Ticket to Ride:
An Auction-based Approach to Real-Time Traffic Management

George Horrell
Stanford University
ghorrell@stanford.edu

Eli Echt-Wilson
Stanford University
eliew@stanford.edu

Abstract

In this paper we pose the problem of congestion caused by ride-sharing companies as a negative externality. We introduce and analyze previous attempts to curb the externality of congestion in urban areas. We determine that existing methods improve social welfare, but do not go far enough. We discuss three different mechanisms that could be used to price access to roads, and how these mechanisms would mitigate the issue of congestion. We discuss various simulations constructed for the purpose of this paper, and how these simulations ratify the suitability of the mechanisms. We conclude that an iterative tolling method, would most accurately price usage of the road, and be the most directly applicable in the real world. Finally, we discuss various implementation strategies and difficulties that might arise with the implementation of these tolls.

1. Introduction

The economic challenge of traffic congestion in cities is an expensive problem, estimated at costing 1 trillion per year. The increasing prevalence of ride-sharing markets raises concerns of congestion, especially in larger cities. While these markets are relatively new, recent studies have concluded that ride-sharing products increase city congestion, and that increases in traffic can partly be attributed to the presence of these on-demand ride-sharing markets. We propose imposing a Pigouvian tax on ride-sharing companies to offset the negative externalities that they incur, specifically increased congestion. In this paper we will explore related work in the field of congestion management (both theory and real-world implementations) and then propose three different real time traffic management systems for ride sharing companies.

2. Externalities of Road Usage

While automobiles and associated infrastructure provide an economic boon and facilitate much of 21st century life, road travel does come with significant negative externalities. The transport sector is the primary contributor of greenhouse gas emissions. Worldwide, transport is responsible for 60% of oil usage and over 25% of energy-related CO_2 emissions. In many cities around the world, especially rapidly urbanizing large urban centres, local air pollution is an alarming problem, and is almost caused predominantly by the transport sector. Furthermore, the overuse of road infrastructure and the issue of congestion is costly and is estimated at \$1 trillion per year. [3]

Most of these negative externalities are a function of road usage; the more that road usage is controlled and decreased, the lesser the externality costs. However, the economic imperative of road travel, especially between highly populated urban centres, demands that there exists some balance between acceptable road usage and prohibition. Therefore, any attempts to curtail road usage must be sensitive to the externality that needs to be addressed, and the costs incurred by that measure/policy. There exists a range of different policies and measures that can be enacted to deal with these externalities, including: fiscal policies such as taxation and subsidies, regulatory policies such as fuel and emissions standards, and planning and investment such as infrastructural expenditure. [3]

As shown, there are multiple externalities that we could seek to correct, however it is unlikely that we will be able to correct all of the issues with the transport sector within this paper, or with a single solution. As a result, we propose to only address the issue of congestion in urban areas, and will leave aside issues of pollution, both local and global. However, it is evident that a solution that mitigates congestion will also pose benefits with respect to pollution. Spreading the automotive load around the transport network will improve the issue of local pollution, preventing certain areas from experiencing dramatically bad air quality due to nearby congested roads. Furthermore, journeys that are completed quickly, without waiting in traffic will be more fuel efficient, and will thereby reduce global pollution. Therefore, we will only seek to address the issue of congestion in this paper.

2.1. Pigouvian Taxes for Controlling Externalities

One of the most well-established and well-researched approaches towards controlling the externalities of human behavior is a Pigouvian tax. A Pigouvian tax is imposed on the individuals / companies who are responsible for the externalities. In 1972, Baumol showed that for public goods (such as road access), a Pigouvian tax alone can create Pareto optimality in the output of a system. To achieve Pareto optimality, the tax imposed must be exactly equivalent to the social marginal cost it imposes on others. Notably, however, he concluded that we can achieve an optimal outcome without taxes on any other party, and without compensation to the industry that suffers damage. [2]

