of the compared probability mass functions. The choice of this cost specification is guided by the economics of the problem. Output, including confidence intervals, may be presented in a way that is typical of regression tables. Our application serves as an example.

We develop two estimators. The first uses the displacement, the shift in the sales distribution following the rationing, in Beijing to estimate the volume of trade. The second estimator uses a difference between the displacement in Beijing and the corresponding displacement in Tianjin. Similar to a difference-in-differences es-

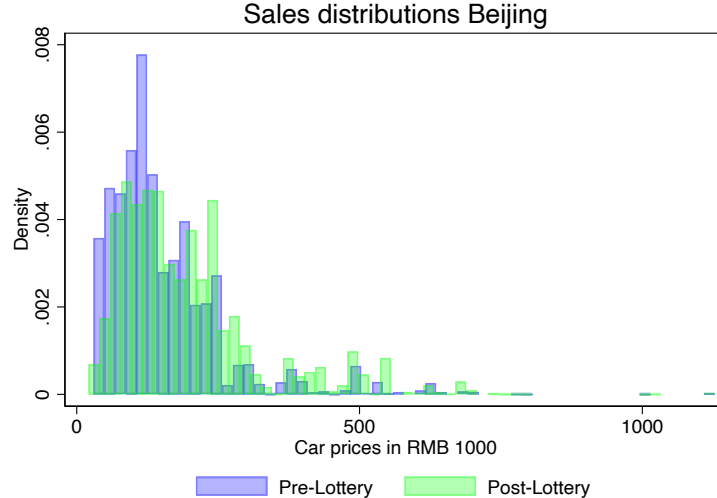


Figure 4

timator, the difference in displacement helps us control for common economic trends.

In the following, we refer to a *buyer* as a household that purchases a car with a license plate it has bought or rented illegally. We refer to a *seller* as a household that won a license plate and then either decided to sell or rent it to a buyer. We refer to the transaction between a buyer and a seller as a *trade*. In the following, it is useful to adopt a potential outcomes notation. Let $\mathbb{P}_{pre}(r, b)$ and $\mathbb{P}_{post}(r, b)$ be the potential population sales distributions in Beijing pre-lottery and post-lottery, respectively. The index $r \in \{0, 1\}$ is one when there is rationing post-lottery, and zero otherwise. The index $b \in \{0, 1\}$ is one when there is a black market post-lottery, and zero otherwise.

5.1 Before-and-after

We start by making making three assumptions that are sufficient to identify a lower bound for the volume of trade. The first rules out anticipatory sales effects of the rationing: the pre-lottery sales distribution does not depend on whether there is rationing and/or a black market in the post-lottery period.

Assumption 1 *No anticipation:*

$$\mathbb{P}_{pre}(r, b) = \mathbb{P}_{pre},$$

for $r \in \{0, 1\}$ and $b \in \{0, 1\}$.

Assumption 1 seems reasonable given the brief period of time between the announcement of the rationing and its implementation (one month). We maintain this assumption throughout. We next assume away time trends that are unrelated to the lottery. Specifically, we assume that the pre-lottery sales distribution is equal to the potential sales distribution post-lottery when there is no rationing and no black market.

Assumption 2 *No time trends:*

$$\mathbb{P}_{pre} = \mathbb{P}_{post}(0, 0).$$

Assumption 2 rules out trends in car preferences such as, say, increasing demand for SUVs. It implies that a household that wins the lottery, and does not trade its license plate, would purchase the same car as in a world without rationing. It also precludes car dealerships from changing the car prices following the rationing, an assumption that is largely verified in the data, see Figure 3, and consistent with the resale price maintenance contracts that the industry used at the time, see the discussion in Section 2. Our next assumption concerns the effect of rationing on car preferences.

Assumption 3 *No general equilibrium effects:*

$$\mathbb{P}_{post}(0, 0) = \mathbb{P}_{post}(1, 0)$$

Assumption 3 forces the potential sales distribution post-lottery, without rationing and with no black market, to equal the potential sales distribution post-lottery, with rationing and with no black market. One example of general equilibrium effects is that rationing license plates, which leads to fewer cars on the roads, increases the

willingness-to-pay for a car. Since the rationing controls the influx of new cars to Beijing (on the order of 260 000 per year), and not the stock (on the order of five million at the time), we believe that such general equilibrium effects are negligible. Assumption 3 also rules out a direct effect of the rationing on the car preferences. This assumption is common in the literature on rationing, but it is untestable, and it has been contested in the past. One example is Tobin (1952) which caters the idea that rationing can change tastes over time:

“Experience under rationing may alter the consumer’s scale of preferences. He may learn to like pattern of expenditures into which rationing forces him, or to dislike it even more intensely than if it had not been forced upon him.” (p. 548).

Shen et al. (2019) similarly explores the hypothesis that the price jump is explained by a sense of luck associated with winning the lottery which in turn changed the car preferences. This hypothesis is inconsistent with Assumption 3. Shen et al.’s hypothesis and our own seem to be identifying assumptions that can not be tested directly in our data, but must be assessed according to their plausibility.

If license plates are reallocated by black market trade, then $\mathbb{P}_{post}(1, 1)$ may not equal $\mathbb{P}_{post}(1, 0)$. It seems plausible that richer buyers buy license plates from poorer sellers. Previous literature has shown that richer buyers buy more expensive cars (Li (2018), Xiao et al. (2017), and Tan et al. (2019)). If a buyer purchases a different car than the seller would have purchased, had he not traded his license, then we can think of that trade as shifting, or transporting, mass from $\mathbb{P}_{post}(1, 0)$ to $\mathbb{P}_{post}(1, 1)$. Summing over the minimum number of such trades required to account for the shift in distributions, we get a lower bound estimate of black market trade. It is a lower bound since, for instance, we will not detect a trade where the buyer purchases the same car as the seller would have purchased, had he not sold his license. Since all license plates can be traded without shifting the sales distribution, the upper bound for the volume of trade is 100% of the quota.⁹ The counterfactual

⁹In Section 6, we impose further assumptions that allow us to bound the volume of trade from above.

distribution $\mathbb{P}_{post}(1, 0)$ is identified since Assumptions 1 and 2 imply that

$$\mathbb{P}_{pre} = \mathbb{P}_{post}(1, 0). \tag{1}$$

We could have given condition (1) directly as an assumption. By distinguishing between Assumptions 2 and 3 and only relaxing Assumption 2 in Section 5.3, where we develop a difference-in-differences like estimator, we get a clear comparison of sufficient conditions for identification.

Under our current assumptions, the minimum amount of mass we need to transport between the (observable) distributions \mathbb{P}_{pre} and $\mathbb{P}_{post}(1, 1)$ estimates a lower bound for the volume of trade. To save on notation, we drop the potential outcome arguments $\mathbb{P}_{post}(1, 1)$ and, abusing notation, let $\mathbb{P}_j(x)$ for $j \in \{pre, post\}$ and $x \in \mathcal{X}$ denote the densities in the rest of this section. If we observed the population distributions \mathbb{P}_{pre} and \mathbb{P}_{post} directly, we could compute the desired lower bound via the following oracle problem

$$\begin{aligned} OT(\mathbb{P}_{pre}, \mathbb{P}_{post}) &= \min_{\gamma \in \Gamma} \int c(x_0, x_1) \gamma(x_0, x_1) dx_0 dx_1 \\ &\text{subject to} \\ &\int \gamma(x_0, x_1) dx_1 = \mathbb{P}_{pre}(x_0), \quad \text{for all } x_0 \in \mathcal{X}, \\ &\int \gamma(x_0, x_1) dx_0 = \mathbb{P}_{post}(x_1), \quad \text{for all } x_1 \in \mathcal{X}. \end{aligned} \tag{2}$$

The optimal transport cost has two components. The first component is $c(x_0, x_1)$, the cost of transport between point x_0 in \mathbb{P}_{pre} and point x_1 in \mathbb{P}_{post} . The second is $\gamma(x_0, x_1)$, which is the amount of mass that is transported between point x_0 in \mathbb{P}_{pre} and point x_1 in \mathbb{P}_{post} . The optimization is over $\gamma \in \Gamma$, where Γ is the set of all bivariate probability distributions. The constraints ensure that the marginal distributions exactly equate the observed sales distributions $\mathbb{P}_{pre}(x)$ and $\mathbb{P}_{post}(x)$.

