Slot-Based Approaches to Airport Congestion Management

by

Jan K. Brueckner
Department of Economics
University of California, Irvine
3151 Social Science Plaza
Irvine, CA 92697
e-mail: jkbrueck@uci.edu

May 2008

Abstract

This paper analyzes slot-based approaches to management of airport congestion, using a model where airlines are asymmetric and internalize airport congestion. Under these circumstances, optimal congestion tolls differ across carriers, and since a slot-sale regime (with its uniform slot price) cannot duplicate this pattern, the equilibrium it generates is inefficient. Flight volumes tend to be too low for large carriers and too high for small carriers. Under a slot-trading regime or a slot auction, however, the existence of a fixed number of slots causes carriers to treat total flight volume (and thus congestion) as fixed, and this difference can lead to an efficient outcome.

Slot-Based Approaches to Airport Congestion Management

by

Jan K. Brueckner*

1. Introduction

Flight delays caused by airport congestion are a growing problem in both the US and Europe, and policymakers have struggled to formulate a response. To address rising congestion at Chicago's O'Hare Airport, the Federal Aviation Administration (FAA) took a micromanagement approach, prevailing on the airport's two major carriers (United and American) to cut their peak flight volumes while prohibiting smaller carriers from adding flights to fill the gap. More-systematic FAA interventions have occurred at New York airports, initially at LaGuardia and most recently at John F. Kennedy and Newark airports, where the FAA capped peak hour flight operations. All of these interventions followed surges in flights spurred by relaxation of long-standing slot constraints at O'Hare, LaGuardia and JFK, three airports where FAA-allocated slots give airlines the right to operate at particular times.[1]

While capping flight quantities, recent FAA interventions also envision a role for prices as a policy tool in attacking the congestion problem. In announcing the New York flight caps, the FAA proposed a system where carriers would each year relinquish a portion of their slots for redistribution via an auction system. This proposal is expected to stimulate the existing secondary slot market, where slots are traded among carriers. Contemporaneously, a position paper issued by the U.S. Department of Justice (Whalen et al., 2007) endorsed slot auctions as a mechanism for addressing airport congestion. Following these New York policy decisions, the FAA took an even more significant step by changing its rules on landing fees, which are charged to carriers for each flight operation at an airport. While landing fees traditionally depended only on aircraft weight, the new rules effectively allow the fees to vary by time of day. This change permits airports to implement congestion pricing, with high landing fees charged during peak hours and lower fees charged in off-peak periods.

With these new developments, price-based solutions to airport congestion have gained credibility, mirroring recent progress in implementing congestion pricing for roads (London and Stockholm are prominent examples). But the FAA's decisions have opened the door to several distinct, and potentially different, pricing approaches. While slot auctions represent one option, a related approach would involve a slot-sale regime, where the airport authority sets a slot price and allows carriers to purchase as many slots as they wish at that price. Under congestion pricing, carriers pay a congestion toll that is analogous to the slot price, but under an ideal structure, tolls are carrier-specific (depending on airport flight shares) rather than uniform. By contrast, under a slot-trading regime, the airport authority distributes slots to the carriers, who trade them at price that is again uniform. The existing secondary markets at slot-constrained airports approximate such a regime, but the current low trading volume suggests a need for institutional improvements (see Whalen et al. (2007)).[2] Given the importance of airport congestion as a policy problem, it is important to understand the potentially different impacts of these price-based regimes. The present paper is designed to achieve such an understanding by comparing the outcomes achieved under congestion pricing, a slot-sale regime, a slot-trading regime, and a slot auction.

In doing so, the paper parallels the analysis of Verhoef (2008)[3] while using a more tractable set of assumptions. The main simplification, which follows Brueckner and Van Dender (2008), is the assumption that carriers face perfectly elastic demands for air travel. This approach eliminates the flight-reducing distortion arising from the exercise of market power, allowing a sole focus on the distortion arising from the congestion externality, which tends to make flight volumes excessive. With market power eliminated and a constant-returns assumption modified, the analysis is able to derive parallel results that are simpler and more clearcut than those of Verhoef (2008). Once the analysis is complete, the two sets of results are compared.

Although a slot-sale regime is less commonly advocated than slot trading or slot auctions, an understanding of its performance is helpful when evaluating these other approaches. As a result, a significant portion of the paper is devoted to analyzing the slot-sale regime and comparing its performance to that of congestion pricing, which is known to generate the social optimum. The difference between these two regimes arises because of internalization of airport congestion, which occurs because carriers at congested airports often operate a large number of flights (road users, by constrast, operate a single vehicle). This fact means that, in scheduling an additional flight, a carrier will take into account the additional congestion costs imposed on the other flights it operates. The appropriate congestion toll then captures only the congestion imposed on other carriers, excluding the congestion the carrier imposes on itself. While Daniel (1995) was the first to recognize the potential for internalization of airport congestion, this pricing rule was advanced by Brueckner (2002, 2005) and further explored by other authors.[4]

A key implication of internalization is that asymmetric carriers should pay different tolls. A carrier with a large flight share at the congested airport, which internalizes most of the congestion from its operation of an extra flight, should pay a low toll, while a small carrier, which internalizes little congestion, should pay a high toll. Because a slot-sale regime, with its uniform price, cannot duplicate this inverse relationship between a carrier's flight share and its charge per flight, the regime is inefficient. By failing to account for differences in the internalization of congestion, the uniform slot price excessively penalizes large carriers and insufficiently penalizes small carriers for the congestion they create.