Other related work explores Pigouvian taxes in the more specific context of routing in networks. Roughgarden defines a system of marginal cost pricing in nonatomic selfish routing games, which is a satisfactory model for road networks. A nonatomic selfish routing game is a game with k players trying to traverse from one location to another in a network. An instance of a game with taxes for traversing each edge is defined by:

$$(G, r, c + \tau)$$

where G is a set of vertices and edges (V, E) and an accompanying set of k source-sink pairs (s_i, t_i), one for each player representing their start location and end destination in the network. The routes chosen by players in the game are denoted by a flow f. r_i denotes the amount of traffic identified with each source-sink pair, and c is a cost function that specifies a cost for traversing each edge in the network. Finally, each edge e possesses a non-negative tax τ_e. An optimal model would tax each player on each edge the amount that its presence on the edge costs the other users of that edge. The principle of marginal cost pricing says that the optimal tax is given by:

$$\tau_e = c'_e(f_e) \cdot f(e)$$

where c'_e is the derivative of c_e, $c'_e(f_e)$ is the marginal increase in cost caused by one user, and f_e is the total amount of traffic that suffers this marginal cost increase. Note that this computation of tax requires that we have well-defined, differentiable cost functions defined for every edge, and does not impose explicit edge capacities on the edges of the network. [8]

In the absence of well-defined cost functions or in the case where we want to explicitly impose a capacity on edges in the network, Roughgarden shows that we can instead frame the problem as a resource allocation game, where we want allocate a single divisible resource (access to a road, for example) to a number of competing players. The game is defined by a capacity $C > 0$, and utility functions for each of the n players of the game. An outcome is an allocation vector $(x_1, ..., x_n)$ such that $\sum_i x_i = C$, where x_i is the amount of the divisible resource allocated to player i. We can frame the road allocation problem in this format, with a road having some capacity C denoting the number of cars allowed to drive on it during a designated time period. Companies wishing to gain access to the road have some utility function specifying the utility gained from having access to the road for a number of cars. The proportional sharing mechanism in this model allocates resources proportional to the bids that players submit. Formally, each player i submits a bid b_i, and player i receives:

$$x_i = \frac{b_i}{\sum_j b_j} \cdot C$$

and each player is charged their bid b_i. In our case, companies would bid for access to a road and in return gain access for a certain number of rides that can use that road. Unlike the mechanism of marginal cost pricing above, this mechanism does not require explicitly defined cost functions, and allows some centralized entity (in our case the government) to specify a capacity and restrict access to the resource of interest rather explicitly. [8]

2.2. Attempts to Curb Negative Externalities

In addition to the theory of congestion prices, we looked at current state-of-the-art implementations of congestion prices to inform our analysis. Congestion prices have been implemented in various places, but the most successful implementations are considered to be in Singapore and London.

In London, congestion prices have been used since 2003 to reduce traffic and simultaneously raise money to improve transportation systems. It is the first example of congestion pricing in a major European city, and suggests that congestion zoning can be implemented successfully in other areas as well. With a few exemptions, motorists driving in central London between the hours of 7:00AM and 6:30PM must pay a flat tax of 8 pounds (increased from 5 in 2005). Residents of the area receive a 90% discount, but anyone else pays a flat daily rate. We notice that the system is far from optimal: the tax is applied equally to all drivers regardless of how long or far they are driving in the applicable area and does not vary with location inside the area or time of day. Interestingly, even this simple system costs approximately 100 million pounds to operate each year, with a yearly revenue of 160 million pounds. However, the system has also proven to be quite effective in numerous ways. It raised the average traffic speed by 37%, and also taxi travel costs dropped significantly because they could run more efficient routes and charge less per ride [5].

Singapore has similarly been implementing Pigouvian congestion taxes since 1975. In the beginning, they implemented a manual pricing system very similar to the system described in London in the previous paragraph: a flat tax (which reduced morning traffic volume by 45% and overall traffic volume by 25%). In 1998, they abandoned the manual pricing system in favor of a newly developed Electronic

Figure 1: Singapore ERP Barrier

Road Pricing system, or ERP. The ERP, much like the system in London, charges a toll per entry for congested zones in Singapore. Unlike London, however, the exact toll is dependent on a variety of factors including time of day, exact roads being taken, and vehicle size. The tolls are updated quarterly (every 3 months), allowing drivers to reevaluate their travel decisions. This varied price system was designed to achieve Pareto optimality in road allocation, and has continually expanded to more locations in Singapore as time passes. It has produced a reduction of traffic volume of approximately 10-15%. While this is notably smaller than the reduction gained from the manual pricing system, it is important to note that the ERP generates lower tolls for the drivers and citizens of Singapore than the manual pricing system, and therefore does a better job of optimizing overall welfare. The system requires the installation of cash cards into all driving vehicles, and then construction of large gantry scanning devices which automatically charge the cash cards when a car drives through the gantry. The initial cost of building the system was approximately S$200 million, with a yearly cost of S$16 million per year. It brings in S$80 million dollars of revenue annually for Singapore. [6, 9, 7, 4]

3. Modelling Transport Networks

We propose three different mechanisms for imposing a Pigouvian tax to correct congestion in the road network, including two auctions and a tolling mechanisms.