For the cost specification, we pick

$$c(x_0, x_1) = \mathbb{1}(|x_1 - x_0| > 0). \quad (3)$$

This function assigns the same cost for transport between any pair of points x_0, x_1 when $x_0 \neq x_1$, and zero cost when $x_0 = x_1$. The objective function of the optimal transport problem therefore represents a volume of trade that transforms the pre-lottery distribution into the post-lottery distribution. The optimal transport cost OT for the problem in (2) can be interpreted as the smallest volume of trade that is consistent with the data. Our estimand is then

$$s(\mathbb{P}_{pre}, \mathbb{P}_{post}) = OT(\mathbb{P}_{pre}, \mathbb{P}_{post}). \quad (4)$$

An example may be instructive. Suppose cars are sold either at a price of 1 or a price of 2, so $\mathcal{X} = \{1, 2\}$. Pre-lottery, eight cars are sold: six at price 1 and two at price 2, such that the population distribution is $\mathbb{P}_{pre} = \frac{1}{8}[6, 2]$. Post-lottery, the quota is set to four cars. One is sold at price 1 and three at price 2, such that the population distribution is $\mathbb{P}_{post} = \frac{1}{4}[1, 3]$. It is immediately clear that at least two license plates must be traded to rationalize the shift in the sales distribution. One solution γ which satisfies the constraints (2) is

$$\gamma = \frac{1}{4} \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}.$$

The diagonal terms imply no transport, and hence no cost. We can therefore calculate the optimal transport by summing all off-diagonal terms of γ . This gives $OT(\mathbb{P}_{pre}, \mathbb{P}_{post}) = 50\%$ as a lower bound of the share of illegal trade. This example assumed away sampling variation in the empirical distributions. We account for sampling variation in Section 5.2 below.

There may be more than one γ that solves (2). Our parameter of interest is however not γ , but $OT(\mathbb{P}_{pre}, \mathbb{P}_{post})$, which is clearly unique. Multiplicity of γ

solutions therefore has no implications for the identification of a lower bound for the volume of trade.

Though richer buyers empirically purchase more expensive models, it is not obvious that trade in license plates shifts the sales distribution towards more expensive car models. One reason is that the transaction price for a license plate serves as a tax on the car purchase which may lead a richer buyer to trade down relative to the car that the same buyer would have purchased if she did not have to first buy a license plate. The estimand in (4) however does not require that buyers purchase more expensive cars than the sellers would have purchased if they could not trade, i.e., that the sales distribution shifts to the right. Any displacement of the sales distributions is counted as trade.

5.2 Before-and-after estimator

We need to account for the fact that we have access to the sample distributions, and not the population distributions. The sample distribution pre-lottery is $\hat{\mathbb{P}}_{pre, n_{pre}}$, where n_{pre} is the number of observations, and $\hat{\mathbb{P}}_{post, n_{post}}$ and n_{post} are defined analogously for the post-lottery period. We approximate the population problem in (2) with its discrete sample equivalent

$$\begin{aligned}
 OT(\hat{\mathbb{P}}_{pre, n_{pre}}, \hat{\mathbb{P}}_{post, n_{post}}) &= \min_{\gamma \in \Gamma} \sum_{i, j} \gamma_{i, j} \mathbf{C}_{i, j} \\
 &\text{subject to} \\
 \sum_i \gamma_{i, j} &= \hat{\mathbb{P}}_{pre, n_{pre}}(j), \quad \text{for all } j, \\
 \sum_j \gamma_{i, j} &= \hat{\mathbb{P}}_{post, n_{post}}(i), \quad \text{for all } i, \\
 \gamma_{i, j} &\geq 0 \quad \text{for all } i, j.
 \end{aligned} \tag{5}$$

where $\mathbf{C} = \mathbf{1}\mathbf{1}^T - \text{diag}(\mathbf{1})$. The observed sample distributions $\hat{\mathbb{P}}_{pre, n_{pre}}$ and $\hat{\mathbb{P}}_{post, n_{post}}$ are plotted in Figure 5. We cannot directly apply the program in (5) to $\hat{\mathbb{P}}_{pre, n_{pre}}$ and $\hat{\mathbb{P}}_{post, n_{post}}$, due to sampling uncertainty which will inflate the transport estimate.

Even in the null case of no trade in license plates, i.e., if both $\hat{\mathbb{P}}_{pre, n_{pre}}$ and $\hat{\mathbb{P}}_{post, n_{post}}$

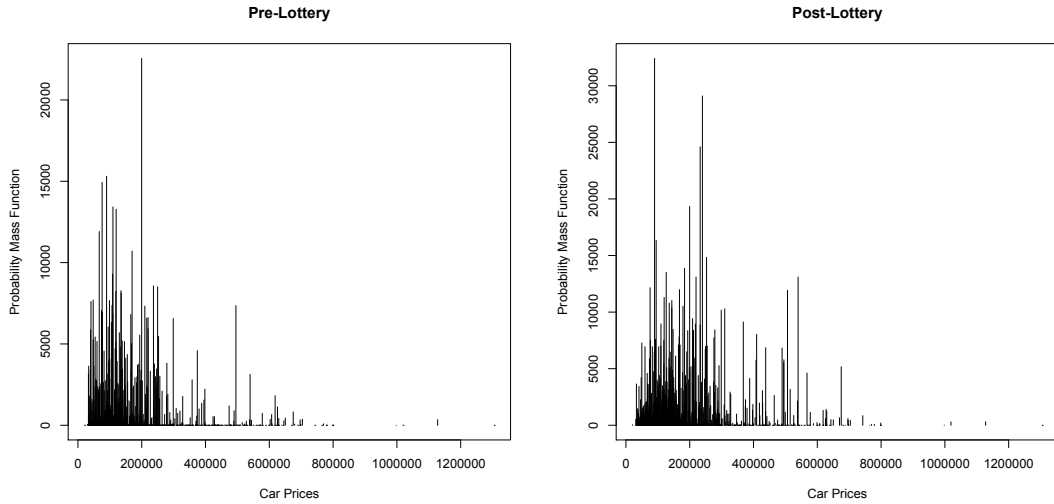


Figure 5: Exact empirical sales distributions before-and-after.

were composed of draws from the same distribution \mathbb{P}_{pre} , sampling variation would lead to differences in the realized distributions.

One transparent approach to smooth out sampling variation is to ignore small moves: we can attribute zero cost to transport mass up to some small distance $d \geq 0$, where the distance is measured in RMB. We want to choose d small enough to detect as many trades as possible, but large enough to not confound sampling variation with trades.

To control the sampling variation using d , we measure the variation in transport cost between two empirical distributions sampled from the same population distribution. We select d such that the transport cost in the placebo problem is at least an order of magnitude smaller than the estimated transport cost for the target problem. This makes the sampling variation statistically and economically negligible.

To implement this smoothing criterion, we replace the cost matrix \mathbf{C} with a

threshold version

$$\mathbf{C}_{i,j}(d) = \mathbb{1}(|x_i - x_j| > d),$$

where x_i is the i^{th} entry of \mathcal{X} . We then solve the sample optimal transport problem

$$\begin{aligned} OT_d(\hat{\mathbb{P}}_{pre,n_{pre}}, \hat{\mathbb{P}}_{post,n_{post}}) &= \min_{\gamma \in \Gamma} \sum_{i,j} \gamma_{i,j} \mathbf{C}_{i,j}(d) \\ &\text{subject to} \\ \sum_i \gamma_{i,j} &= \hat{\mathbb{P}}_{pre,n_{pre}}(j), \text{ for all } j, \\ \sum_j \gamma_{i,j} &= \hat{\mathbb{P}}_{post,n_{post}}(i), \text{ for all } i, \\ \gamma_{i,j} &\geq 0 \text{ for all } i, j. \end{aligned} \tag{6}$$

We choose the parameter d such that the estimated transport cost between two distributions sampled from \mathbb{P}_{pre} is approximately zero, that is, when we know that the cost in the population problem is exactly zero. In order to calibrate d on pairs of empirical distributions drawn from \mathbb{P}_{pre} , we use $\hat{\mathbb{P}}_{pre,n_{pre}}$ as an estimate of \mathbb{P}_{pre} . We then obtain new simulated empirical distributions by drawing $X_{j,1}, \dots, X_{j,n_{pre}}$ from $\hat{\mathbb{P}}_{pre,n_{pre}}$, for $j = 1, 2$, and by collecting their corresponding empirical distributions $\tilde{\mathbb{P}}_{pre,n_{pre}}^{(j)}$.