Given this pattern, large carriers will operate too few and small carriers too many flights under a slot-sale regime, provided the number of slots sold is close to the socially optimal flight volume. But since a welfare-maximizing airport authority, who sets the number of slots sold in a second-best fashion, may choose a slot total that diverges from the optimal flight volume, these relationships are not guaranteed in general. The analysis offers some partial results and a complete characterization in one special case, where the slot total is socially optimal and the flight volumes of large carriers are indeed too small and those of small carriers too large. Numerical results show the likely robustness of this result outside the special case.

If the model's large and small airlines are replaced by symmetric carriers, then the slot-sale regime's common price does not constitute an inefficient constraint, and the regime is efficient. The analysis shows that efficiency also obtains when carriers do not internalize congestion, behavior that Daniel (1995) and Daniel and Harback (2008) claim is realistic. They argue that non-internalizing behavior emerges in the presence of competitive-fringe carriers, who offset through their own flight increases any attempt by large carriers to limit self-imposed congestion.[5]

The analysis then turns to a slot-trading regime, where slots are distributed to carriers and then traded among them at a fixed price (also known as a "secondary market" for slots). Remarkably, the analysis shows that such a regime is efficient, overcoming the slot-sale regime's limitations, provided that the optimal number of slots is distributed. The key difference between the regimes is that carriers participating in a slot-trading regime understand that the total flight volume is fixed by the number of distributed slots, while slot-sale participants perceive no such constraint, expecting total flights (and airport congestion) to be affected by their slot purchases. The differing view of carriers under a slot-trading regime generates an efficient outcome.

The analysis relies on a highly stylized model, but the main conclusions should be robust to generalizations that offer greater realism. Adapting the framework of Brueckner and Van Dender (2008), the model portrays a congested airport served by two asymmetric carriers, with peak and off-peak periods collapsed into a single period that is always congested. In the analysis, carriers treat congestion tolls and slot prices as parametric and uninfluenced by their chosen flight volumes. This view is consistent with the usual approach to Pigouvian taxation, where the government, faced with a market distortion, computes the social optimum and levies taxes at a fixed rate to reach it. Economic agents, even if they otherwise enjoy market power, treat such Pigouvian taxes as parametric. Analogously, the airport authority in the present model has the information necessary to reach the first-best optimum via tolls or a second-best optimum via an appropriate slot price, and the carriers, even though they are nonatomistic, view these charges as immutable and unaffected by their own choices.

In an actual implementation of congestion tolls or a slot-sale regime, different behavior could emerge. For example, implementation of tolls might rely on an iterative approach, where peak-period tolls are initially computed based on current traffic volumes and then adjusted downward as traffic shifts toward off-peak periods. The carriers, perceiving a connection between flight volumes and tolls, would then have an incentive to manipulate the system, acting on the basis of false, understated demands for airport usage with the goal of depressing the toll. Similarly, under a slot-sale regime, the airport authority might take a trial-and-error approach in setting the slot price, encouraging the airlines to view the price as endogenous and thus subject to manipulation. The same incentive might arise under a slot-trading regime. If such manipulative behavior occurs, then the results of the present analysis are not strictly relevant, calling into question their usefulness as a guide for public policy. However, if the extent of manipulation is "small," then the results may still have some practical value. Whatever view of carrier behavior is correct, the urgency of the airport congestion problem makes any analysis of price-based remedies, including one based on standard Pigouvian assumptions, a high-priority undertaking.

Slot auctions are the final focus of the analysis. The discussion assumes that slots are allocated via a uniform-price, multi-unit auction, and in keeping with the non-manipulative behavior assumed in the prior analysis, strategic bidding is ruled out, with carriers assumed to make bids based on their true valuations of slots. This assumption, while strongly at variance with the huge auction literature, may provide an approximation to the actual outcome under a slot auction. The analysis shows that, without strategic behavior, the auction generates the same efficient outcome as the slot-trading regime.

The plan of the paper is as follows. Section 2 characterizes the social optimum and the laissez-faire equilibrium, where partially uninternalized congestion leads to an excessive total flight volume, and derives the congestion tolls required to support the optimum. Section 3 analyzes the slot-sale regime, providing numerical examples to supplement the analytical results, and then analyzes the slot-trading regime and a slot auction. Section 4 offers conclusions.