We first consider the simplified model of a single road. In this model, the capacity of the single divisible resource of a road $C > 0$ is to be allocated between $n > 1$ bidders, representing ride-sharing companies. Each bidder i possesses a concave, strictly increasing and continuously differentiable utility function U_i. We consider a resource allocation game, defined by the utility functions and the road capacity. The outcome of the game is an allocation x_i of the total capacity C for each bidder, such that:

$$\sum_i x_i = C$$

We are interested in maximizing the utility of all the bidders $\sum_i U_i(x_i)$, and ensuring that all bidders are incentivized to bid truthfully.

3.1. Proportional Sharing

As discussed above, one way of distributing access to the road is via proportional sharing. We will consider an ascending proportional sharing auction (APSA) motivated by these design criteria. For the purpose of this auction, we will consider that access to the road for a given time period t can be considered as an infinitely divisible finite resource C_t, with n bidders.

The format of this auction is as follows:

Ascending Proportional Sharing Auction

while(TRUE):

- Accept open bids from each bidder i, such that the proportion of C_t assigned to i is:

$$\frac{b_i}{\sum_{j=1}^{n} b_j}$$

- The bidder can also choose to leave the auction at this point, thereby bidding 0 and receiving no proportion of capacity. These bidders will not be polled again for subsequent bids, and their proportion will be distributed among the remaining bidders.
- The bidder can also choose to bid the same as the previous round. If in one round, all bidders choose to maintain their bid, the auction ends.

This mechanism has appreciable advantages. It is a simple mechanism, with easy participation. Unlike some auction mechanisms, bidders aren't required to specify their entire demand curve. To show that this auction would produce reasonable prices if utilized in the real world, we constructed a simulation of a road auction with utility-maximizing agents. We considered each agent i as possessing a demand for d_i capacity on the road, and a valuation of q_i per unit capacity. The valuation of agents is linear, such that their valuation for x units is xq_i. Agents use the BFGS optimization algorithm to maximize their value over the following function:

$$v_i(b) = \min\left(\frac{b}{\sum_{j=0}^{n} b_j} \cdot C_t, d_i\right) \cdot q_i - b$$

Below, we ran a simulation of two agents on a unit capacity road, each with a demand for a single unit, at a valuation of one unit.

