INTERNATIONAL ECONOMIC REVIEW
Vol. 44, No. 2, May 2003

AN ANALYSIS OF THE MARKET FOR TAXICAB RIDES IN NEW YORK CITY*

BY RICARDO LAGOS[1]

New York University

In the last few years, the city of New York has increased taxicab fares and relaxed a 59-year-old cap on the number of licenses. This article uses a dynamic equilibrium model of meeting frictions to quantify the impact of these policies on medallion prices and on the process that rules the meetings between passengers and taxicabs in New York City.

1. INTRODUCTION

As in most U.S. cities, the taxicab market is regulated in New York City: Both fares and the number of licences are set by the regulating agency. Regulatory changes and their effects on the taxicab market have often attracted the attention of economists. In order to solve an exercise posed in the Appendix to Milton Friedman's (1962) *Price Theory: A Provisional Text*, Orr (1969) theoretically studied the impact that increasing fares and relaxing entry would have on medallion prices.[2] De Vany (1975) contrasted the differences in output, capacity, and utilization of capacity when the market is organized as a medallion system and when there is free entry. Schroeter (1983) presented a model of taxicab service in a regulated market where auto dispatch and airport cabstand are the primary modes of operation, and empirically studied the effects of regulatory reforms.

In this article, I use a dynamic equilibrium model of search and meeting frictions to quantify the impact of a recent change in taxicab fares and entry restrictions on the process that rules the meetings between passengers and taxicabs in New York City (NYC). In particular, the objective is to assess the effects of these policies on medallion prices and the frequency of meetings. Meeting frictions are a key feature of taxicab markets in most cities, so ideally, the analysis would be based on a model that explicitly accounts for the nature of these frictions and their implications for

* Manuscript received October 2000; revised March 2002.

[1] I am grateful to the New York City Transit Authority, Urbitran, Vollmer Associates, and especially to Bruce Schaller for having provided access to the Origin and Destination Survey of the Route 9A Reconstruction Project. I thank two anonymous referees for useful comments and suggestions. I also thank the C.V. Starr Center for Applied Economics at the New York University for financial support. All errors are mine. Please address correspondence to: Ricardo Lagos, Department of Economics, New York University, 269 Mercer Street, New York, NY 10003, USA. Tel: +212 998-8937. E-mail: *ricardo.lagos@nyu.edu.*

[2] The term "medallion" is used in reference to the actual aluminum medallion that is affixed to the hood of every taxicab in NYC representing a taxi license.

market outcomes. For these reasons, I follow the theoretical treatment of meeting frictions in Lagos (2000).[3]

I find that, everything else constant, the fare increase of March 1996 would have led to a 24% increase in medallion prices. On the other hand, also keeping all else constant, the price of a license would have dropped almost by 5% in response to the 3% expansion in the number of medallions that took place between 1996 and 1997. The exercise indicates that the combined effect of these policies was to increase medallion prices by around 19%, which would account for more than 60% of the observed increase in medallion prices from their pre-policy level.

The remainder of the article is organized as follows. Section 2 briefly outlines the theoretical model. Section 3 presents a geographical abstraction of the street map of NYC, analyzes an "Origin and Destination" data set, and uses them to parameterize the model. The policy experiments are conducted in Section 4. Section 5 concludes.

2. THE MODEL

The model is a special case of the one in Lagos (2000). There are large populations of taxicabs and passengers, with sizes v and u, respectively, who live on a grid with a finite number of locations n. Time is discrete and the horizon infinite. The focus is on steady-state outcomes. Every period, passengers get a wish to move to some other location. This is modeled by a Markov matrix $\mathbf{A}$, where a_{ij} denotes the probability that a passenger in location i wishes to reach j; hence $a_{ii} = 0$ and $\sum_j a_{ij} = 1$. For simplicity, assume $\mathbf{A}$ has a unique invariant distribution μ. Passengers must rely on taxicabs to reach their desired destination and each cab can take at most one passenger. No meeting frictions are assumed within each location, so if there are u_i passengers and v_i cabs in location i, then $m_i = \min\{u_i, v_i\}$ meetings result. Thus, letting $\theta_i = v_i/u_i$ and assuming meetings are random means that $p_i = \min\{1/\theta_i, 1\}$ is the probability a cab meets a passenger and $p_i\theta_i$ is the converse. Note that if there are more passengers than cabs in location i, then some passengers end up rationed and have no choice but to remain in i. In the following period, they draw a new desired destination and try, once again, to find a cab to get there.