2. Basic Analysis

2.1. The setup

The analysis focuses on a single congested airport served by two airlines, denoted 1 and 2, who interact in Cournot fashion. Following Pels and Verhoef (2004), the model combines the peak and offpeak periods from Brueckner's (2002) analysis into a single congested period, an assumption that rules out reallocation of traffic between periods as a response to price-based congestion remedies. While the airlines experience common congestion, they are assumed to serve separate markets out of the congested airport, thus charging different fares. This assumption serves to generate asymmetry between the carriers, a crucial component of the ensuing analysis.

In order to maintain the simplest possible focus on the congestion phenomenon, the analysis suppresses the market-power element found in many previous models, including Verhoef (2008). In these models, a reduction in a carrier's flight volume reduces the level of airport congestion while also raising fares through a standard market-power effect. As a result, airline choices involve both the exploitation of market power and the desire to limit congestion. To focus solely on the congestion issue, market power is eliminated from the model by assuming that carriers face perfectly elastic demands for air travel.[6]

Accordingly, it is assumed that the passengers of airlines 1 and 2 are willing to pay fixed "full prices" of p_1 and p_2 for travel in and out of the congested airport, reflecting horizontal demand curves in the two markets. Airline 1 is assumed to serve the higher-price market, so that $p_1 \geq p_2$. Since passengers dislike airport congestion, which imposes additional time costs, the actual fares that the airlines charge must be discounted below these full prices.[7] To derive the discount, let f_1 and f_2 denote flight volumes for the two carriers, and let $t(f_1 + f_2)$ denote the extra time cost per passenger due to congestion and the resulting delays, a cost that depends on total flights at the congested airport. The function t satisfies $t(0) = 0$, $t' \geq 0$ (equality may hold over a range of low traffic levels), and $t'' \geq 0$ over the function's positive range. Taking account of passenger congestion cost, airline 1 is then able to charge a fare equal to $p_1 - t(f_1 + f_2)$, with airline 2 charging $p_2 - t(f_1 + f_2)$. When congestion cost is added to these fares, the resulting full prices are p_1 and p_2.

Letting s denoted the fixed seat capacity of an aircraft and assuming that all seats are filled, the total number of seats sold by carrier i is sf_i, $i = 1, 2$. For simplicity, s is normalized to unity, so that revenue for airline i is

$$[p_i - t(f_1 + f_2)]f_i, \quad i = 1, 2. \tag{1}$$

Note that, with the normalization of s, p_i becomes the full price per flight.

In addition to raising passenger time cost, airport congestion raises an airline's operating cost by $g(f_1 + f_2)$ for each flight. Like $t(\cdot)$, the function g satisfies $g(0) = 0$ and $g', g'' \geq 0$.

An airline also incurs operating costs per flight that depend on its own flight volume but are unrelated to airport-level congestion. These costs, given by $\tau(f_1)$ and $\tau(f_2)$, are assumed to increase with a carrier's flight volume, reflecting decreasing returns to scale ($\tau' > 0,\ \tau'' \geq 0$ hold along with $\tau(0) > 0$).

While the analogous τ function in Brueckner and Van Dender (2008), Verhoef and Pels (2004), and Verhoef (2008) is constant, reflecting constant returns to scale (a constant cost per flight), the assumption of decreasing returns is needed to generate sensible results in the presence of perfectly elastic demands when full prices differ across carriers. If the cost per flight were instead constant, the social optimum would involve a degenerate solution in which only the carrier serving the high-price market operates. Decreasing returns may, in any case, be a plausible assumption for a carrier operating at a congested airport. While intense use of runways and other airport infrastructure used jointly by both carriers is the source of airport congestion (an effect captured by the $t(\cdot)$ and $g(\cdot)$ functions), a busy airport will also have intense usage of carrier-specific facilities such as gates and baggage systems (captured by $\tau(\cdot)$). Such usage may well be subject to decreasing returns at high levels.

Using the above functions, total costs for the two airlines are given by $[\tau(f_1)+g(f_1+f_2)]f_1$ and $[\tau(f_2)+g(f_1+f_2)]f_2$. Airline 1's profits can then be written

$$\pi_1 \;=\; [p - t(f_1+f_2)]f_1 \;-\; [\tau(f_1) + g(f_1+f_2)]f_1 \tag{2}$$

and rewritten as

$$\pi_1 \;=\; [p - \tau(f_1)]f_1 \;-\; c(f_1+f_2)f_1 \tag{3}$$

where

$$c(f_1+f_2) \;\equiv\; t(f_1+f_2) \;+\; g(f_1+f_2) \tag{4}$$

gives passenger plus airline congestion cost per flight (note that t is multiplied by $s = 1$). Given the properties of the t and g functions, $c(0) = 0$ holds and $c' > 0,\ c'' \geq 0$. Analogously, carrier 2's profit is given by