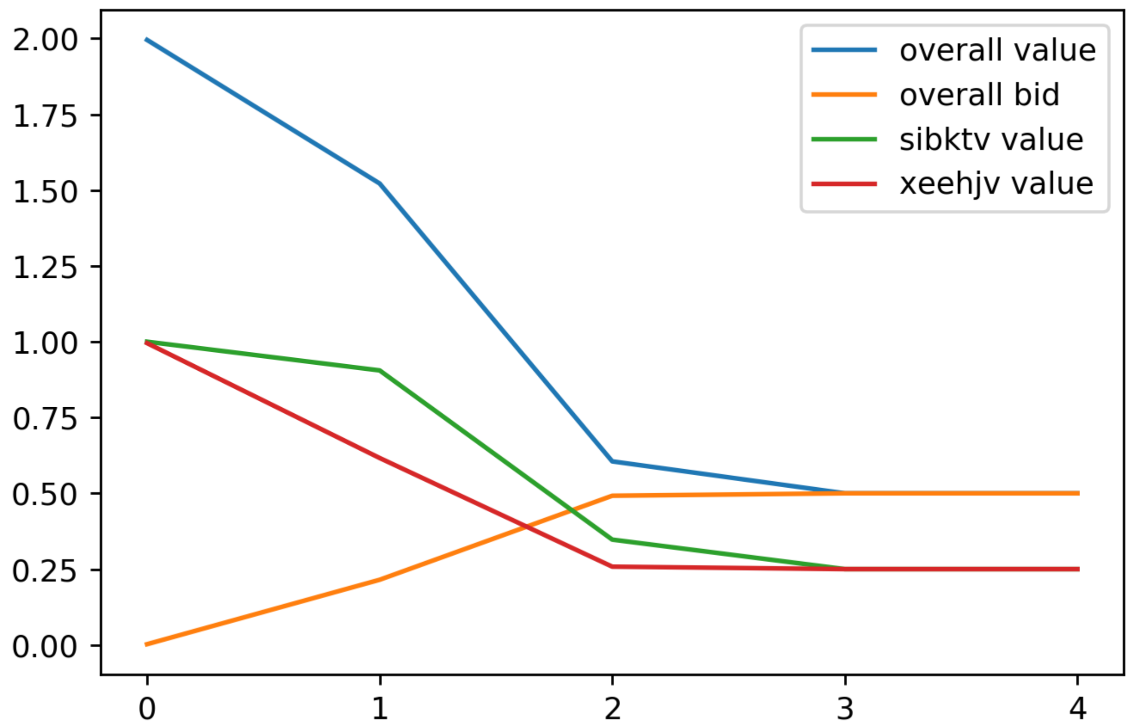

As shown in the plot above, both agents reach an equilibria of bids $(0.25, 0.25)$, where neither bidder can alter their bids without lowering their value. The simulation shows that similar equilibra are reached for more bidders, and arbitrary valuations. The plot below shows six bidders converging to a result where all but two of the participants have dropped out, with a range of randomly generated valuations and demands. This result is reasonable, given that there is far more demand for capacity than there is supply.

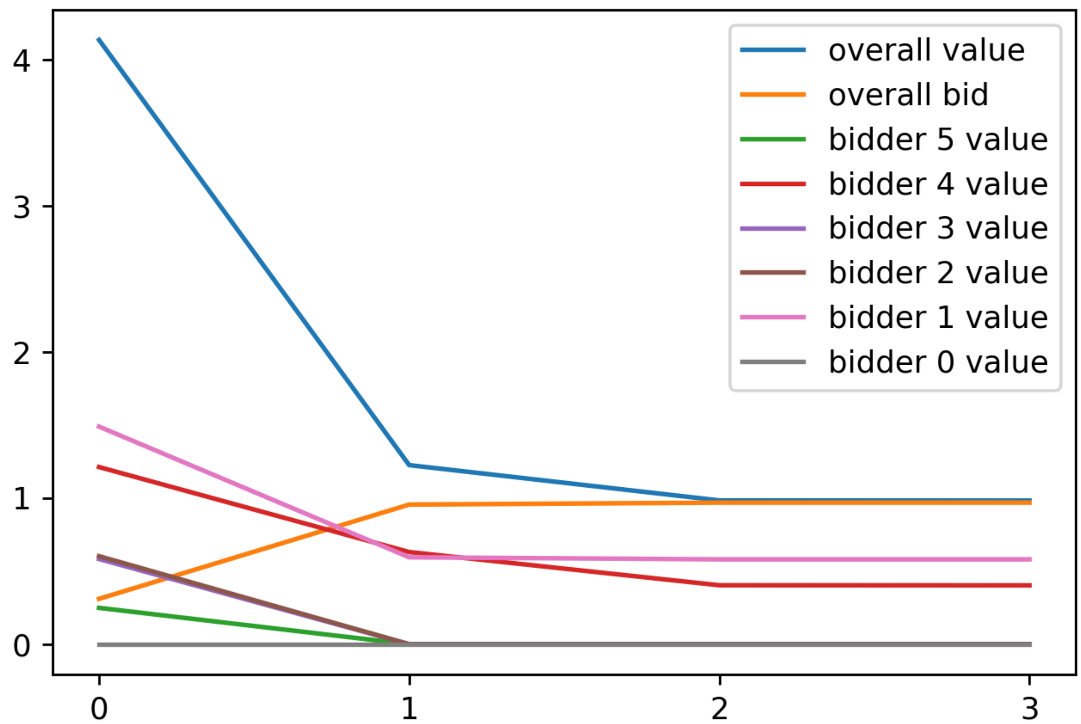

This auction is revenue-maximizing on behalf of the auctioneer (government). However, this revenue-maximizing property is sub-optimal. Without some real investment, companies would not accurately report their valuations for road usage, and spots could not be distributed evenly, but that does not imply that the government should seek to maximize revenue. Since this is a new tax being levied upon ride-sharing companies, it is probably in the best interests of the government to minimize revenue extracted, to ensure support for the new measure. We do not factor the revenue raised by the toll in our calculation of social welfare.