A cab selling a trip from i to j earns profit $\pi_{ij} = b + \pi\delta_{ij}$, where δ_{ij} denotes the distance between locations i and j, π the per-mile profit, and b the "flag-drop" rate. It is convenient to think of the timing of events within a period as follows. At some point in every period there is a "meeting session" where all the

[3] Passengers often wait for taxicabs in some parts of the city, while at the same time, taxicabs tend to spend long periods of time waiting for passengers in some other parts (airports, for instance). For example, La Croix et al. (1992) report (in Table 1, p. 151) the "typical" waiting times borne by the taxicab drivers in five major U.S. airports to be between two and five hours. Although, some of this "spatial mismatch" may be due to standard search-type imperfect information problems, this evidence mainly suggests that cab drivers are often knowingly looking for passengers in locations where chances of finding them are low. The main objective of Lagos (2000) is to provide an explicit model of the meeting process that can give rise to an aggregate matching function.

period's meetings take place. Let V_i denote the value function for a cab in location i before this meeting session. If unable to contact a passenger, a cab must wait for a period until the next meeting session for another chance to find one; let U_i be the value function corresponding to this state. Cabs that did not meet passengers in location i are able to relocate before the next meeting session. Therefore, by driving empty, the cab is able to participate in next period's meeting session at some other location j. Formally, using modulo n and letting β be the discount factor, $U_i = \beta \max(V_i, V_{i+1}, \ldots, V_{i+n-1})$ and

$$V_i = p_i \sum_j a_{ij} \max\left(\pi_{ij} + \beta V_j, U_i\right) + (1 - p_i)\, U_i, \quad \text{for} \quad i, j = 1, \ldots, n$$

A steady-state equilibrium is a time-invariant distribution of passengers and cabs across locations together with a "no arbitrage" condition specifying that the latter have no incentive to relocate. Formally, an equilibrium is an allocation $\{u_i, v_i\}_{i=1}^n$ such that $\sum_i u_i = u$ and $\sum_i v_i = v$ (the distribution is consistent with the aggregate numbers of cabs and passengers); $V_i = V_n$ for all i (no arbitrage for cabs); and $\sum_j a_{ij} m_i = \sum_j a_{ji} m_j$ (steady state).

It is easy to show that $(1-\beta)V_i = p_i \pi_i$ in equilibrium, where $\pi_i = \sum_j a_{ij}\pi_{ij}$ is a cab's conditional (on having contacted a passenger) expected (over all possible destinations the passenger may desire) profit from selling a trip originating in location i. Let $\phi_i = \pi_i / (\sum_{j=1}^n \mu_j \pi_j)$, $\phi = \min(\phi_1, \ldots, \phi_n)$, and assume $\phi = \phi_n$ (i.e., label locations so that n denotes the "worst one" according to a cab's conditional expected profit). Then Lagos (2000) shows that the equilibrium allocations depend crucially on $\theta = v/u$ (usually known as aggregate "market tightness"). In particular, if $\theta \geq 1/\phi$, then $\{u_i, v_i\}_{i=1}^n = \{\mu_i u, \mu_i \phi_i v\}_{i=1}^n$, where the vector μ is the invariant distribution of $\mathbf{A}$. Conversely, if $\theta < 1/\phi$, then $\{v_i\}_{i=1}^n = \{\mu_i \phi_i v\}_{i=1}^n$, $\{u_i\}_{i=1}^{n-1} = \{\mu_i \phi v\}_{i=1}^{n-1}$, and $u_n = u - (1-\mu_n)\phi v$. The aggregate number of meetings is always given by $M(u, v) = \min\{u, \phi v\}$. It is easy to check that locations 1 through $n-1$ always exhibit excess supply of cabs, whereas location n has excess supply if $\theta > 1/\phi$, excess demand if $\theta < 1/\phi$, and the market there clears when $\theta = 1/\phi$.[4]

Summarizing, the equilibrium depends on $\mathbf{A}$, the matrix that determines trip destinations; on $[\delta_{ij}]_{i,j=1}^n$, the full set of bilateral distances; on u, a measure of demand for taxicab rides; and on the policy parameters b, π, and v. I now turn to the task of assigning reasonable values to these parameters.

3. THE MARKET FOR TAXICAB RIDES IN NEW YORK CITY

NYC is composed of five boroughs: Brooklyn, the Bronx, Queens, Staten Island, and Manhattan. This section uses the model summarized in the previous one to compute an approximation to the matching process according to which cabs and passengers meet in Manhattan, below 130th St.

[4] To simplify notation, here I argue as if there was a unique location that is relatively worse than the rest from a cab's perspective; but in general, there could be $k < n$ of them.