$$\pi_2 \;=\; [p - \tau(f_2)]f_2 \;-\; c(f_1+f_2)f_2. \tag{5}$$

2.2. Social optimum

Consider first the social optimum. With perfectly elastic demands, consumer surplus is zero, which means that the social optimum maximizes the combined profits of the carriers. After adding the profit expressions in (3) and (5), differentiation with respect to f_1 and f_2 yields the first-order conditions

$$p_1 - \tau(f_1) - f_1\tau'(f_1) - c(f_1+f_2) - (f_1+f_2)c'(f_1+f_2) = 0 \quad (6)$$

$$p_2 - \tau(f_2) - f_2\tau'(f_2) - c(f_1+f_2) - (f_1+f_2)c'(f_1+f_2) = 0. \quad (7)$$

Computation of the Hessian determinant of total profit shows that satisfaction of the second-order condition is not guaranteed and must be assumed.[8]

From (6) and (7), a carrier's flight volume is optimal when the full price p_i per flight equals the marginal social cost of a flight, which is given by $\tau + f_i\tau' + c$ plus the marginal congestion damage from an extra flight. This latter cost is computed taking into account the congestion cost imposed on both carriers when an extra flight is operated. In particular, when f_1 is increased, passenger plus airline congestion costs for airline 1 (given by f_1c) increases by $c + f_1c'$, while these costs for airline 2 (given by cf_2) increase by f_2c'. The sum of the terms involving c', equal to $(f_1+f_2)c'(f_1+f_2) \equiv \text{MCD}$, gives the marginal congestion damage from an extra flight.

Inspection of (6) and (7) shows that airline 1, which serves the high-price market, operates more flights than airline 2 at the optimum. Denoting the social optimally values with an asterisk, $f_1^* > f_2^*$ then holds. This conclusion follows because $\tau(f) + f\tau'(f)$ is increasing in f under the maintained assumptions, implying that $f_1 > f_2$ must hold for both (6) and (7) to be satisfied given $p_1 > p_2$. For future reference, let $q = f_1 + f_2$ denote the total flight volume, and let $q^* = f_1^* + f_2^*$ denote its optimal value.

2.3. The laissez-faire equilibrium and congestion tolls

Consider next the laissez-faire equilibrium. Each airline, behaving in Cournot fashion, maximizes profit viewing the other airline's flight volume as fixed, yielding the first-order conditions

$$p_1 - \tau(f_1) - f_1\tau'(f_1) - c(f_1+f_2) - f_1c'(f_1+f_2) = 0 \tag{8}$$

$$p_2 - \tau(f_2) - f_2\tau'(f_2) - c(f_1+f_2) - f_2c'(f_1+f_2) = 0. \tag{9}$$

The carriers' second-order conditions are satisfied, and it is easily seen that (8) and (9) generate downward-sloping reaction functions. Airline 1's reaction function has a slope between -1 and 0 (f_1 is on the vertical axis) and is thus flatter than 2's function, which has a slope less than -1. As a result, the laissez-faire equilibrium is unique and stable. As in the case of the social optimum, $f_1 > f_2$ holds in the equilibrium.[9]

Focusing on airline 1, the difference between conditions (8) and (6) is the absence of f_2c' in the last term. This absence shows that, in scheduling an extra flight, airline 1 takes into account the additional congestion costs imposed on its own flights (f_1c'), ignoring the congestion imposed on airline 2 (f_2c'). Thus, while the airline internalizes some of the congestion from an extra flight, it ignores the impact on the other carrier. Airline 2 behaves in analogous fashion.

With both carriers ignoring a portion of the congestion they create, the total flight volume in the laissez-faire equilibrium is excessive relative to the social optimum. This conclusion follows from noting that the locii generated by (8) and (9) in (f_2, f_1) space, whose intersection determines the optimum, are both lower than the reaction functions generated by (6) and (7), a consequence of the larger multiplicative factor in the last term. As result, the socially optimal point must lie below both reaction functions. This conclusion in turn implies that the optimum lies below the line where $f_1 + f_2$ is constant at the equilibrium level, a line that lies between the reaction functions on either side of the equilibrium (it passes through the equilibrium point and has a slope of -1). However, even though the socially optimal point must lie below this line (yielding a smaller flight total), both individual flight volumes need not be smaller than the equilibrium levels, as would occur in the symmetric case. For example, the socially optimal point could lie to the northwest of the equilibrium (making f_1 larger and f_2 smaller than the equilibrium levels). A location to the southwest of the equilibrium is ensured if the degree of asymmetry between the carriers (the difference between p_1 and p_2) is sufficiently small.[10] Summarizing yields

Proposition 1. *The laissez-faire equilibrium has a larger total flight volume than the social optimum. If the degree of carrier asymmetry is small, flight volumes for the two carriers are individually larger than the optimal levels.*

