Dynamic Taxi Pricing

Cheng Zeng[1] and **Nir Oren**[2]

Abstract. Taxi journeys are usually priced according to the distance covered and time taken for the trip. Such a fixed cost strategy is simple to understand, but does not take into account the likelihood that a taxi can pick up additional passengers at the original passenger's destination. In this paper we investigate dynamic taxi pricing strategies. By using domain knowledge, such strategies discount trips to locations containing many potential passengers, and increase fares to those areas with few potential passengers. Identifying a closed form optimal dynamic pricing strategy is difficult, and by representing the domain as an MDP, we can identify an optimal strategy for specific domains. We empirically compare such dynamic pricing strategies with fixed cost strategies, and suggest future extensions to this work.

1 Introduction

Taxi fares are commonly priced based on the distance and time taken for a trip. This model is near-universal, and normally regulated by a licensing authority. However, and particularly for long journeys, a driver will often negotiate the price for a journey with a potential passenger while ignoring the official pricing rules. Literature in the field of economics has examined this phenomenon (e.g. [1, 3]), and suggests that in negotiating a price, the driver considers not only their own costs, but also the likelihood that another passenger would require a ride near the original passenger's destination, providing the driver with more work — and additional income — at little cost. Unlike the formal pricing structure, where prices are set in advance, such negotiations appear to offer the driver additional utility. In this paper, we empirically investigate how the best price for such a pricing strategy (which we call *dynamic taxi pricing*) can be determined.

Our hypothesis is that the dynamic pricing strategy will provide a driver with increased utility when compared to the the regulated approach. We also compare the effects of the two techniques on passengers. We describe our model of the system in the next section, and undertake an empirical evaluation of it in Section 3. Related and future work are discussed in Section 4.

2 The Model

We consider a simple system containing only a single taxi, which can transport a single passenger at a time from their origin to a destination. The environment in which taxis and passengers exist is a weighted graph $G = (N, E, W)$, where the set of nodes N identify points at which passengers can be picked up and dropped off; edges $E : N \times N$ identify routes between pick-up and drop-off points; and weights $W : E \to \mathcal{Z}$ identify the time and cost taken to travel between nodes. We denote the shortest distance between two nodes i, j as $|i, j|$.

We adopt a discrete time representation. During each time interval, there is a likelihood that a passenger will appear at node n_i and desire to travel to some other node j, written as p_i^j. We require that for any $i \in N, \sum_{j \in N} p_i^j \leq 1$. We assume that a node will only generate a single passenger at any point in time, and that a passenger who is not immediately transported will not remain following the current time step.

The price a taxi will suggest for a journey is clearly dependent on the passenger's willingness to accept this price. We therefore model passengers via a continuous *demand function* $df : \mathbb{Z} \times \mathbb{Z} \to [0, 1]$. This function takes the proposed price and distance between nodes, and returns the probability that a passenger will accept the proposed price. We assume that this function is monotonically decreasing with increasing price for the same distance, and that all passengers have an identical demand function. We also introduce a reflexive cost function $c(i, j)$ for travelling between nodes.

During each time interval, a taxi can take one of two actions — move; or bid-and-move. The former allows the taxi to travel to another node, while the latter allows the taxi to propose a price to a passenger and then travel to another node. If the passenger accepts the taxi's proposal, it will board the taxi, and the taxi will start moving to its destination, otherwise, it will move to its original target node.

We identify an optimal pricing and movement strategy for our taxi by modelling the system as an Markov Decision Process (MDP). To do so, we discretise our action space, making the problem amenable to off-the-shelf value or policy iteration solvers. This discretisation is done by noting that distances are bounded (and integer), and by allowing only integers to be bid as prices for trips, again bounding these by an upper value (we assume that there is some finite maximum price for a trip that a passenger is willing to pay). We also implicitly assume that our graph contains a finite number of nodes.

Our MDP is generated such that if a node i within our system has n non-zero p_i^j, then it will be associated with $n + 1$ *normal* states. Each such state represents the presence of a passenger desiring to travel from i to j, with an additional state representing no passengers being present at this node. We label each such state s_i^j with $s_i^\emptyset$ representing the "no passenger present" state. We also create a set of additional *reward states* associated with each possible price - a price of x

for moving from i to j will lead to a state s_i^{jx}.

We consider two types of actions - *move* actions, and *bid and move* actions. A move action allows a taxi to move to a neighbouring node (or remain at the current node), and transitions from a state s_i^j to a state s_j^k with probability p_j^k. Transitioning to $s_j^{\emptyset}$ occurs with probability $1 - \sum_k p_j^k$. The reward for this transition is equal to $-c(i,j)$. A bid and move action can be informally interpreted as "bid x for the journey, if successful, transition to the passenger's destination node n_d and drop them off, otherwise, move to node j". Such an action transitions from s_i^d to s_i^{dx} with a probability $df(x, |i,d|)$, and to a state s_j^k with probability $(1 - df(x|i,d))p_j^k$. It transitions to a state $s_j^{\emptyset}$ with probability $(1 - df(x|i,d))(1 - \sum_k p_j^k)$. The reward for the transition to a reward state s_i^{dx} is $x - c(n_i, n_d)$, while the reward for a transition to state s_j^k is $-c(i,j)$.