3.1. *Why NYC, Why Manhattan, and Why below 130th St.* A reason why it is appealing to apply the analysis to NYC's market for taxicab rides is that the city has recently changed the market's regulations. Within the last few years, the New York City Taxi and Limousine Commission (TLC), the regulating entity, has raised the fare and increased the number of medallions for the first time in 59 years.[5] The remainder of the article will attempt to quantify the effects of these policies on the frequency of meetings and medallion prices.

Choosing to focus on Manhattan, is natural because NYC's taxicab rides center on Manhattan. Only about 1% of NYC's taxi trips serve the "outerboroughs" (i.e., Brooklyn, the Bronx, Queens, Staten Island or Manhattan above 130th St.), and a negligible number of all trips both begin and end in the outerboroughs. Moreover, 70% of all trips transport Manhattan residents. With 78% of its households not owning a car, public transportation is a vital part of life in Manhattan. And, transporting 34% of all fare-paying bus, subway, taxi, or for-hire vehicle passengers traveling within Manhattan, taxis are a vital part of Manhattan's transportation network. Finally, the fact that at least 98% of all NYC's trips begin and end in Manhattan south of 130th St. suggests that the analysis should focus on this particular area.[6]

3.2. *A Simplified Map of Manhattan.* I proceed to simplify the map of Manhattan by dividing it into seven "locations." As can be seen in Figure 1, all locations have roughly the same size, and are bounded by the Hudson River on the West and the East River on the East.[7] Distances between locations are measured from a particular point (the "central point") chosen roughly in the geographical center of the location. The street corner chosen to be a location's central point is indicated with a dot in Figure 1.[8] Table 1 reports the distances between each pair of locations, measured in miles.

Having worked out a geographical abstraction for the area of interest (i.e., having assigned values to the δ_{ij}'s), I proceed to choose the parameter values that in equilibrium determine each location's fraction of passengers wanting to go to any other location. The following section uses data on the origin and the destination of trips in NYC to discipline **A**, the Markov matrix that in the model rules, passengers' wishes to move between locations.

[5] Until 1995, NYC's law used to limit the number of (yellow) taxicabs to 11,787. Legislative approval was granted to issue 400 additional taxicab licenses, breaking the 59-year-old cap. One hundred and thirty three licenses were auctioned in early May 1996. The remaining 267 licenses were auctioned in October 1996 and September 1997. The TLC increased the taxi fare, effective from March 1, 1996 to a new rate consisting of $2.00 for the initial 1/5 mile and $0.50 per additional 1/5 mile.

[6] In a slight abuse of terminology, I will sometimes refer to this particular area as "Manhattan". All facts reported in this section are documented in *The New York City Taxicab Fact Book* (NYC Taxi and Limousine Commission, 1994b) or obtained from the Origin and Destination data set described below.

[7] The north–south boundaries are as follows. For location 1: 130th St.–79th St.; for location 2: 79th St.–59th St.; for location 3: 59th St.–39th St.; for location 4: 39th St.–19th St.; for location 5: 19th St.–Houston; for location 6: Houston–Chambers; and for location 7: Chambers–Battery Park.

[8] The central points for locations 1 through 7, respectively, are: 89th St. and Lexington Av.; 69th St. and Lexington Av.; 49th St. between 6th and 7th Av.; 29th St. between 6th and 7th Av.; 8th St. and University Place; Canal and Lafayette; Wall St. and Broadway.