The divergence between the laissez-faire equilibrium and the optimum can be eliminated by imposition of congestion tolls. The toll per flight is equal to that portion of the congestion damage from an extra flight not internalized by a carrier. The toll is thus equal to f_2c' for carrier 1 and f_1c' for carrier 2, with both expressions evaluated at the optimum. Thus, the tolls are given by

$$T_1 \;=\; f_2^*c'(f_1^*+f_2^*) \;=\; (1-\phi)(f_1^*+f_2^*)c'(f_1^*+f_2^*) \;=\; (1-\phi)\mathrm{MCD}^* \qquad (10)$$

$$T_2 \;=\; f_1^*c'(f_1^*+f_2^*) \;=\; \phi(f_1^*+f_2^*)c'(f_1^*+f_2^*) \;=\; \phi\,\mathrm{MCD}^* \qquad (11)$$

where

$$\phi \;\equiv\; \frac{f_1^*}{f_1^*+f_2^*} \;>\; \frac{1}{2} \qquad (12)$$

is airline 1's airport flight share and MCD^* is marginal congestion damage, both evaluated at the optimum. With imposition of these tolls, f_1T_1 and f_2T_2 are subtracted from the profit expressions in (3) and (5). Assuming that the carriers view the tolls as parametric, as discussed in the introduction, T_1 and T_2 are then subtracted from the expressions in the airline first-order conditions (8) and (9), and the solutions to these modified conditions coincide with the social optimum.

The key feature of the toll structure is that, because of airline 1's higher flight share, it pays a lower toll than airline 2. In other words, $T_1 = (1-\phi)\mathrm{MCD}^* < \phi\mathrm{MCD}^* = T_2$. The reason is that, since airline 1 has more flights, it internalizes more of the congestion damage from its operation of an extra flight than does airline 2. Since less congestion damage then goes uninternalized, airline 1 can be charged a lower toll. While this toll pattern is required for efficiency, the inverse association between a carrier's size and the toll it pays would generate controversy and political opposition from smaller carriers.

3. Slot-Based Regimes

3.1. The slot-sale regime

Alternate price-based approaches to reducing congestion rely on airport slots, with carriers needing to acquire a slot for each flight operated at the airport. Under one approach, the airport authority sells slots, allowing carriers to purchase as many as they wish at an announced price. This slot price is set by the authority to generate the desired total flight volume, and as discussed above, both carriers treat the price as parametric.

Let z denote the price of a slot and n denote airport authority's target flight volume. Then, since the terms f_1z and f_2z are subtracted from the profit expressions in (3) and (5), the slot-sale regime's equilibrium is characterized by the following conditions:

$$p_1 - \tau(f_1) - f_1\tau'(f_1) - c(f_1+f_2) - f_1c'(f_1+f_2) = z \qquad (13)$$

$$p_2 - \tau(f_2) - f_2\tau'(f_2) - c(f_1+f_2) - f_2c'(f_1+f_2) = z. \qquad (14)$$

$$f_1 + f_2 = n. \qquad (15)$$

Eqs. (13) and (14) are the carriers' first-order conditions, while (15) indicates that total slot purchases (equal to the total flight volume) equals the target level n. Note that (13) and (14) differ from the laissez-faire conditions (8) and (9) only in the appearance of z (rather than zero) on the RHS.

The equilibrium conditions in (13)–(15) generate solutions for f_1, f_2, and z conditional on n. Let the flight-volume solutions be denoted $f_1(n)$ and $f_2(n)$. As before, $f_1(n) > f_2(n)$ holds, so that carrier 1 operates more flights than carrier 2 for any given n.

Taking into account the dependence of f_1 and f_2 on the number of slots sold, a welfare-maximizing airport authority would select n in an optimal fashion, with the goal of maximizing total airline profit. The resulting n, denoted $\widehat{n}$, generates the second-best social optimum, conditional on use of the slot-sale regime. The corresponding second-best optimal f_1 and f_2

values are denoted $\widehat{f}_1 \equiv f_1(\widehat{n})$ and $\widehat{f}_2 \equiv f_2(\widehat{n})$. The conditions characterizing $\widehat{n}$ are developed below.

When carriers are asymmetric, the slot-sale regime is inefficient, as can be seen by contrasting (13) and (14) with the analogous conditions for the toll regime. Under that regime, z in (13) is replaced by T_1 and z in (14) is replaced by T_2. Since these tolls have different magnitudes, whereas carriers pay a common slot price, a slot-sale regime will not be able to generate the social optimum. By not taking into account airline 1's greater internalization of congestion, a slot-sale regime will tend to penalize airline 1 too much and airline 2 not enough for the congestion they create. Thus, the regime will tend to make the flight volume too small for the large carrier and too large for the small carrier.