For a particular set of simulated distributions, the placebo transport costs are $OT_d(\tilde{\mathbb{P}}_{pre,n_{pre}}^{(1)}, \tilde{\mathbb{P}}_{pre,n_{pre}}^{(2)})$ from program (6), with the constraints on the marginal distributions replaced with $\tilde{\mathbb{P}}_{pre,n_{pre}}^{(1)}$ and $\tilde{\mathbb{P}}_{pre,n_{pre}}^{(2)}$. Our estimator of the placebo transport costs is $\hat{s}_{placebo}(d) = \mathbb{E} \left[OT_d(\tilde{\mathbb{P}}_{pre,n_{pre}}^{(1)}, \tilde{\mathbb{P}}_{pre,n_{pre}}^{(2)}) \right]$, where the average is over the set of simulated distributions.

We next probe the sensitivity of the estimates over a range of values of d 's. Figure 7 plots $\hat{s}(d) = OT_d(\hat{\mathbb{P}}_{pre,n_{pre}}, \hat{\mathbb{P}}_{post,n_{post}})$ and $\hat{s}_{placebo}(d)$ against d . While $\hat{s}(d)$ estimates the lower bound for the market share, $\hat{s}_{placebo}(d)$ estimates the sampling uncertainty as a function of d .

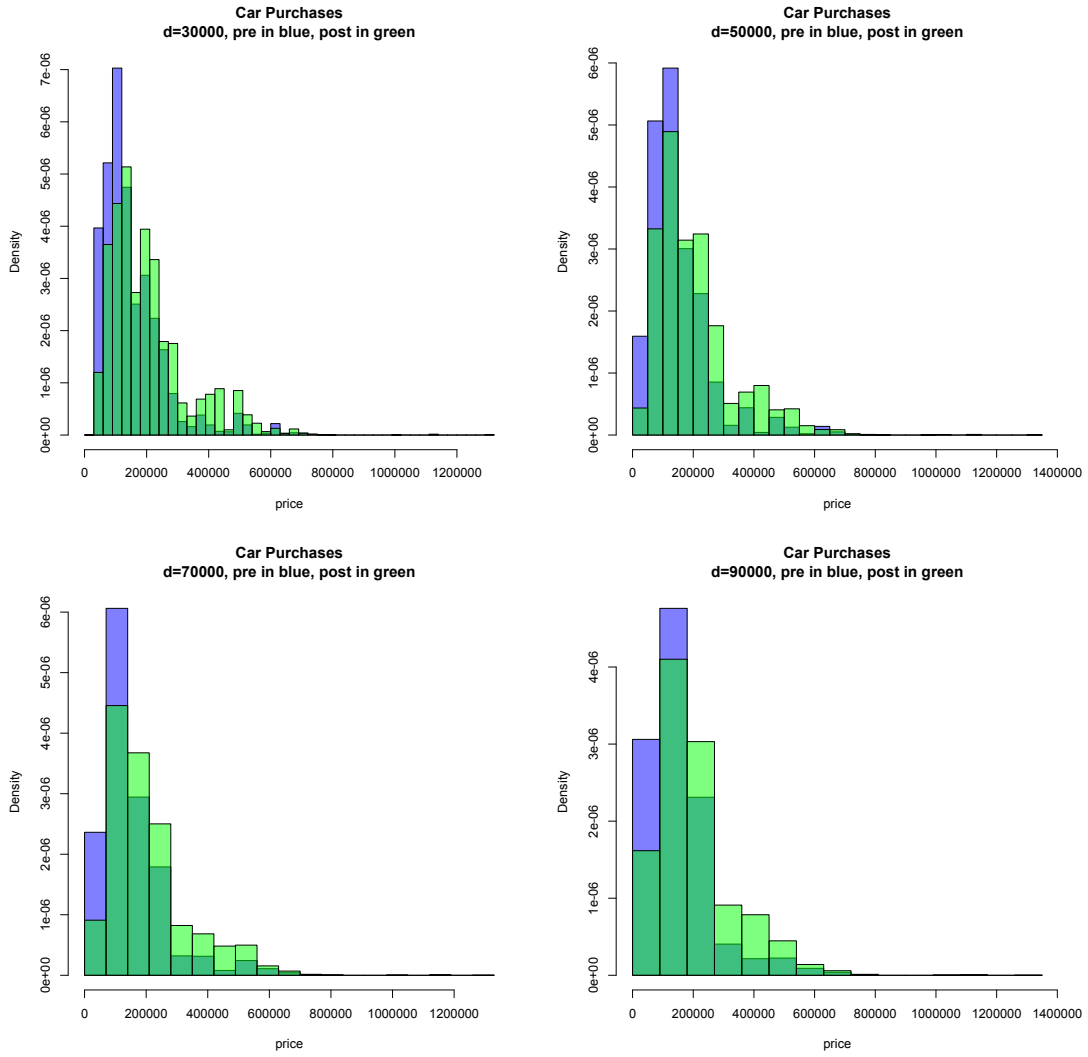


Figure 6: Distributions smoothed at candidate d 's.

d	10 000	20 000	30 000	50 000	70 000	90 000
$\hat{s}(d)$	14%	11%	8%	5%	4%	3%
$\hat{s}(d)_{placebo}$	0.05%	0.03%	0.03%	0.001%	0.001%	0.001%

Table 4: Beijing placebo transport costs

The results in Table 4 suggest that using a d larger than 30 000 gives placebo sampling uncertainty less than two orders of magnitudes smaller than the estimate, which seems overly conservative. The placebo transport cost $\hat{s}_{placebo}(10\,000)$ is already small in absolute terms and still two orders of magnitude smaller than $\hat{s}(10\,000)$. This tells us that sampling uncertainty is negligible at this d , so we choose $\hat{s}(10\,000) = 14\%$ as our preferred before-and-after estimate.

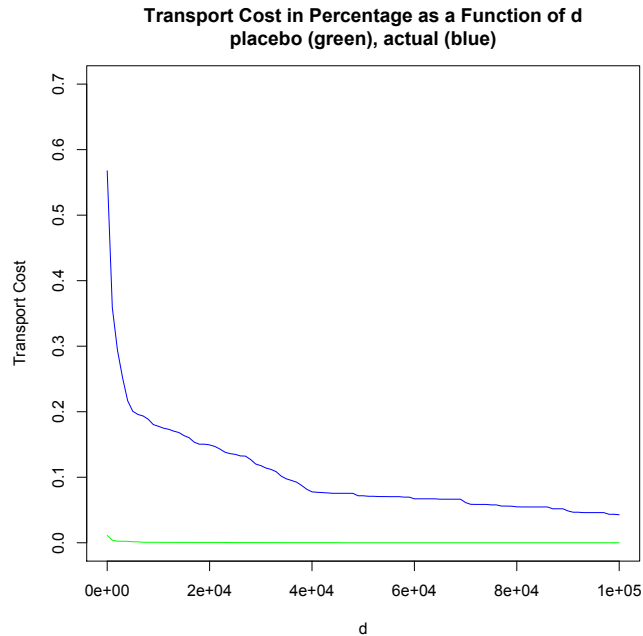


Figure 7: True and placebo costs as a function of d

5.3 Difference-in-transports

The results in the previous section relied on Assumption 2 which precludes time trends. This assumption may not hold. For instance, income growth or changes in car preferences could lead the sales distributions to shift for reasons that are unrelated to black market trade. In Section 4, we used the sales data from Tianjin to control for common trends using difference-in-differences regressions. We draw on