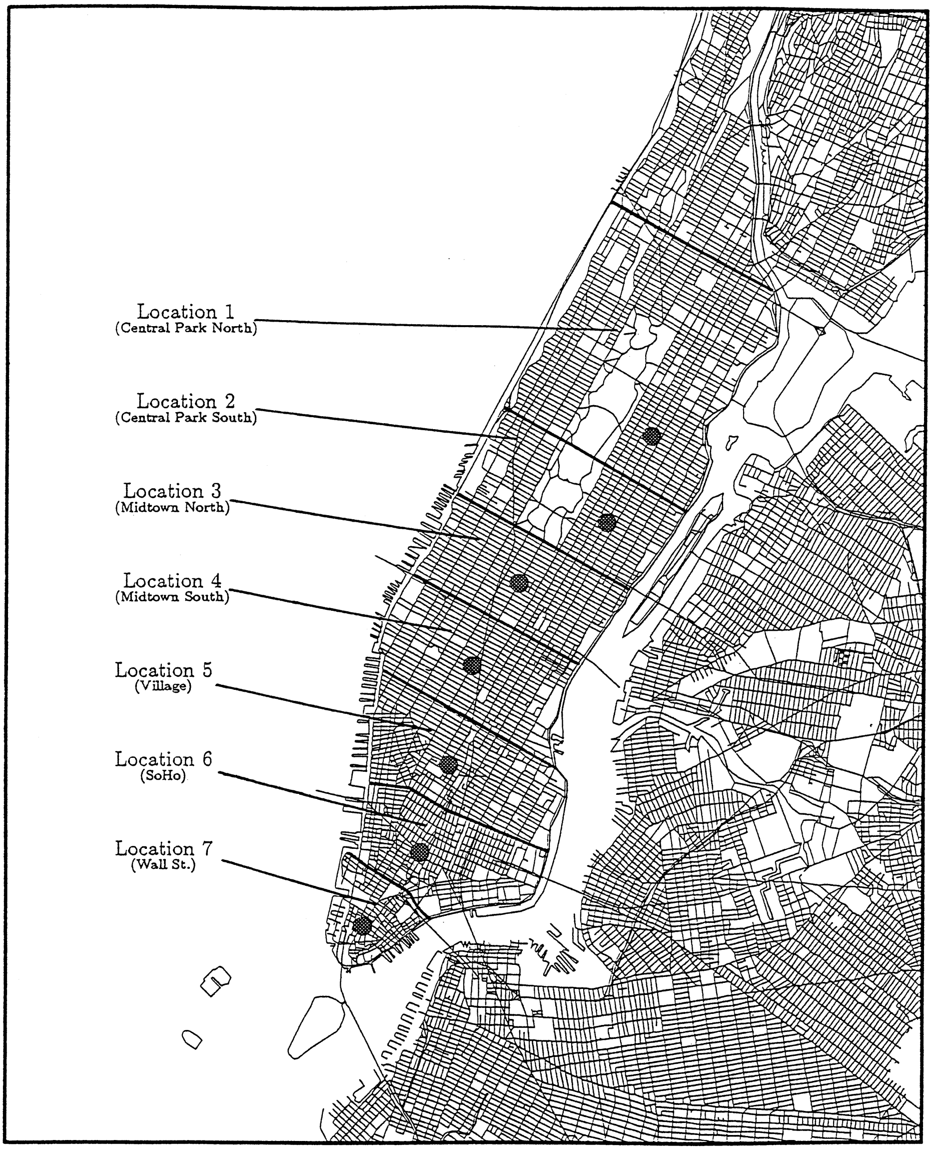

FIGURE 1

THE SEVEN LOCATIONS AND THEIR "CENTRAL POINTS"

3.3. *Origin and Destination Data.* A survey of taxicab trip sheets, conducted in November of 1988, recorded the origins and destinations of 22,604 trips in NYC. A typical trip sheet has 30 trips, so the sample consists of approximately 753 cabs.

In the model, a_{ij} represents the fraction of trips originating in location i that have location j as the destination. The calibration strategy followed is to construct

TABLE 1
PAIRWISE DISTANCES BETWEEN LOCATIONS (IN MILES)

$i \backslash j$	1	2	3	4	5	6	7
1	0	1.7	3.4	4.9	6.4	7.9	9.4
2	1.7	0	1.7	3.2	4.7	6.2	7.7
3	3.4	1.7	0	1.5	3.0	4.5	6.0
4	4.9	3.2	1.5	0	1.5	3.0	4.5
5	6.4	4.7	3.0	1.5	0	1.5	3.0
6	7.9	6.2	4.5	3.0	1.5	0	1.5
7	9.4	7.7	6.0	4.5	3.0	1.5	0

the empirical counterparts of the a_{ij}'s and then use them to parameterization the a_{ij}'s from the model. This procedure will guarantee that the model is consistent with the observed frequencies of trips between locations.

As a first step, the origin and destination data need to be aggregated to fit the geographical abstraction given by the seven locations described in Section 3.2.[9] For the majority of the trips, the aggregation process is straightforward since most of them originate in one of the seven locations and have one of the other six locations as destination. However, since the geographical abstraction adopted does not cover the whole of NYC, and is certainly not detailed enough to individualize all possible origins and destinations, some trips require special consideration. These trips can be classified into two types: (a) trips that involve an unmodeled location (for example, a trip from Brooklyn to Staten Island or from Manhattan to Queens), and (b) trips that start and end within one of the seven modeled locations.

The number of trips that involve at least one unmodeled location is small as a fraction of the total number of recorded trips (around 1%). Furthermore, the number of trips that originate in one of the seven locations and have destinations outside the modeled area is small even as a fraction of that location's trips. This fact is illustrated in Table 2, where t_{io} and t_i denote the number of trips that originate in location i and have a destination outside the modeled area, and the total number of trips from location i, respectively. Since the trips with outside destinations are too few to be quantitatively important, they are simply excluded from the analysis.