Whether these tendencies end up making $\widehat{f}_1$ smaller and $\widehat{f}_2$ larger than the first-best optimal values f_1^* and f_2^* depends on the relationship between $\widehat{n}$, the optimal number of slots sold, and q^*, the socially optimal total flight volume. As will become clear below, the relationship between $\widehat{n}$ and q^* is ambiguous in general. If $\widehat{n}$ happens to equal q^*, then the above tendencies will indeed make $\widehat{f}_1$ too small and $\widehat{f}_2$ too large under the slot-sale regime. But when $\widehat{n} \neq q^*$, these relationships could be disrupted. For example, if $\widehat{n} \geq q^*$, then airline 2's insufficient congestion penalty combined with an excessive total number of flights will again make $\widehat{f}_2$ too large. But with the flight total excessive, airline 1's overly severe congestion penalty could lead to an $\widehat{f}_1$ value that is either larger or smaller than f_1^*. A formal statement is as follows:

Proposition 2. *If the optimal number of slots sold ($\widehat{n}$) is less than or equal to the socially optimal total flight volume (q^*), then the large carrier operates too few flights ($\widehat{f}_1 < f_1^*$). If $\widehat{n} \geq q^*$ holds, then the small carrier operates too many flights ($\widehat{f}_2 > f_2^*$). Combining this information, $\widehat{n} = q^*$ implies $f_1^* > \widehat{f}_1 > \widehat{f}_2 > f_2^*$, so that the slot-sale regime inefficiently narrows the difference between the flight volumes of the large and small carriers.*

The proposition thus shows that at least one of the inequalities $\widehat{f}_1 < f_1^*$ and $\widehat{f}_2 > f_2^*$ must hold, providing a partial characterization of the slot-sale regime's inefficiency.

To establish these results, first note that, after combining (8) and (9) and eliminating f_2,

q^* and f_1^* must satisfy

$$p_1 \; - \; \tau(f_1^*) \; - \; f_1^*\tau'(f_1^*) \;\; = \;\; p_2 \; - \; \tau(q^* - f_1^*) \; - \; (q^* - f_1^*)\tau'(q^* - f_1^*). \qquad (16)$$

Similarly, after combining (13) and (14), $\widehat{n}$ and $\widehat{f}_1$ must satisfy

$$\begin{aligned} p_1 \; - \; \tau(\widehat{f}_1) \; - \; \widehat{f}_1\tau'(\widehat{f}_1) \;\; &= \;\; p_2 \; - \; \tau(\widehat{n} - \widehat{f}_1) \; - \; (\widehat{n} - \widehat{f}_1)\tau'(\widehat{n} - \widehat{f}_1) \; + \; (\widehat{f}_1 - \widehat{f}_2)c'(\widehat{n}) \\ &> \;\; p_2 \; - \; \tau(\widehat{n} - \widehat{f}_1) \; - \; (\widehat{n} - \widehat{f}_1)\tau'(\widehat{n} - \widehat{f}_1), \qquad (17) \end{aligned}$$

where the inequality uses $\widehat{f}_1 > \widehat{f}_2$. Now suppose that $\widehat{n} \leq q^*$ holds while $\widehat{f}_1 \geq f_1^*$. Since $\tau(f) + f\tau'(f)$ is increasing in f, the LHS of (17) is then no larger than the LHS of (16). Similarly, the last expression in (17) is then at least as large as the RHS of (16). But the equality in (16) then implies that the last expression in (17) should be at least as large as the LHS expression, and the resulting contradiction establishes that $\widehat{f}_1 < f_1^*$ must hold when $\widehat{n} \leq q^*$. The remainder of Proposition 2 is proved in similar fashion.

3.2. Choosing the optimal n

As seen in Proposition 2, whether the flight volumes of the individual carriers are too large or too small under the slot-sale regime depends in part on the relationship between $\widehat{n}$ and q^*. To investigate this relationship, the conditions determining the optimal number of slots sold must be derived. The first step is to substitute the solutions $f_1(n)$ and $f_2(n)$ into the welfare function, which equals total airline profit (net of slot purchases) plus the airport authority's slot revenue. Since slot revenue cancels, the objective function reduces to $\pi_1 + \pi_2$ and can be written

$$W(n) \;\; = \;\; (p_1 - \tau[f_1(n)])f_1(n) \; + \; (p_2 - \tau[f_2(n)])f_2(n) \; - \; nc(n). \qquad (18)$$

The airport authority chooses n to maximize (18), and resulting the first-order condition is

$$[p_1 \; - \; \tau(f_1) \; - \; f_1\tau'(f_1)]f_1' \; + \; [p_2 \; - \; \tau(f_2) \; - \; f_2\tau'(f_2)](1 - f_1') \;\; = \;\; c(n) \; + \; nc'(n), \;\; (19)$$

where the n arguments of f_1 and f_2 are suppressed. Note that (19) requires a weighted average of $p_i - \tau(f_i) - f_i\tau'(f_i)$, $i = 1, 2$, to equal the RHS expression, where the weights are the derivatives f_1' and $1 - f_1' = f_2'$.