This mechanism is also not truthful. Bidders can "bully" other bidders out of the auction. If they raise their bidding price above their true valuation to the point where another bidder has a negative value for any bid that they make, then that bidder will be forced to drop out of the auction. Their proportion of C_t is then distributed among the remaining bidders, and so the bully bidder gains value. While this behaviour could be prevented by forcing bidders to remain in the auction, this would lead to non-CE results, as some bidders lose value in the auction. Therefore, we see that this auction format incentivizes a non-truthful strategy.

Since this mechanism neither fulfills our requirement for truthfulness, or maximum social utility, we must consider other mechanisms.

3.2. Multi-unit Auction

We can consider the resource of road access in another manner - as a set of n identical "spots for vehicles" on the road for a set time period. In this formulation, the government auctions off the n spots to bidders (ride-sharing companies). Each bidder i has a private marginal valuation function $\mu_i(j)$ for the jth spot.

For this multi-unit auction, a traditional starting point mechanism is the "uniform-price" auction. Consider an auction for k units and n bidders:

Uniform Price Auction

1. Accept sealed bids from each bidder i, where each bid contains a price b she is willing to pay for quantity of units $D(b)$. Relabel the bids so that $b_1 \geq b_2 \geq b_3 \geq \ldots \geq b_n$.

2. Assign the first bidder to the first $D(b_1)$th spots, the second bidder to spots $D(b_1)+1, \ldots, D(b_1)+D(b_2)$. Continue to assign spots in this manner until k spots have been assigned. Denote the set of bidders assigned spots as P.

3. Charge each bidder $i \in P$ the clearing price times q_i.

However, with this auction format, bidders are heavily incentivized to perform demand reduction, which leads to non-truthful bidding, and it is therefore not suitable for our road access auction.

While this format is un-ideal, there are other multi-unit auction mechanisms that are more suitable. We can make the further assertion that valuations are downward-sloping, such that $\mu_i(1) \geq \mu_i(2) \geq \ldots \geq \mu_i(m)$. Since the valuations of individual rides are independent and ride-sharing companies would prioritize filling the first spot with their highest value ride, and the second spot with their second highest value ride, and so on, this is a reasonable assumption to make. Given this assertion, it would be preferable for our mechanism to simulate the allocation of the VCG mechanism, thereby maximizing welfare [8]. We consider the clinching auction (which simulates the VCG allocation), with k units and n bidders [1]:

Clinching Auction

1. Initialize $p = 0$.

2. while (TRUE):

   (a) Ask each bidder for the number of units they would buy at price p: $D_i(p)$.

   (b) If $\sum_{i=1}^{n} D_i(p) \leq k$, then halt the auction and allocate units and prices as detailed below.

   (c) Otherwise, increment p by ϵ.

Allocation of units

If p is the price at termination: allocate $x_i \in [D_i(p), D_i(p-\epsilon)]$ units to bidder i. Ensure that $\sum_{i=1}^{n} x_i = k$.

Allocation of prices

For $j = 1, \ldots, k$, allocate the price of each unit as follows:

$$q_i(j) = -\epsilon + \text{min}_{t \in \mathbb{Z}+} \left\{ \epsilon t : \sum_{s \neq i} D_s(\epsilon t) \leq s - j \right\}$$

This mechanism has the attractive properties of being both truthful and social welfare maximizing. However, as with all the mechanisms discussed thus far, this approach requires discretizing the road resource, and placing a hard cap on the number of cars allowed on the road for a time period. This isn't optimal for the following reasons:

1. Calculating the ideal number of cars is challenging, and would vary based on road conditions (rain, snow), time of day, time of year, etc...

2. There is an advantage to letting ride-sharing companies bid on road access in advance. This would let companies more accurately plan how they will route their drivers. However, whatever period of time in advance you let companies bid on access, is a lag in reaction to these changing road conditions.