On the other hand, "intralocation" trips (ILTs) amount to 28% of all recorded trips, and are a sizable fraction of each location's trips (as can be seen in the last row of Table 2, where t_{ii} denotes the number of trips that stayed within location i). Since ILTs are likely to be trips too short to get out of a single location, disregarding them could make locations with high t_{ii}/t_i ratios look more attractive than they

[9] The grid of NYC on which the origin and destination survey was based consists of more than 500 cells (with two blocks in a typicall cell). Although the model can in principle be applied to any finite number of locations, the computational burden grows very fast with the number of locations. As an example, notice that with a large number of locations, n, even the simple task of measuring distances becomes demanding; as—even under symmetry—there are $\frac{n!}{2(n-2)!}$ pairwise distances to be measured.

TABLE 2
INTRALOCATION TRIPS AND TRIPS WITH DESTINATIONS OUTSIDE THE MODELED AREA

	i						
	1	2	3	4	5	6	7
t_{io}	15	15	52	26	17	12	10
t_{ii}	1370	996	2293	1166	400	82	77
t_i	3697	4151	7073	4384	2033	619	530
t_{io}/t_i	0.004	0.003	0.007	0.005	0.008	0.02	0.02
t_{ii}/t_i	0.37	0.24	0.32	0.27	0.20	0.13	0.14

TABLE 3
THE a_{ij}'S IMPLIED BY THE DATA

$i \backslash j$	1	2	3	4	5	6	7
1	–	0.60	0.26	0.08	0.03	0.01	0.02
2	0.35	–	0.46	0.12	0.04	0.01	0.02
3	0.16	0.32	–	0.41	0.06	0.02	0.03
4	0.09	0.10	0.56	–	0.19	0.03	0.03
5	0.09	0.07	0.24	0.43	–	0.13	0.04
6	0.06	0.05	0.22	0.23	0.32	–	0.12
7	0.13	0.07	0.27	0.16	0.10	0.27	–

really are from a taxicab's perspective, since it would make the fraction of long trips desired by passengers appear larger than it really is. ILTs are incorporated in the analysis by assigning them a destination outside their location of origin. For example, a fraction of the ILTs of location 3 are assigned location 2 as a destination, and the remainder are treated as trips from 3 to 4. More explicitly, the number of ILTs for 3 that are *imputed* as trips from 3 to 4 is denoted by $\hat{t}_{34}$ and given by $\hat{t}_{34} = \frac{t_{34}}{t_{32}+t_{34}} t_{33}$, where t_{34} and t_{32} denote the recorded number of trips from 3 to 4, and from 3 to 2, respectively. After imputing ILTs in this way, I obtain the fractions of trips from i to any of the other seven locations that are reported in Table 3.

If, as in the model, meetings between cabs and passengers with different desired destinations in the data occur randomly, then Table 3 provides the empirical counterpart of $\mathbf{A}$. Using the geographical abstraction of Section 3.2 and the parameter values reported in Tables 1 and 3, I next characterize all equilibria and quantify the meeting process generated by Manhattan's market for taxicab rides.

3.4. *Manhattan in Equilibrium.* As discussed in Section 2, when the number of cabs is "very large" relative to the number of potential passengers, there will be more cabs than people needing cabs in every location throughout the city: Every passenger is able to find a cab within the first model period. The parameter values reported in Tables 1 and 3 imply that conditional on having found a