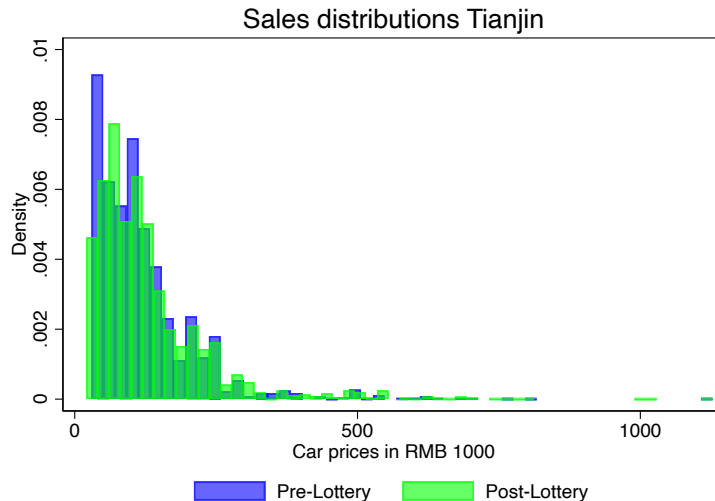


Figure 8

the same idea for our second estimator. We want to use the observable displacement of the Tianjin sales distribution to account for external factors (changes in income, preferences etc) that would have shifted the sales distribution in Beijing, had it not introduced a rationing. We call this analog to the standard difference-in-differences estimator a difference-in-transports. Figure 8 shows that the sales distribution in Tianjin, where there was no rationing, may have shifted slightly to the right. The displacement in Tianjin is clearly less pronounced than the corresponding displacement in Beijing, which is seen in Figure 4.

The target quantity is the displacement between the observed, post-lottery Beijing sales distribution, with rationing and a black market, and the counterfactual, post-lottery Beijing sales distribution, with rationing and no black

market,

$$OT(\mathbb{P}_{Beijing,post}(1, 0), \mathbb{P}_{Beijing,post}(1, 1)). \quad (7)$$

Its oracle difference-in-transports proxy is given by the transport cost for the observed Beijing sales distributions before and after the lottery with rationing and a black market, net of the transport cost that would have been incurred if there was rationing, but no black market.

$$OT(\mathbb{P}_{Beijing,pre}, \mathbb{P}_{Beijing,post}(1, 0)) - OT(\mathbb{P}_{Beijing,pre}, \mathbb{P}_{Beijing,post}(1, 1)) \quad (8)$$

By the triangle inequality of Proposition 55 in [Peyré \(2018\)](#), the population quantity s_{dit} with $\mathbf{C} = \mathbf{C}(0)$ produces a lower bound on the quantity of interest,

$$\begin{aligned} & OT(\mathbb{P}_{Beijing,pre}, \mathbb{P}_{Beijing,post}(1, 0)) - OT(\mathbb{P}_{Beijing,pre}, \mathbb{P}_{Beijing,post}(1, 1)) \\ & \leq OT(\mathbb{P}_{Beijing,post}(1, 1), \mathbb{P}_{Beijing,post}(1, 0)). \end{aligned} \quad (9)$$

This parameter depends on the unobserved quantity $OT(\mathbb{P}_{Beijing,pre}, \mathbb{P}_{Beijing,post}(1, 0))$. In contrast to the before-and-after case, we now relax [Assumption 2](#) which allows the sales distribution in Beijing to shift between the pre-and post lottery period for reasons other than the lottery, i.e. $\mathbb{P}_{Beijing,post}(0, 0) \neq \mathbb{P}_{pre}$. We instead use shifts in the Tianjin sales distribution to control for external factors. We replace [Assumption 2](#) with the following parallel trends-like assumption.

Assumption 4 *Equal displacement:*

$$OT(\mathbb{P}_{Beijing,pre}, \mathbb{P}_{Beijing,post}(0, 0)) = OT(\mathbb{P}_{Tianjin,pre}(0, 0), \mathbb{P}_{Tianjin,post}(0, 0)).$$

This assumption requires the observed displacement in Tianjin to equal the corresponding, counterfactual displacement in Beijing if there was no rationing and no

black market trade. Assumptions 3 and 4 together imply

$$OT(\mathbb{P}_{Beijing,pre}, \mathbb{P}_{Beijing,post}(1, 0)) = OT(\mathbb{P}_{Tianjin,pre}, \mathbb{P}_{Tianjin,post}(0, 0)). \quad (10)$$

We can therefore write (8) in terms of optimal transports between observable distributions

$$s_{dit} = OT(\mathbb{P}_{Beijing,pre}, \mathbb{P}_{Beijing,post}(1, 1)) - OT(\mathbb{P}_{Tianjin,pre}(0, 0), \mathbb{P}_{Tianjin,post}(0, 0)). \quad (11)$$

which is the analog of the traditional difference-in-differences estimand. Although s_{dit} under Assumption 4 is well estimated by

$$OT_d\left(\hat{\mathbb{P}}_{Beijing,pre}, \hat{\mathbb{P}}_{Beijing,post}(1, 1)\right) - OT_d\left(\hat{\mathbb{P}}_{Tianjin,pre}(0, 0), \hat{\mathbb{P}}_{Tianjin,post}(0, 0)\right)$$

for a small, but non-zero value of d , we would like our difference-in-transport estimate to give a lower bound on $OT_d\left(\hat{\mathbb{P}}_{Beijing,post}(1, 1), \hat{\mathbb{P}}_{Beijing,post}(1, 0)\right)$, the sample version of the target quantity, and thus maintain the desirable population inequality in (14) in sample. We show below that the following estimator delivers this property.

$$\hat{s}_{dit}(d) = OT_{2d}\left(\hat{\mathbb{P}}_{Beijing,pre}, \hat{\mathbb{P}}_{Beijing,post}(1, 1), \right) - OT_d\left(\hat{\mathbb{P}}_{Tianjin,pre}(0, 0), \hat{\mathbb{P}}_{Tianjin,post}(0, 0)\right). \quad (12)$$

Under Assumption 4, and for d sufficiently large to smooth out sampling variation, $\hat{s}_{dit}(d)$ is a proxy for

$$OT_{2d}\left(\hat{\mathbb{P}}_{Beijing,pre}, \hat{\mathbb{P}}_{Beijing,post}(1, 1)\right) - OT_d\left(\hat{\mathbb{P}}_{Beijing,pre}, \hat{\mathbb{P}}_{Beijing,post}(1, 0)\right), \quad (13)$$

which we can show lower bounds the sample analog of the population quantity of interest.

Theorem 1 *Let $OT_d(\hat{\mathbb{P}}, \tilde{\mathbb{P}})$ be the discrete optimal transport problem described in*

d	1 000	5 000	10 000	20 000	30 000
Beijing	31.1%	17.2%	14.4%	11.5%	8.3%
95% CI	[30.7, 32.7]%	[16.7, 18.9]%	[13.5, 16.5]%	[10.5, 12.5]%	[7.3, 9.7]%
Tianjin	22.5%	5.8%	2.8%	1.7%	0.9%
95% CI	[22.4, 24.2]%	[5.1, 7.2]%	[2.4, 3.5]%	[1.2, 2.1]%	[0.7, 1.4]%
Difference	8.6%	11.4%	11.6%	9.8%	7.4%
95% CI	[7.2, 9.8]%	[10.2, 13.0]%	[10.8, 13.0]%	[8.7, 10.6]%	[6.2, 8.7]%

Table 5: The first and third rows report the estimated displacements in Beijing and Tianjin. For Beijing, these are the before-and-after estimates of s . The fifth row reports the difference-in-transports estimate of s_{dit} . Confidence intervals are calculated as the quantiles from subsampling.