3. Perhaps the most significant issue is other drivers on the road, such as individuals, service vehicles, freight, etc... It is unclear how this system would accommodate for these drivers. If we imagine a total capacity of T for the road, and a predicted demand of individuals I, then the proportion allocated for auctioning to ride-sharing companies $C = T - I$. However, if the number of individual drivers unexpectedly swells, then congestion would persist and ride-sharing companies would be paying an unfair tax.

In light of these disadvantages, we should consider a mechanism that permits an arbitrary number of drivers.

3.3. Iterative Toll Updating

While the previous two proposed mechanisms have a solid foundation in theory, we also consider the potential benefits of a mechanism that instead leverages sophisticated existing technologies to achieve the desired outcome. In both of the previous mechanisms, we are forced to discretize our resource (access to a road) and define an explicit limit on the amount of the resource we are willing to dole out (the capacity C for in the proportional sharing mechanism and the n identical spots auctioned in clinching auction). This can be difficult to do in reality (see above). As an alternative, we consider a mechanism which imposes a toll on the road and updates the toll continuously based on previous data in order to achieve a desired traffic flow.

For this mechanism we begin by representing time as a repeated sequence of time intervals of size n: $t_1, ..., t_n, t_1, ..., t_n, ...$ For example, we could divide a week into $n = 168$ single-hour time intervals, and assume similarity between the same time intervals in two different weeks (assume that 5-6PM on Monday this week will look similar to 5-6PM on Monday next week). The road has a toll c_i which is the cost charged for entry to the road at time interval t_i. There is also a demand curve $d(c_i)$, which specifies the number of people seeking entry to the road at time interval t_i given toll c_i. Lastly, we assume a throughput function for the road $f(d)$ which gives the estimated speed of traffic as a fraction of the maximum speed based on demand for the road (number of people seeking entry). The government has some desired throughput $f*$.

We define the mechanism as follows:

1. Initialize tolls $c_1, ...c_n$ arbitrarily to some value > 0.

2. After initialization, at every time step t_i compute $f(d(c_i))$, and perform the following update for c_i (for the next time we hit t_i):

$$c_i = c_i \cdot \begin{cases} 1+\alpha & f(d(c_i)) < f* \\ 1-\alpha & \text{otherwise} \end{cases}$$

where α is the learning rate of the mechanism.

The mechanism effectively performs iterative updates of the toll (for a time period t_i) based on whether the previous toll successfully reduced traffic to the desired level for t_i. Over time, the tolls will converge toward values that produce the desired traffic flow rate $f*$, eventually oscillating around the optimal value.

The challenge associated with this mechanism is approximating the throughput function f. The throughput of a given road (given demand) may be a complicated function of a variety of factors: the number of lanes, number of traffic lights / stop signs, etc. However, we wish to treat the throughput function as a black box which simply returns some fraction of the maximum speed on the road. Existing technologies such as Google and Apple maps have leveraged large amounts of data available to them to approximate throughput in this way. Google maps shows a road in green if it is moving at or near the optimal speed, yellow if there is some traffic, and red if there is significant traffic. We propose using these available technologies to approximate our throughput function. This would amount to imposing a toll c_i, observing the resulting traffic rate using some existing technology, and then updating the toll according to whether the traffic rate is above or below our desired rate. This system has a number of advantages / disadvantages worth considering:

1. The system does not require a centralized agency to determine a finite amount of road access to distribute. It instead relies on continually updating the tolls and allowing the market to resolve itself over time. This would be far more convenient for the government agency imposing the toll.

2. The system assumes similarity between different instances of the same time interval. While this assumption should generally hold (assuming that the time intervals are laid out reasonably), it fails to account for special circumstances that may periodically cause a significant increase / decrease in traffic, such as accidents, holidays where usage peaks more than usual, etc. The mechanism will likely fail to set an appropriate toll when circumstances like this occur because it can't update quickly enough.

3. There may be times where, regardless of ride-sharing companies and their usage of roads, other entities using the roads cause the traffic speed to be worse than the desired rate (especially in concentrated urban areas such as Los Angeles). During these time intervals, the tolls will continually increase without bound, because tolling the ride-sharing companies won't be able to improve traffic to the desired state. Ride-sharing companies would be effectively pushed out of the market by ridiculously high tolls. This issue would have to be mitigated by a larger congestion management system that didn't only apply to ride-sharing companies.