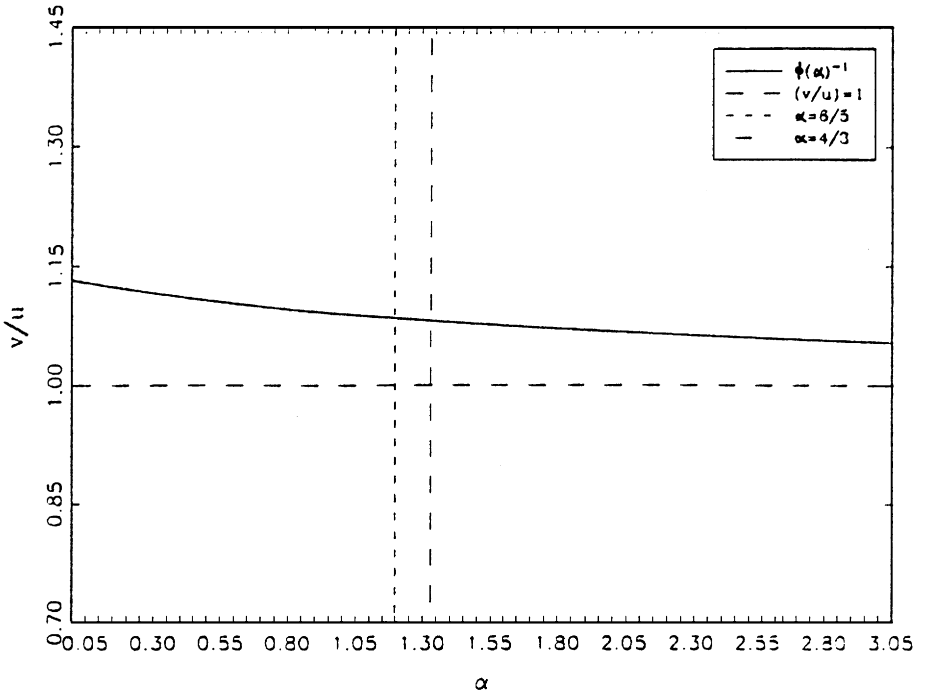

FIGURE 2

MANHATTAN'S "NO-FRICTIONS FRONTIER"

passenger, the expected profit for a cab in location 4 is lower than elsewhere in the city, and hence if the number of cabs is "small enough" relative to the number of people wishing to find a cab, there will be a shortage of cabs in this area.[10] The parameterization implies $\phi(\alpha) = \phi_4 = (\alpha + 2.11)/(\alpha + 2.39)$, where $\alpha = b/\pi$ is a measure of the attractiveness of short trips vis-à-vis long ones. The equilibrium allocations discussed in Section 2 imply that: (a) if $\theta > 1/\phi(\alpha)$, then there exists a unique equilibrium and it exhibits excess supply of cabs in all locations; (b) if $\theta = 1/\phi(\alpha)$, then there exists a unique equilibrium and it exhibits market clearing in location 4 and excess supply elsewhere; and (c) if $\theta < 1/\phi(\alpha)$, then there is a unique equilibrium and it exhibits excess demand in location 4 and excess supply elsewhere. Figure 2 depicts the "no-frictions frontier" $1/\phi(\alpha)$ implied by the parameterization of Manhattan's taxicab market.[11] The policy parameter α is on the

[10] The parameterization implies that $\pi_1 = b + 2.75\pi$, $\pi_2 = b + 2.16\pi$, $\pi_3 = b + 2.15\pi$, $\pi_4 = b + 2.11\pi$, $\pi_5 = b + 2.58\pi$, $\pi_6 = b + 3.12\pi$, and $\pi_7 = b + 4.80\pi$. So location 4 is the "worst location" from a cab's perspective. All the numerical results we report below are robust to this ranking in the sense that none of them would change if we had $\pi_2 = \pi_3 = \pi_4$ instead of $\pi_4 < \pi_3 < \pi_2$.

[11] There are (meeting) frictions below the solid line in Figure 2 in the sense that in that region there are *both* vacant taxicabs and unserved passengers who are unable to meet. The equilibrium exhibits vacant cabs but no unserved passengers when parameters lie anywhere above the solid line, and hence

TABLE 4

FLAG-DROP AND PER-MILE CHARGES: NEW AND OLD

	b	π	α
Old fare (1990–1996)	1.50	1.25	6/5
New fare	2	1.50	4/3

horizontal axis and aggregate tigthness $\theta = v/u$ on the vertical. Location 4 will have excess demand for cabs iff v/u lies below the solid line.

3.5. *Aggregate Market Tightness.* Taxi fares in NYC were last increased in March 1996. The fare structures before and after are reported in Table 4.[12] Notice that the flag-drop charge was increased proportionately more than the per-mile charge, causing their ratio (denoted by α) to rise.

The effect of this policy on the city-wide matching function depends on the values of the parameters v and u. To see this, notice that if $\theta \geq 1/\phi(6/5)$, then (given that v and u remain unchanged by the policy) the fare increase has no effect on the number of meetings between cabs and passengers. Conversely, if $\theta < 1/\phi(6/5)$, because this new rate structure increases the profits from short trips relatively more than it does for long ones, cabs find locations such as location 4 now more attractive than before. They respond by relocating toward the location(s) with excess demand, so the change in fare increases the number of contacts.[13] Therefore, values for v and u must be chosen before the effects of the policy can be assessed.

At the moment of the fare change, there were 11,787 licensed taxicabs in NYC. According to *The New York City Taxicab Fact Book* (NYC Taxi and Limousine Commission, 1994b), approximately 490,000 trips take place on an average weekday in NYC (roughly 340 contacts per minute), and the average trip was 11 minutes long. This means that in an average minute, only $11{,}787 - 3{,}740 = 8{,}047$ cabs are seeking passengers. Hence, $v = 8{,}047$. The average number of passengers u, will be chosen so that the model is consistent with the fact that on average there are 340 meetings per minute. The aggregate number of meetings in a model period is $\min\{u, \phi v\}$, so setting a period to be 1 minute, it must be that $\min\{u, 7{,}419\} = 340$ for the model to be consistent with the empirical frequency of meetings. Notice that this parametrization puts market tightness *above* the no-frictions frontier depicted in Figure 2.

there are no meeting frictions in that region (since the number of meetings is as large as it can be, namely equal to the "short side of the market").