(6) with $\mathbf{C} = \mathbf{C}(d)$, and let $\hat{\mathbb{P}}_a$, $\hat{\mathbb{P}}_b$ and $\hat{\mathbb{P}}_c$ be three probability mass functions. Then, the following inequality holds

$$OT_{2d}(\hat{\mathbb{P}}_a, \hat{\mathbb{P}}_b) - OT_d(\hat{\mathbb{P}}_c, \hat{\mathbb{P}}_b) \leq OT_d(\hat{\mathbb{P}}_a, \hat{\mathbb{P}}_c). \quad (14)$$

Proof: This follows immediately from the argument of Proposition 55 in [Peyré \(2018\)](#), and the analogous triangle inequality for the pseudo-distance $\mathbf{C}(d)$,

$$\mathbb{1}(|x_i - x_j| > 2d) - \mathbb{1}(|x_j - x_k| > d) \leq \mathbb{1}(|x_i - x_k| > d),$$

where x_l is the l^{th} entry of \mathcal{X} . \square

In particular, Theorem 1 stipulates that if Assumption 4 holds in sample, i.e., $OT_d(\hat{\mathbb{P}}_{Beijing,pre}, \hat{\mathbb{P}}_{Beijing,post}(0, 0)) = OT_d(\hat{\mathbb{P}}_{Tianjin,pre}(0, 0), \hat{\mathbb{P}}_{Tianjin,post}(0, 0))$, then

$$\hat{s}_{dit}(d) \leq OT_d(\hat{\mathbb{P}}_{Beijing,post}(1, 1), \hat{\mathbb{P}}_{Beijing,post}(1, 0)). \quad (15)$$

This is a desirable property: our estimator satisfies (15), the sample analog to bound (7), which is the quantity of interest in the population.

The difference-in-transports estimator does not lend itself as readily to choosing d using the sampling uncertainty in placebo transports as the before-and-after estimator in Section 5.1 does. We therefore use a different criterion. As before, we want d small enough to detect as many trades as possible from displacement in the

sales distributions, yet large enough that we do not mistake sampling uncertainty for trades. We report the difference-in-transport results for a range of values of d in Table 5. The difference-in-transport estimates are stable for d in the range between 5 000 and 10 000. We pick $d = 10\,000$ as the value for the tuning parameter. At $d = 10\,000$, the difference-in-transport estimate of the proportion of trades is 11.4%, with a 95% confidence interval of [10.8, 13.1]. The confidence intervals are computed by subsampling.

The parameter d can be interpreted as an upper bound for the types of change in prices which we do not want to count as trade in order to produce the lower bound for the volume of trade. Interpreting d as such suggests reinterpreting our cost function $C_{i,j}(d)$ as applying $C_{i,j}(0)$ after correcting for idiosyncratic variation in the distributions.

Table 5 shows that the estimated differences in displacements are stable even up to negligible values of d , although the transportation costs decrease in d for both Beijing and Tianjin. This fact has the comforting implication that our difference-in-transport estimates are not particularly sensitive to the choice of the tuning parameter d . It suggests that the displacement of distributions sampled from the same data generating process is captured and corrected by the difference-in-differences when they are not discarded by a large enough d . The low sensitivity to d also means that we can use the difference for transport costs corresponding to an arbitrarily small value of the tuning parameter d as an estimate. The difference-in-transport estimate is close to, but slightly lower than the estimate that only uses the Beijing distribution, which suggests that there are common external factors shifting the sales distribution.

Our difference-in-transport estimator bears some resemblance to [Athey and Imbens \(2006\)](#)'s non-parametric difference-in-differences estimator (change-in-changes). Under an assumption of rank invariance on the unobserved heterogeneity, the change-in-changes gives the full distribution of treatment effects, is invariant

to transformations of the outcome variable, and controls for a common trend at the same quantile level. In comparison, our difference-in-transports estimator does not make a rank invariance assumption and it does not deliver a distribution of treatment effects.

6 Inferring transaction costs and prices

We infer the unobserved transaction costs and prices from a market equilibrium model that combines our estimated volume of trade with empirical results from the literature. We derive our model from a common version of the Coase Theorem: If the initial allocation leaves gains from trade, as is expected when license plates are allocated by lottery, households will transact until the gains from trade are exhausted. A market price will form that equates the number of lottery winners willing to sell their license plates to the number of prospective buyers willing to buy a license. Transaction costs, which may be both pecuniary and non-pecuniary, are frictions that reduce the volume of trade. We bound the transaction costs by treating these as a tax on trade necessary to rationalize our estimated volume of black market trade.

One likely important transaction cost component is potential legal liabilities. For instance, a new car must be registered in the name of the legal owner of the license plate, and it is therefore the formal owner, and not the driver, which in some cases will be liable for traffic accidents. The fact that trade in license plates is illegal may also give rise to non-pecuniary transaction costs, such as feelings of culpability. Search costs are another component: buyers and sellers need to meet in the market. The reported existence of online market places for license plates however suggests that search costs are small.¹⁰

Potential enforcement, and perhaps moral hazard of the kind associated with

¹⁰Bhave and Budish (2018) notes that online market places for ticket resale decreases transaction costs and increases the volume of resold tickets in the secondary market.

leaving the formal car ownership in someone else’s name, are likely the key transaction cost components. While there may be further transaction cost components, we do not attempt to distinguish between these. The transaction costs in our model are therefore a simple quantitative measure of the frictions in the black market.

We derive a demand and supply curve using Li’s estimated willingness-to-pay for license plates. These estimates are represented by a function $v(n)$ with range RMB $[0,280000]$.¹¹ In the following, we take Li’s estimated willingness-to-pay for a license plate $v(n)$ to be known without sampling variation. It is useful in the following to write $v(n)$ in terms of a cumulative distribution function $F(v)$ which returns the share of households in the market with a willingness-to-pay less than v , i.e. $F(v) = \frac{1}{N} \sum_{n=1}^N \mathbb{1}(v(n) \leq v)$. Since $v(n)$ is strictly decreasing, we can recover F from v as

$$\frac{v^{-1}(v)}{N} = 1 - F(v).$$

We assume that only prospective car buyers enter the lottery.

Assumption 5 *The lottery draws q winners with equal probability from the population of N prospective car buyers, whose willingness-to-pay for a license plate is distributed according to the cumulative distribution function F .*

It follows immediately from Assumption 5 that both the q sellers’ (winners) and the $N - q$ buyers’ willingness-to-pay are distributed according to cdf F . We set $q = 260000$, the quota in 2011, and the market size to $N = 700000$, the total sales of new cars in 2010 before the lottery. Assumption 5 rules out speculators, which we define to be those who enter the lottery with no intention of buying a car if they win a license plate. An influx of speculators is consistent with the sharp decline in the lottery odds, which went from 10% in the first auction in January 2011, to a monthly average of 4% for 2011, and dropped further to a monthly

¹¹Shanjun Li has generously shared his estimates with us, which we have slightly modified. While Li’s Figure 3 shows positive willingness-to-pay for quantities far beyond the unrationed, pre-lottery market equilibrium, we require that the marginal willingness-to-pay for a license plate at the unrationed market equilibrium is zero, i.e. that $v(N) = 0$. Our results are however not very sensitive to this modification.

average of 2% in 2012. We relax Assumption 5 and allow for speculators in Section D.

We next make an assumption about how prices and transaction costs form.

Assumption 6 *Each trade generates transaction costs of $2t$ which are equally borne by the buyer and the seller. The transaction price p equates demand to supply, given the transaction costs.*

A buyer is willing to buy a license plate if $v > p + t$, and a seller is willing to sell a license plate if $v \leq p - t$. There are $N - q$ rationed prospective buyers that demand license plates, while there are q lottery winners that supply licenses. This gives the demand and supply functions

$$\begin{aligned} D(p, t) &= (N - q)(1 - F(p + t)) \\ S(p, t) &= qF(p - t) \end{aligned} \tag{16}$$

Figure 9 plots the demand and supply curves derived from Li's estimates. Suppose first that there are no transaction costs. Then demand equals supply at price p_{notc} and the quantity is q_{notc} . The volume of trade that occurs without transaction costs q_{notc} is less than the quota q since lottery winners with valuation in excess of the market clearing transaction price prefer not to trade. Adding Assumptions 5 and 6 brings the upper bound of the volume of trade down from 100% of the quota to $s_{upper} = \frac{q_{notc}}{q} = \frac{N-q}{N} = 62\%$ of the quota. The estimate of s_{upper} implies that though our lower bound estimate is $\hat{s} = 11\%$ of the quota, it is $\frac{11\%}{62\%} = 18\%$ of the largest number of trades that this market can support in equilibrium.