To compute f_1', (13) and (14) are combined, eliminating z, and f_2 is eliminated using (15). Total differentiation of the resulting condition yields

$$f_1'(n) \;=\; \frac{2\tau'(f_2) \;+\; f_2\tau''(f_2) \;+\; c'(n) \;-\; (2f_1 - n)c''(n)}{2\tau'(f_1) \;+\; f_1\tau''(f_1) \;+\; 2\tau'(f_2) \;+\; f_2\tau''(f_2) \;+\; 2c'(n)}, \tag{20}$$

where n arguments on the RHS are again suppressed. Since $f_1(n) > f_2(n)$, it follows that the expression $(2f_1 - n)c''(n)$ in the numerator of (20) is nonnegative. Inspection of (20) then establishes $0 < f_1'(n) < 1$, as required for the weights in (19) to be positive.

The first-order condition (19) in conjunction with (20) yields $\widehat{n}$, the optimal value of n, telling the airport authority the optimal number of slots to sell. Given the complexity of (19) and (20), a general comparison of $\widehat{n}$ to the socially optimal flight volume q^* appears to be infeasible. However, a simple statement is available in a particular special case. This is the case where the $\tau(\cdot)$ and $c(\cdot)$ functions are linear, with $\tau(f_i) \equiv \theta + \alpha f_i$ and $c(f_1 + f_2) \equiv \beta(f_1 + f_2)$. In this case, the second derivatives in (20) are zero and $f_1' = (2\alpha + \beta)/(4\alpha + 2\beta) = \frac{1}{2}$. As a result, the weighted averaging in (19) involves equal weights of one-half.

The implications of this fact for the relationship between $\widehat{n}$ and q^* can be seen by adding the social optimality conditions (6) and (7) and dividing by two, which yields (after inserting asterisks)

$$[p_1 \;-\; \tau(f_1^*) \;-\; f_1\tau'(f_1^*)]/2 \;+\; [p_2 \;-\; \tau(f_2^*) \;-\; f_2^*\tau'(f_2^*)]/2 \;=\; c(q^*) \;+\; q^*c'(q^*). \tag{21}$$

With $f_1' = \frac{1}{2}$, it is clear that, aside from notation, (19) and (21) are the same condition. Moreover, with linearity of τ, the LHS expression in (21) reduces to $\frac{1}{2}(p_1 + p_2 - 2\alpha q^*)$, while the LHS expression in (19) is $\frac{1}{2}(p_1 + p_2 - 2\alpha n)$. With the individual flight volumes dropping out, the two equations directly determine the optimal values of q and n, and the solutions are the same, with $\widehat{n} = q^*$. Summarizing yields

Proposition 3. *When operating and congestion costs per flight depend linearly on flight volumes, the optimal number of slots $\widehat{n}$ sold under a slot-sale regime equals the socially optimal flight volume q^*. As a result, $f_1^* > \widehat{f}_1 > \widehat{f}_2 > f_2^*$.*

Making use of Proposition 2, Proposition 3 thus generates a full characterization of the inefficiency of the slot-sale regime. Explicit solutions for the linear case are presented in the appendix.

A final question concerns the magnitude of the slot price z. Substituting (13) and (14) into (19) and rearranging yields

$$z \;=\; [(1 - f_1')f_1 \;+\; f_1' f_2]c'(n). \tag{22}$$

Thus, the slot price equals a weighted average of the flight volumes times marginal congestion cost, although the weights are reversed from (19). With linearity and the weights thus equal to one-half, (22) reduces to $z = nc'(n)/2$, and since $\widehat{n} = q^*$, the slot price associated with $\widehat{n}$ is $\widehat{z} = \frac{1}{2}\text{MCD}^*$. Recall that the optimal tolls for carriers 1 and 2 are respectively smaller and larger than this value (see (10) and (11)). However, $\widehat{z}$ is equal to the average of the two optimal tolls.

3.3. Numerical examples

Since further general analysis of the difference between the slot-sale and toll regimes is not feasible, this section presents numerical examples to determine whether some of the patterns seen with linear cost functions persist under nonlinearity. Consideration of two examples is sufficient to establish several points. In the examples, $\tau(f_i) \equiv \theta + 4f_i^4$ and $c(f_1 + f_2) \equiv 4(f_1 + f_2)^4$. While variation in the multiplicative constants in these functions (which are arbitrarily set equal to the quartic exponents) has little qualitative effect, use of quadratic rather than quartic functions generates an insufficient degree of decreasing returns unless the multiplicative constants are very large.