3.3.1 Iterative Toll Mechanism Simulation

We explored the viability of this approach in a simulation with synthetic data, to illustrate that tolls eventually converge towards optimal levels. We used 24 single-hour time intervals and on each hour of each day sampled a total population demand for the road from a normal distribution. The mean and standard deviation of population demand for the same time interval on two different days were equivalent in order to achieved the assumed similarity (with some noise). We assumed a demand curve which specified a percentage of the population demand that would still want to use the road as a function of the current toll, and a throughput function which estimates traffic speed (a proportion of maximum) as a function of number of people using the road. Finally, we initialized toll values for the 24 time intervals and simulated the mechanism, finding that the tolls relatively quickly converged on and then oscillated around their optimal value.

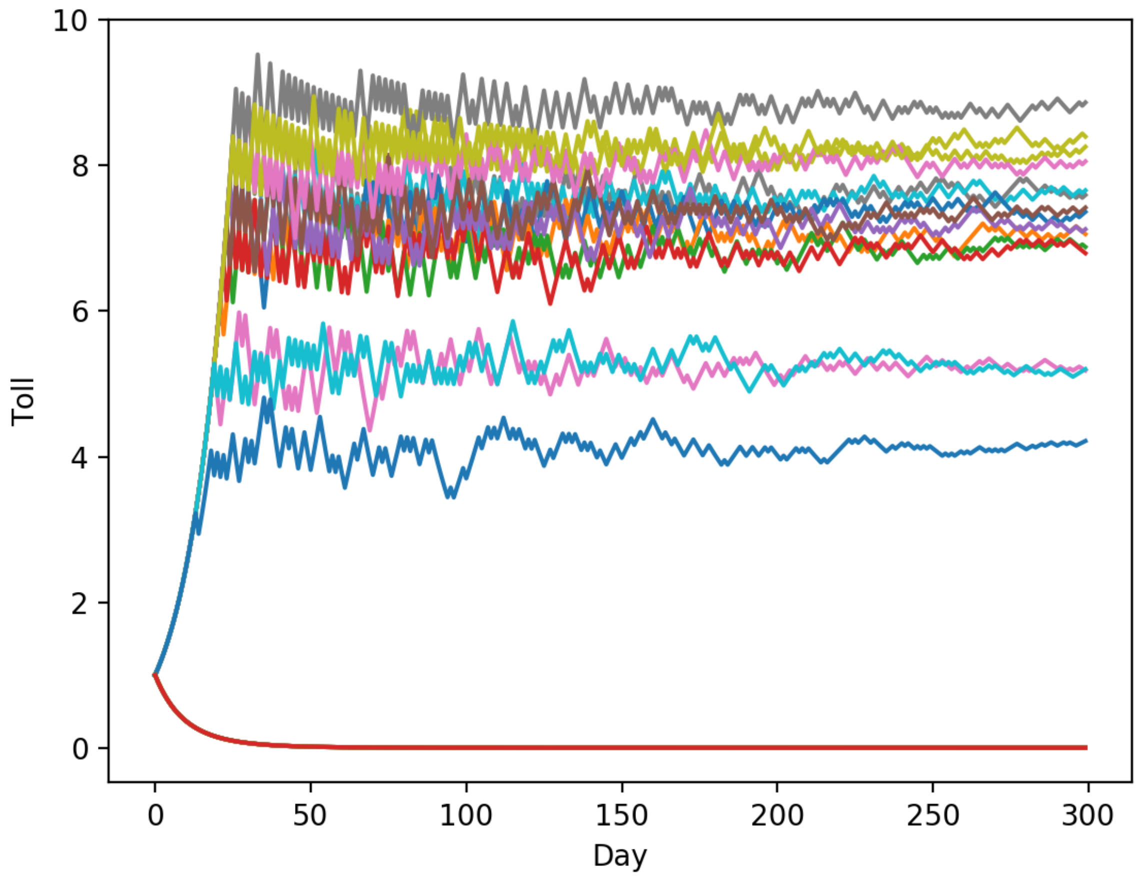

While the simulation is relatively naive and makes a good number of simplifying assumptions, it serves as a baseline proof of concept for the iterative toll generating mechanism.

3.4. Network Model

While we chose not to formally define a model that operates over larger networks of roads for this paper, we include considerations of how these different single-road mechanisms would translate to a larger scale.