The transaction costs t in Figure 9 are seen to drive a wedge between the supply and demand that lowers the volume of illegal trade to sq . The total transaction costs, summed over buyers and sellers, that rationalize the trade are hence $2tsq$.

Theorem 2 *Suppose that F is known, continuous, and strictly increasing. Suppose*

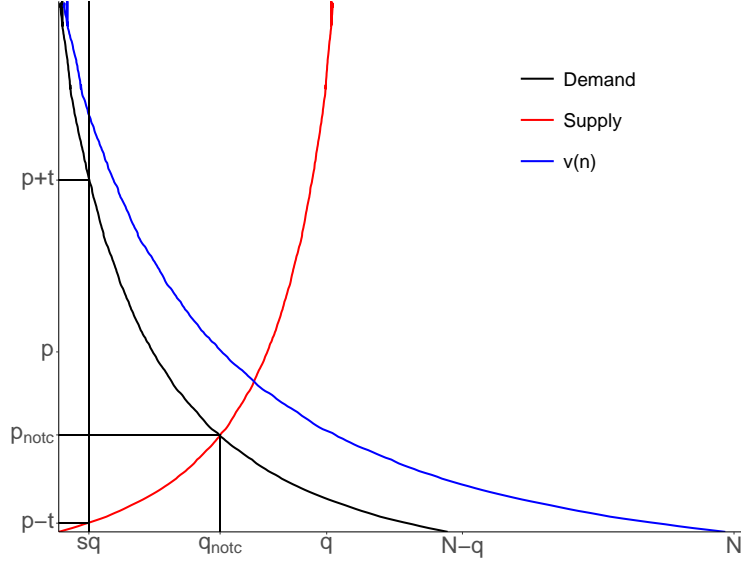


Figure 9: Demand and supply derived from $v(n)$. The figure is drawn to scale with quantity in 10 000 and price in RMB 1000.

that the market clears at transaction price p and transaction cost t for a known volume of sq black market trades. Then the transaction prices and costs are identified.

Proof:

For any pair of transaction costs and prices, the valuation of the marginal seller is $v_{seller} = p - t$ and the valuation of the marginal buyer is $v_{buyer} = p + t$. Together, we get

$$\begin{aligned}
 p &= \frac{1}{2} (v_{seller} + v_{buyer}) \\
 t &= \frac{1}{2} (v_{buyer} - v_{seller}).
 \end{aligned}
 \tag{17}$$

Equating demand to supply using (16) at the estimated lower bound of trades sq , the marginal valuations are uniquely recovered by inverting $F^{-1}(s) = v_{seller}$ and $F^{-1}\left(1 - \frac{sq}{N-q}\right) = v_{buyer}$. It follows immediately that p and t are uniquely determined given sq . ■

We can alternatively dispense with the homogenous transaction cost assumption and allow buyers and sellers to bear different transaction costs, i.e. $t_{buyer} \neq t_{seller}$. Though we can not jointly identify the buyer and seller specific transaction costs

along with the transaction price, a pair $\tilde{p}_{buyer} = p + t_{buyer}$ and $\tilde{p}_{seller} = p - t_{seller}$ is identified. This interpretation does not affect the identification of the total transaction costs, which are $sq(\tilde{p}_{buyer} - \tilde{p}_{seller})$ either way. Since these interpretations are observationally equivalent, we can not infer which side of the market bears the majority of the transaction costs. In the following, we therefore maintain the homogenous transaction cost interpretation. The willingness-to-pay F , which we derived from Li’s estimates, satisfies the continuity and monotonicity restrictions.

The implied transaction costs and prices at \hat{s} are given in Table 6, along with 95% confidence intervals.¹² We show in the appendix that \hat{t} is an upper bound for the transaction cost at \hat{s} , and that the implied \tilde{p}_{buyer} and \tilde{p}_{seller} at \hat{s} are upper and lower bound estimates, respectively. In the model, trades take place between a selection of buyers with particularly high valuations and sellers with particularly low valuations. This pattern is consistent with the literature on the resale market of tickets where trades are observed (Bhave and Budish (2018), Leslie and Sorensen (2013)) and where search costs are relatively small.

The implied transaction price at the lower bound estimate of the volume of trade is RMB 105 000, which is higher than Li (2018)’s estimated market price of about RMB 75 000 (see its Figure 3), but is on par with our lower bound. The two prices are however not directly comparable for three reasons. Our model is one of a market where each one of q households offers its license in a market with $N - q$ buyers and where each of these households has a reservation price $v \sim F$. Li (2018)’s market clearing price is computed for a (counterfactual) market where all q licenses are offered to N households by auction, without a reservation price and with no transaction costs. In our model, that corresponds to a supply curve which equals zero until q and is vertical at q . In Section D of the appendix, we interpret this model as one where the market is exclusively supplied by speculators. Secondly, we consider a market with transaction costs, while Li’s market has none. In the case of no transaction costs, the implied transaction price is $p_{notc} = \text{RMB } 59$

¹²We use the bootstrapped values of \hat{s} for the estimated share of trade in Table 5 calculation for $d = 5000$.

Estimates	at $\hat{s} = 11\%$	95% CI
\hat{p}	105	[91, 121]
\hat{t}	100	[83, 118]
total transaction costs	5.7	[3.8, 7.3]
net gains from trade	1.3	[0.5, 2.5]

Table 6: Estimates of transaction prices and costs in RMB 1000. Gains from trade and transaction costs are in RMB billion.

000.

6.1 Transaction costs and net gains from trade

We compare estimates of the total transaction costs and net gains from black market trade to two benchmarks: one where non-transferability is strictly enforced and one where there are no transaction costs. The latter can be interpreted as the case with no enforcement. We follow the approach in Li closely, but where Li studied the welfare effects of the lottery, which includes congestion and pollution costs, our interest instead lies in the incentives to trade, which are unaffected by externalities. We therefore ignore externalities.¹³

We consider transaction prices to be transfers between sellers and buyers which do not affect gross gains from trade (the area between the demand and the supply curve up to $\hat{s}q$). Table 6 shows that the gross gains from trade is RMB 7.0 billion at $s = 11\%$, but transaction costs sum up to RMB 5.7 billion. The net gains from trade are the gross gains from trade minus the transaction costs. The lower bound estimate for the net gains from trade is RMB 1.3 billion, while the upper bound, for the case with no transaction costs at a black market share $q_{notc} = 62\%$ and $p_{notc} = \text{RMB } 59\,000$, is RMB 18.8 billion.

¹³Li estimates the net welfare loss of the lottery relative to an auction to be about RMB 30 billion (\$5 billion). The allocative inefficiency is the largest component, about RMB 33 billion, while the external cost savings are about RMB 3 billion. Our gains from trade concept is different from Li's welfare concept.

6.2 Further bounds on transaction costs

Under the assumptions of the equilibrium model, the data are consistent with a volume of trade that ranges from 11% to 62 % of the quota. This range implies transaction costs from zero to RMB 5.7 billion and transaction prices from RMB 59 000 to RMB 105 000. We can narrow the bounds with information on the transaction prices. Table 7 shows how the implied transaction costs and the volume of trade vary conditional on a range of known transaction prices. We can therefore tighten the bounds further if we are willing to use the news reports from Section 2 as informative about transaction prices.

Though we found a wide range of reported transaction prices, most are above RMB 70 000. It appears that rentals is the most common sales format. The most frequently reported rental price range is from RMB 500 to 1 000 per month. We may take the lower end of that range as a plausible lower bound. A 20 year NPV of RMB 6 000 in annual rental income is RMB 73 000, which we may think of as a conservative, yet plausible, lower bound for the transaction price.