Finally, we create a *virtual action*, enabling an agent to move from a reward state to a normal state. This is the only action possible in a reward state. Starting in a reward state s_i^{jx}, this action leads to a normal state s_j^k with probability p_j^k, and to $s_j^{\emptyset}$ with probability $1 - \sum_k p_j^k$. The reward for moving between these two states is 0. Since reward states capture the utility gain of our agent, such virtual actions are needed to return the MDP to a normal state.

With regards to size, in the worst case, there are $O(n^2 + b^2)$ states in the MDP, where n is the number of nodes and b is the number of possible bids. In such a system, there are $O(n^4)$ edges. This case occurs when p_i^j is non-zero for all nodes, and the graph is fully connected.

We can now use standard finite MDP solving techniques to identify an optimal strategy for a taxi given a specific road network, demand function and cost function.
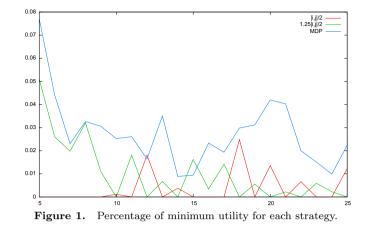
3　Evaluation

The following table summarises the experimental parameters used in our empirical evaluation.

Parameter	Value	Parameter	Value
Discount	0.1	$df(x, |i,j|)$	$1 - x/|i,j|$
$c(|i,j|)$	$1 + |i,j|/10$	Distance range	[5,15]
β	0.3	k	3

Note that the maximum expected utility for travel between nodes i and j for the given df occurs at a bid value of $|i,j|/2$. Therefore, having constructed our system, our MDP considers only (integer) bids ranging from $\lfloor 0.75|i,j|/2 \rfloor$ to $|i,j|$, and performed policy iteration over this MDP to identify an optimal strategy. The results of this strategy were compared with two fixed bid models, which bid $|i,j|/2$ and $1.25|i,j|/2$ respectively. Each experiment was averaged over 5 runs, with experiments being run for 5 to 25 nodes in the graph respectively. Fig. 1 plots the utility of each strategy as the percentage improvement over the worst performing strategy. The use of a dynamic strategy yields small improvements over the optimal fixed price approach in this situation.

Turning our attention to the number of passengers successfully transported, the dynamic approach results in around 10% more passengers being transported than the best fixed price approach. In turn, the lowest utility fixed strategy transports around 10% more passengers than the dynamic approach. The dynamic approach therefore appears to not only maximise taxi utility, but still perform well in terms of passen-

Figure 1.　Percentage of minimum utility for each strategy.

gers transported when compared to the highest utility fixed price approach.

4　Discussion and Conclusions

The dynamic approach yielded only a small improvement over fixed price approaches. However, these improvements were consistent, and were obtained with a (relatively) small discounting value, meaning that the taxi did not pay much attention to future rewards. Surprisingly, we found that when this discounting value was increased, the strategy performed worse, and we believe that this is an artefact of our domain. However, as future work, we intend to investigate this issue more closely. Furthermore, we intend to increase the complexity of our domain by adding passenger queues; considering multi-taxi scenarios; situations where a taxi can pick up more than one passenger at a time; and examine how an MDP model can affect passenger pricing structures. We intend to investigate what formal guarantees with regards to performance can be given. Finally, we will investigate how our proposed pricing structure interacts with subsidies in realistic domains.

There has been surprisingly little work on the topic of dynamic taxi pricing strategies. [2] examined a time based pricing strategy to incentivise drivers to work during peak times, and [4] gives an overview of taxi models from the transport literature. Finally, [5] consider non-linear fare structures which bear passing similarity to the structures learned by our model. Our work departs from these techniques by proposing a pricing structure based on optimisation techniques first studied in Artificial Intelligence.

REFERENCES

[1] H.S. Farber, 'Is tomorrow another day? The labor supply of New York City cabdrivers', *J. Political Economy*, **113**(1), 46–82, (2005).
[2] J. Gan, B. An, H. Wang, X. Sun, and Z. Shi, 'Optimal pricing for improving efficiency of taxi systems', in *IJCAI*, (2013).
[3] D. Orr, 'The "taxicab problem": A proposed solution', *J. Political Economy*, **77**(1), 141–47, (1969).
[4] J. Salanova, M. Estrada, G. Aifadopoulou, and E. Mitsakis, 'A review of the modeling of taxi services', *Procedia - Social and Behavioral Sciences*, **20**(0), 150 – 161, (2011).
[5] H Yang, C. S. Fung, K. I. Wong, and S. C. Wong, 'Nonlinear pricing of taxi services', *Transportation Research Part A*, **44**, 337–348, (2010).