[12] The figures for π ignore per-mile costs of operation (for example, gas) which are usually either small or not directly borne by the cab *driver.*

[13] The policy also induces a reallocation of cabs when $\theta \geq 1/\phi(6/5)$, but in that case the reallocation has no effect on the number of meetings because—since the total number of cabs in the city was "large enough" relative to the total number of movers—there were already enough cabs to serve all passengers in all locations before the policy took place. Throughout, we work under the presumption that passenger's relative demands for trips of different lengths (summarized by **A**) are not significantly altered by the policies. We provide some indirect evidence for this in Section 4.

4. POLICY

In this section, the calibrated model is used to quantify the impact of the fare increase and of the expansion in the number of medallions on the aggregate frequency of meetings and medallion prices.

All available evidence seems to suggest that u and v (which are taken parametrically in the analysis) do not respond much to fare increases. Two independent studies argue that the demand for taxicab rides in NYC is perfectly inelastic. According to *The New York City Taxicab Fact Book* (NYC Taxi and Limousine Commission, 1994b, p. 16), the 12% average fare increase of January 1990 induced no ridership loss. Similarly, Parsons et al. (1989, p. VIII12) report that the 22% average fare increase of May 1987 "brought about nearly 22% in extra revenue."

Since the number of medallions is fixed by law, the number of cabs cannot increase. A fare increase could, in principle, induce the existing number of licensed cabs to be operated more hours, with effects similar to those of an increase in the actual number of cabs. However, there are good reasons why a significant increase in the number of operating hours due to rising fares seems an unlikely outcome.

Of the 11,787 outstanding licenses, 6,818 are "corporate licenses" which are required by TLC regulation to be run for two shifts of nine to twelve hours' duration each day. Furthermore, around 33% of the 4,969 individual taxis are double-shifted (Parsons et al. 1989, p. VIII16). This means that 70% of the existing cabs are already running practically full time. In addition, there is also some evidence suggesting that drivers' preferences would work against the increase in operating hours.[14] Based on this and other evidence, the extensive study by Parsons et al. (1989, pp. VIII12–VIII16) concludes that most likely, "a fare increase of 23% would not appreciably increase the number of taxis double-shifting, nor cause a change in the basic operating pattern of the existing taxis." To be consistent with all this evidence, the effects of the recent fare increase will be analyzed keeping θ constant.

4.1. *Aggregate Meeting Frequency.* Figure 2 shows that whatever be the value of v/u, the key parameters were lying on the $\alpha = 6/5$ line before, and on the $\alpha = 4/3$ locus after the fare increase. Our calibration indicates that the city was above the no-frictions frontier before the fare was increased and entry relaxed. Thus the market will remain above the frontier after these changes (since both have the effect of raising ϕv). Increases in ϕv have no effect above the no-frictions frontier, so the policy has no effect on the aggregate number or frequency of meetings.[15]

[14] "According to taxi industry sources, some drivers will work fewer hours if they make more per hour, preferring the leisure time" (Parsons et al., 1989; p. VIII13).

[15] The effect of the fare increase alone on the aggregate meeting process would be very small even for the parametrization that generates the biggest possible impact. Let $\bar{u} = v\phi(4/3) \approx 7455$. The fare change posed has the biggest impact when $u > \bar{u}$, since in that case the market starts and ends below the no-frictions frontier and the number of matches increases from $v\phi(6/5)$ to $v\phi(4/3)$. For this case, the percentage change in the aggregate number of contacts is $[\phi(4/3) - \phi(6/5)]/\phi(6/5) \approx 0.003$. Thus independently of the parametrization of θ, the meeting process barely responds to the fare increase: The 11% increase in α will *at most* induce a 0.3% increase in the number of meetings.

4.2. *Medallion Prices.* Changes in the fare structure affect a cab driver's expected discounted income stream, and hence will affect the price of a medallion. The fare structure reported in the first row of Table 4 was in place from 1990 until the beginning of 1996. Over that period, the price of a (corporate) medallion stood, on average, at around \$196,000.[16] Let P denote the price of a medallion. Then the model implies that[17]

$$P(\theta, b, \pi, \beta) = \frac{(1/\theta)(b + 2.38\pi)}{1 - \beta}$$

The discount factor can be calibrated by solving $P(23.67, 1.5, 1.25, \beta^*) = 196{,}000$ for β^*. This procedure ensures that the model is consistent with the price of a medallion before the policy change. Given β^*, the pricing function P can be used to predict the change in medallion prices induced by the policies.