Assumption 7 *RMB 73 000 is a lower bound for the transaction price.*

Table 7 shows that Assumption 7 tightens the upper bound for s that is consistent with the data from 62% to 27%, and shift the lower bound for the share of transaction costs from zero to 61%. We distinguish the realized gross gains from trade from the potential gains from trade, which are those that could be realized in the black market if there were no transaction costs. A black market without transaction costs would in our framework theoretically be as efficient as an auction. The potential gains from trade are RMB 18.8 billion, which imply that the share of the realized net gains from trade lie in the range from 7% to 28%.

Table 7 allows readers who prefer a more conservative lower bound to the transaction prices to trade weaker assumptions against wider bounds for the quantities of interest.

p	t	s	Net gains from trade	Share total transaction costs
59	0	62%	18.8	0%
57	21	49%	12.8	30%
64	42	37%	8.2	49%
73	59	27%	5.3	61%
105	100	11%	1.3	82%

Table 7: The calculations are conditional on s ranging from the estimated lower bound \hat{s} to the upper bound s_{upper} that the market can support, and assume no speculators ($z = 0$). Prices and costs are in RMB 1000. Gains from trade are in RMB billion.

Data	Assumptions	p	s	t	share transaction costs
car sales	1, 2, 3	\mathbb{R}_+	[14, 100]%	\mathbb{R}_+	[0,100]%
car sales	1, 3, 4	\mathbb{R}_+	[11, 100]%	\mathbb{R}_+	[0,100]%
car sales, F	1, 3, 4, 5, 6	[59, 105]	[11, 62]%	[0, 100]	[0,82]%
car sales, F , news reports	1, 3, 4, 5, 6, 7	[73, 105]	[11, 27]%	[59, 100]	[61,82]%

Table 8: Transaction prices and transaction costs in RMB 1000s. Assumptions 1: No anticipation, 2: No time trends, 3: No general equilibrium, 4: Equal displacement, 5: Market participants, 6: Market equilibrium, 7: Lowest plausible transaction price.

Table 8 summarizes the assumptions and the corresponding inferred bounds for the market performance measures p , s , t , and the share of gains from trade lost to transaction costs. We started with inferring a lower bound for the volume of trade using comprehensive car sales data on millions of car sales under weak assumptions that are common in the program evaluation literature. Under these assumptions, the car sales data, which are indirectly related to our quantities of interest, gave an informative lower bound for s , but did not deliver useful bounds on transaction prices and transaction costs. We then imposed increasingly stronger assumptions on black market transactions (5, 6 & 7). These assumptions are seen to be powerful in tightening the bounds, in particular Assumption 7. We note that this assumption, which serves like a highly informative prior on a variable that is directly related to our quantities of interest, is supported by only a handful of news reports and we do not know their sampling distribution.

7 Discussion

It is puzzling that the Beijing government seems to allow a substantial volume of black market trade. The government can close the black market. One reason may be that the scale of trade in the license plates makes the costs of strict enforcement prohibitive. Another reason may be that some illegal trade plausibly takes place between family members, e.g. a daughter wins a license and lets her parents register a car in her name. Such trades may be less politically expedient to enforce.

Lax enforcement may also strike a balance between allocative efficiency and equity concerns in a society where only 8% of the population approves of auctions (Li (2018)). Black market trade may be viewed as a second-best solution that ameliorates the worst misallocations resulting from the lottery allocations. More transparent mechanisms are however available. Other Chinese cities that later introduced rationing (Guangzhou, Tianjin, Hangzhou, Shenzhen, and Shijiazhuang) chose a hybrid lottery/auction mechanism which offer a different balance between efficiency and equity. Huang and Wen's estimates that the hybrid mechanism in Guangzhou, which allocates about 50% of the license plates by lottery and the other half by auction, preserves 83% efficiency. Our results give that the Beijing black market preserves at most 30%, down to as little as 10%, of efficiency.

There is historical precedent for a policy of tacitly accepting black markets. Chinn (1977) notes that during the rationing of rice in Japan following World War II, the Japanese government not only allowed a sizable black market, estimated to be about half of the total market, but also collected data on prices and quantities on illegal trade to monitor the market performance. The government price controls were set to control industry wages in a time of rapid industrialization, while the black market served as a slackness control to dampen the most severe inefficiencies.

Perhaps the enforcement is not so lenient after all. The high willingness-to-pay for license plates in the market creates strong incentives for illegal trade. If

interest is in an equitable allocation, then the size of the transaction costs (up to 82% of the gains from trade) suggests that enforcement is in fact quite effective in precluding illegal trade of high private value.¹⁴ If interest is instead in an efficient black market, the size of the transaction costs suggests relaxing enforcement.

8 Summary

Black markets may serve as an alternative allocation mechanism to auctions in places where auctions are either practically or politically infeasible. We adopted a partial identification approach in the spirit of [Manski \(2003\)](#) to analyze the performance of the black market for Beijing license plates. We estimated a set of increasingly narrow bounds for the volume of trade, the gains from trade, and the transaction costs in this market corresponding to a set of increasingly strong assumptions. In order to estimate these bounds, we developed a non-parametric difference-in-differences-like estimator using optimal transport methods. We found that the black market plausibly reallocates between 11% and 27% of the rationed license plates and conservatively estimated the net gains from trade in the black market to be between RMB 1.3 billion and RMB 5.3 billion. Yet, between 61% and 82% of the gross gains from trade are lost to transaction costs and the black market realizes between 7% and 28% of the potential gains from trade. The size of the transaction costs suggests that enforcement is effective and that the black market realizes modest efficiency gains compared to hybrid lottery/auction allocation mechanisms that are used in other large Chinese cities.

References

Athey, S. and G. W. Imbens (2006). Identification and inference in nonlinear difference-in-differences models. *Econometrica* *74*(2), 431–497. [28](#)

¹⁴Note that we use enforcement here to mean not only active enforcement, but any deterrent effect that codifying the license plates as non-transferable has.

- Bhave, A. and E. Budish (2018). Primary-market auctions for event tickets: Eliminating the rents of 'bob the broker'? Working Paper w23770, NBER. 29, 33
- Carrier, G., V. Chernozhukov, and A. Galichon (2016, 06). Vector quantile regression: An optimal transport approach. *The Annals of Statistics* 44(3), 1165–1192. 13
- Chiappori, P.-A. and B. Salanie (2016). The econometrics of matching models. *Journal of Economic Literature* 54(3), 832–861. 5
- Chinn, D. L. (1977). Staple food control and industrial development in postwar japan, 1950-1957: The role of the black market. *Journal of Development Economics* 4(2), 173 – 190. 37
- Chiong, K., A. Galichon, and M. Shum (2016, 3). Duality in dynamic discrete-choice models. *Quantitative Economics* 7(1), 83–115. 13
- Davidson, C., L. Martin, and J. D. Wilson (2007). Efficient black markets? *Journal of Public Economics* 91(7), 1575 – 1590. 2
- Drèze, J. H. (1975). Existence of an exchange equilibrium under price rigidities. *International Economic Review* 16(2), 301–320. 4
- Dye, R. A. and R. Antle (1986). Cost-minimizing welfare programs. *Journal of Public Economics* 30(2), 259 – 265. 4
- Feenstra, R. C. (1988). Quality change under trade restraints in japanese autos*. *The Quarterly Journal of Economics* 103(1), 131–146. 9
- Galichon, A. (2016). *Optimal transport methods in Economics*. Princeton University Press. 5
- Galichon, A., S. D. Kominers, and S. Weber (2018). Costly concessions: An empirical framework for matching with imperfectly transferable utility. *Journal of Political Economy* 0(ja), null. 13