One advantage of the iterative toll mechanism on a larger scale is that it self adjusts, not requiring the government to specify a capacity for every road in the network. Doing so would be difficult, especially if the capacities are changing with time.

On the other hand, the auction-based systems put less accountability on the shoulders of the government in terms of imposing prices, because the prices are determined by the companies bids rather than the government itself. In the case of a larger network, companies would be at liberty to strategize and route with respect to predicted tolls. We believe that the market of ride-sharing companies would select for success in best handling these tolls.

4. Additional Considerations

There are a variety of logistical details that we must consider in looking at how these mechanisms would be implemented in the real world.

4.1. Enforcement

A critical consideration to address is the question of how road access will be enforced. We suggest enforcing road access similarly to the system used in Singapore, with cash cards installed in any registered ride-sharing vehicle and scanning devices installed on roads that the government wants to limit access to. This would require a large initial investment in the system, but the success of the systems in Singapore and London suggest that the tax income generated from the system will eventually pay off that investment. [7]

For mechanisms where the government only wants to allow companies a finite number of rides on a particular road, violations of a companies allotted limit would penalize the company by charging them a fine much larger than the amount they paid per road access.

For the tolling mechanism, enforcement can work almost identically to the way it works in Singapore, except for the detail that cash cards would only need to be installed in ride-sharing vehicles.

4.2. Cost Distribution

We assume that the cost of our proposed systems will be distributed between the ride-sharing companies and the users. Ride-sharing companies will likely be forced to increase prices (especially during times of high traffic), but also will try to keep costs as low as they can to maintain high demand for their product. Ultimately, that decision will be left in the hands of the ride-sharing companies, as they are free to set prices as they wish. They can either transfer the cost of the system entirely to the users (increasing prices by the additional amount they pay), or absorb a portion of the cost by keeping prices as low as they can while making a profit. In either case, we would predict that costs will increase and some users will be incentivized to switch to a more affordable option such as public transportation.

5. Conclusion

In this paper we analyzed multiple different approaches towards controlling congestion through a Pigouvian tax on ride-sharing companies. Two of the mechanisms are auction-based mechanisms motivated by existing theory and related work in the fields of congestion management and resource allocation. We find that, while these mechanisms have nice theoretical properties, there are implementation details that may be difficult to achieve in reality. The third proposed mechanism aims to make congestion management more feasible by leveraging existing technologies to measure levels of traffic and iteratively update road tolls based on whether the traffic levels are acceptable or not. We used simulations to show basic proof-of-concept for two of the proposed mechanisms. For all of the mechanisms, we notice that it may be infeasible to achieve desired levels of traffic by only imposing taxes on ride-sharing companies, for traffic has existed long before the presence of ride-sharing companies in the market. That said, the proposed mechanisms can either be extended to coexist alongside other congestion management systems (such as those used in Singapore and London), and the iterative toll update system could be extended to apply to non ride-sharing entities as well.

References

[1] L. M. Ausubel. An efficient ascending-bid auction for multiple objects. *American Economic Review*, 94(5):1452–1475, 2004.

[2] W. J. Baumol. On taxation and the control of externalities. *The American Economic Review*, 62(3):307–322, 1972.

[3] H. Dulal and G. Timilsina. Urban road transportation externalities: Costs and choice of policy instruments. *World Bank Research Observer*, 26:162–191, 02 2010.

[4] M. Goh. Congestion management and electronic road pricing in singapore. *Journal of transport geography*, 10(1):29–38, 2002.

[5] T. Litman. London congestion pricing–implications for other cities. *CESifo DICE Report*, 3(3):17–21, 2005.

[6] S.-Y. Phang and R. S. Toh. From manual to electronic road congestion pricing: The singapore experience and experiment. *Transportation Research Part E: Logistics and Transportation Review*, 33(2):97 – 106, 1997.

[7] S.-Y. PHANG and R. S. TOH. Road congestion pricing in singapore: 1975 to 2003. *Transportation Journal*, 43(2):16–25, 2004.

[8] T. Roughgarden. Algorithmic game theory. *Commun. ACM*, 53(7):78–86, July 2010.

[9] L. Xie and P. Olszewski. Modelling the effects of road pricing on traffic using erp traffic data. *Transportation Research Part A: Policy and Practice*, 45(6):512 – 522, 2011.