Keeping the number of medallions constant, the change in the fare structure implies $P(23.67, 2, 1.5, \beta^*) = 243{,}960$. That is, all else equal, the fare change alone would have induced a 24% increase in medallion prices. On the other hand, keeping the fare structure unchanged, increasing the number of medallions by 400 lowers their price to $P(24.84, 1.5, 1.25, \beta^*) = 186{,}727$. So on its own, the 3% increase in the number of licenses induces a 4.7% reduction in medallion prices.[18] Combined, the policies result in $P(24.84, 2, 1.5, \beta^*) = 232{,}408$, roughly a 18.6% increase relative to the 1990–1995 average. During the period 1996–September 2000, the average price of a medallion was \$254,000, namely 29.6% higher than in the 1990–1995 period. So according to the model, 63% of this increase was the result of the fare increase, whose impact was somewhat dampened by the simultaneous relaxation of entry restrictions.

5. CONCLUDING REMARKS

Within the last few years, the NYC has raised the taxicab fare and, for the first time in 59 years, increased the number of medallions. In this article, I parameterized a dynamic equilibrium model of meeting frictions and used it to quantify the impact of these policy changes on medallion prices and on the city-wide frequency of meetings. The model suggests that the policies are unlikely to have had significant effects on the city-wide number or frequency of meetings. On the other

[16] Monthly data on medallion transfers is readily available from the TLC. All medallion prices are CPI-deflated and reported in 1999 dollars.

[17] To see this, recall that $(1-\beta)V_i = p_i\pi_i$ in equilibrium. Since the parametrization indicates that $v\phi > u$, we can use the equilibrium allocations to write $(1-\beta)v\phi V_i = u\pi_i$. In equilibrium $V_i = V_n$ for all i, so we can let $i = n$ (where n denotes the worst location from a cab's perspective) and substitute the expressions for ϕ and π_i to arrive at $(1-\beta)V_n = (1/\theta)[b + \pi(\sum_i \mu_i \sum_j a_{ij}\delta_{ij})$. The parameterization from the previous section implies that $\sum_i \mu_i \sum_j a_{ij}\delta_{ij} \approx 2.38$.

[18] To re-calibrate v, notice that above the no-frictions frontier the number of meetings per minute is still 340. Hence given that the average duration of a trip remains unchanged and equal to 11 min, there now are 12,187 − 3,740 = 8,447 cabs seeking passengers, which is the value we use for v after the increase in the number of medallions.

hand, the model predicts that considered independently, the fare increase caused a 24% increase in medallion prices, whereas the 3% expansion in the number of medallions led to an almost 5% reduction. Medallion prices have been higher since these policies were implemented. According to the model, more than 60% of this increase can be attributed to the combined effect of the policies considered.

REFERENCES

De Vany, A. S., "Capacity Utilization under Alternative Regulatory Restraints: An Analysis of Taxi Markets," *Journal of Political Economy* 83(1) (February 1975), 83–94.

Friedman, M., *Price Theory: A Provisional Text* (Chicago: Aldine Publishing Co., 1962), 282.

La Croix, S. J., J. Mak, and W. Miklius, "Evaluation of Alternative Arrangements for the Provision of Airport Taxi Service," *Logistics and Transportation Review* 28 (June 1992), 147–66.

Lagos, R., "An Alternative Approach to Search Frictions," *Journal of Political Economy* 108(5) (October 2000), 851–73.

NYC Taxi and Limousine Commission, "Taxi Trip and Fare Data: A Compendium," October 1991.

———, "Hailing from New York: Results of an In-Cab Survey of Passengers Show the Role of Taxicabs in New York's Transportation Network," November 1993.

———, "Should the Taxi Fare Go Up?: Financial and Service Issues Findings and Recommendations," March 1994.

———, "The New York City Taxicab Fact Book," May 1994b, 3rd edition.

Orr, D., "The Taxicab Problem: A Proposed Solution," *Journal of Political Economy* 77(1) (February 1969), 141–47.

Parsons, Brinckerhoff, Quade & Douglas Inc, "Final Environmental Impact Statement on Additional Taxicab Licenses," Manuscript (prepared for the City of New York), June 1989.

Schroeter, J. R., "A Model of Taxi Service Under Fare Structure and Fleet Size Regulation," *Bell Journal of Economics* 14(1) (Spring 1983), 81–96.