- Huang, Y. and Q. Wen (2019). Auctionlottery hybrid mechanisms: Structural model and empirical analysis. *International Economic Review* 60(1), 355–385. 2, 5, 37
- Kantorovitch, L. (1958). On the translocation of masses. *Management Science* 5(1), 1–4. 12
- Leslie, P. and A. Sorensen (2013, 10). Resale and Rent-Seeking: An Application to Ticket Markets. *The Review of Economic Studies* 81(1), 266–300. 33
- Li, S. (2018, 10). Better Lucky Than Rich? Welfare Analysis of Automobile Licence Allocations in Beijing and Shanghai. *The Review of Economic Studies* 85(4), 2389–2428. 2, 3, 4, 5, 9, 16, 30, 31, 33, 34, 37, 45
- Manski, C. (2003). *Partial Identification of Probability Distributions*. New York: Springer-Verlag New York. 38
- Peyré, G. (2018). Mathematical foundations of data sciences. <https://mathematical-tours.github.io>, 267–297. 25, 27
- Schneider, F., A. Buehn, and C. E. Montenegro (2010). New estimates for the shadow economies all over the world. *International Economic Journal* 24(4), 443–461. 2, 4
- Shen, L., M. Hu, and D. C. (2019). The luck celebration hypothesis. Working paper, CUHK. 16
- Sommerfeld, M. and A. Munk (2018). Inference for empirical wasserstein distances on finite spaces. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 80(1), 219–238.
- Stahl, D. O. and M. Alexeev (1985). The influence of black markets on a queue-rationed centrally planned economy. *Journal of Economic Theory* 35(2), 234 – 250. 4
- Tan, J., J. Xiao, and X. Zhou (2019). Market equilibrium and welfare effects of a fuel tax in china: The impact of consumers’ response through driving patterns. *Journal of Environmental Economics and Management* 93, 20 – 43. 5, 11, 16

Tobin, J. (1952). A survey of the theory of rationing. *Econometrica* 20(4), 521–553.

4, 16

Xiao, J., X. Zhou, and W.-M. Hu (2017). Welfare analysis of the vehicle quota system in China. *International Economic Review* 58(2), 617–650. 3, 5, 11, 16

A Difference-in-differences regressions

We run the following difference-in-differences specifications in logs for three control groups: Tianjin, Shijiazhuang, and Tianjin and Shijiazhuang combined.

$$p_{j,c,t} = \alpha_0 + \alpha_1 \text{Beijing}_{j,c,t} + \alpha_2 \text{post}_{j,c,t} + \alpha_3 \text{Beijing}_{j,c,t} \times \text{post}_{j,c,t} + \epsilon_{j,c,t} \quad (18)$$

where $p_{j,c,t}$ is log of the j th price observation in city c , in month t , and where $\text{Beijing}_{j,c,t}$ is an indicator which is one for price observations from Beijing and zero for cities in the relevant control group. The indicator $\text{post}_{j,c,t}$ is one in all months after the introduction of the lottery in Beijing, zero otherwise. Table 9 shows that the estimated price jumps, $\hat{\alpha}_3$, are similar, and around 17% across the specifications.

Table 9: Diff-in-diff regressions in logs for different control groups.

	(1)	(2)	(3)
	Tianjin	Shijiazhuang	Both
Beijing	0.266*** (172.32)	0.233*** (116.76)	0.0829*** (74.23)
post	0.0572*** (35.30)	0.0438*** (19.97)	0.0618*** (61.75)
Beijingxpost	0.169*** (88.06)	0.182*** (75.52)	0.165*** (114.54)
Constant	-2.322*** (-1713.50)	-2.288*** (-1233.68)	-2.139*** (-2563.36)
r2	0.08	0.07	0.03
N	2116640	1795962	3214675

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

B Rules of the lottery and restrictions on driving in Beijing

The lottery eligibility criterions are given http://www.gov.cn/gzdt/2010-12/24/content_1771921.htm. To be eligible for the lottery, the applicant must satisfy at least one of the following criteria.

- Be a Beijing citizen.
- Be an army force resident in Beijing.
- Have a valid permanent residence certificate but not Beijing citizen.
- Have a valid temporary residence certificate but not Beijing citizen, and pay Social security and income tax to Beijing for five consecutive years.
- Be a citizen of Hong Kong or Macau who have lived in Beijing for more than a year.

The official restrictions are given http://www.gov.cn/gzdt/2010-12/24/content_1771991.htm .

- Cars without a Beijing plate that wants to enter with the five rings must get a “Beijing entering permit” .
- Cars that have a “BJ entering permit” can not drive within the five rings on week days between 7-9 am and 5-8 pm.
- Between 9am and 5pm, cars with a BJ entering permit that want to drive within the five rings must rotate according to the even/odd number of the last digit of the auto plate (the same as the cars with Beijing plate).

The permit lasts from two to seven days. People need to make appointment one to four days ahead of time to get the digital permit. The permits can only be renewed within three days after the current one expires, and only for five days. And the renewed ones only last for five days. Each driver can only have one permit at a time.

C Comparative statics in transaction costs and prices

We derive the comparative statics for the transaction costs and price estimators. We drop the hat-notation for expositional convenience. Taking derivatives of the equilibrium conditions in (17), we get

$$2 \frac{\partial p(s)}{\partial s} = \frac{1}{f(v_{seller})} - \frac{q}{N-q} \frac{1}{f(v_{buyer})}$$

$$2 \frac{\partial t(s)}{\partial s} = -\frac{q}{N-q} \frac{1}{f(v_{buyer})} - \frac{1}{f(v_{seller})}$$

It is immediately clear that the implied transaction costs decrease with the volume of trade s . The same is not necessarily true for the transaction price, which may increase or decrease, depending on the shape of f . However, both $\tilde{p}_{buyer} = p + t$ and $\tilde{p}_{seller} = p - t$ are monotonic in s . These statics show that, conditional on a lower bound estimate s , t and $\tilde{p}_{buyer} = p + t$ are upper bound estimates, and that $\tilde{p}_{seller} = p - t$ is a lower bound estimate.

D Extension to speculators

We now relax Assumption 5 and allow for speculators. Suppose speculators have zero willingness-to-pay for a license plate if non-transferability is strictly enforced, i.e. a speculator will never purchase a car if she wins a license. Speculators have two effects in the market: they crowd out car buyers on the supply side (shift the supply curve down) and they increase the number of car buyers on the demand side (shift the demand curve out). Suppose that a share z of the license plates are allocated to speculators, and suppose that a speculator's reservation price is zero. Then the demand and supply is

$$D(p, t) = (N - q(1 - z))(1 - F(p + t)) \tag{19}$$

$$S(p, t) = zq + \frac{(s - z)}{s} q F(p - t) \tag{20}$$

	Estimates	95% CI
\hat{p}	110	[98, 125]
\hat{t}	105	[98, 125]
Total transaction costs	6.0	[4.0, 7.9]
Net gains from trade	1.1	[0.4, 2.1]

Table 10: Estimates of transaction prices and costs at $\hat{s} = 11\%$ and $z = 0.11\%$. Prices and costs are in RMB 1000. Gains from trade are in RMB billion.

Speculators give a supply curve that is flat from zero to zq , and increasing thereafter to q . As z goes to s , the supply curve becomes vertical at q . On the demand side, the demand curve shifts out towards $v(n)$ as z goes to one and all prospective car buyers must turn to the black market for license plates.

In Table 10, we report results assuming that all 11% illegal trades are by speculators. Since the supply curve is now flat from zero to sq , all speculators sell at $\tilde{p}_{seller} = 0$. We see that now $\hat{p} = \hat{t}$. At the same time, the demand curve shifts out to $v(n)$. The outward shift in the demand curve dominates the downward shift in the supply curve: both transaction prices and costs increase, from 105 to 110 and from 100 to 105, respectively. The net gains from trade go down by 15%, largely due to further taxation of buyers.

One limit case of interest is when speculators completely crowd out prospective car buyers and there are no transaction costs, i.e. when $z = q$ and $t = 0$. Then our model gives the same allocation as Li (2018)'s counterfactual auction market, where the supply curve is vertical at q and demand is $v(n)$. It follows that Li (2018)'s analysis applies.¹⁵

¹⁵Except for the congestion and pollution externalities, which we have ignored.