Table 1 shows numerical results in two cases with different values of $p_i - \theta$. In the first case, $p_1 - \theta = 9$ and $p_2 - \theta = 5$, while in the second case $p_1 - \theta$ equals 7 instead of 9. Starting with a comparison between q values, the upper panel shows that $\widehat{n} < q^*$ holds when $p_1 - \theta$

is high, with the reverse inequality holding in the second panel where $p_1 - \theta$ is lower (the values are close in each case, however). This reversal shows that optimal number of slots sold may be either larger or smaller than the socially optimal flight volume, with the linear case representing a particular instance where the two values are equal.

In each panel, comparison of the flight volumes reveals a relationship seen in the linear case. In both cases, the difference between f_1 and f_2 is narrower under the slot-sale regime than under the toll regime, with f_1 smaller and f_2 larger than the respective socially optimal values. This tendency may therefore be fairly general, holding under a variety of functional specifications.

Table 1 also shows the laissez-faire equilibrium, illustrating its excessive total flight volume and lower welfare level. Note that slot-sale regime achieves around 70 to 75 percent of the welfare gain generated by the first-best toll regime. Observe also that, in both cases, the socially optimal f_1 is larger than the equilibrium level, with the opposite relationship holding for f_2. Thus, the social optimum lies to the northwest of the equilibrium point, a possibility recognized in the earlier discussion. Additional computations confirm that a more-natural southwesterly location for the optimum (with both flight volumes smaller than the equilibrium levels) arises when the difference between p_1 and p_2 is sufficiently small.

The numerical results highlight the main source of inefficiency of a slot-sale regime relative to a toll regime: its tendency to overpenalize large carriers, whose high internalization of congestion is not taken into account, while underpenalizing small carriers. As long as the total flight volumes are not too different under both regimes, as in the present examples, this tendency will generate flight volumes that are smaller than optimal for large carriers and larger than optimal for small carriers.

3.4. When is the slot-sale regime efficient?

Although the results have given a negative picture of the efficiency of slot sales, such a regime is efficient under some circumstances. One such case is where carriers are symmetric, with $p_1 = p_2$. Since each carrier pays the same toll under symmetry, the slot-sale regime (with its common price) can perfectly replicate the toll regime and thus generate the first-best optimum.[11]

Slot sales are also efficient when carriers do not internalize congestion, behavior that Daniel (1995) and Daniel and Harback (2008) view as realistic. They argue that, when a Stackelberg leader interacts with competitive-fringe airlines, the offsetting behavior of these carriers (which increase their flights in one-for-one fashion as the leader reduces its flights) eliminates that carrier's incentive to take account of self-imposed congestion. To incorporate such behavior, suppose that both carriers, despite their nonatomistic sizes, treat airport congestion as parametric, ignoring the fact that the $c(\cdot)$ function depends on both flight volumes.[12]

With congestion viewed as parametric, the terms containing c' on the LHS of the laissez-faire conditions (8) and (9) vanish. The tolls required to generate the optimum are then the same across carriers and equal to $q^*c'(q^*) = \text{MCD}^*$, so that each carrier is charged for the full congestion damage done by an extra flight (including the damage done to its existing flights). Since tolls are uniform even though carriers are asymmetric, a slot-sale regime can again generate the social optimum.[13]

Summarizing yields

Proposition 4. *A slot-sale regime is efficient when* (*i*) *carriers do not internalize congestion or* (*ii*) *carriers internalize congestion but are symmetric, with* $p_1 = p_2$*. To achieve the first-best optimum, the airport authority sets* $\widehat{n}$ *equal to the relevant* q^* *value.*

3.5. Slot trading

Instead of slot sales, consider now a slot-trading regime, where the airport authority distributes slots free of charge to the carriers, who then trade them at a mutually agreed price in order to adjust flight volumes (this setup, also known as a secondary market, follows Verhoef (2008)). Let n_1 and n_2 denote the numbers of slots allocated to the two carriers, with $n_1 + n_2 = n$. Letting w denote the price at which the carriers agree to trade slots, profits for carriers 1 and 2 are now equal to

$$\pi_1 \;=\; [p - \tau(f_1)]f_1 \;-\; c(n)f_1 \;-\; w(f_1 - n_1) \tag{23}$$

$$\pi_2 \;=\; [p - \tau(f_2)]f_2 \;-\; c(n)f_2 \;-\; w(f_2 - n_2). \tag{24}$$

Several features of these expressions deserve note. First, if $f_i > n_i$, then carrier i is a buyer of slots and makes an outlay of $w(f_i - n_i) > 0$, while $f_i < n_i$ means that the carrier is a slot seller, earning revenues equal to the negative of the previous expression. Second, because a carrier understands that total flight volume remains constant when it trades a slot with the other carrier, the $f_1 + f_2$ argument of the congestion-cost function c remains constant at n, the fixed slot total.

The equilibrium conditions for slot trading are

$$p_1 - \tau(f_1) - f_1\tau'(f_1) - c(n) = w \tag{25}$$

$$p_2 - \tau(f_2) - f_2\tau'(f_2) - c(n) = w \tag{26}$$

$$f_1 + f_2 = n. \tag{27}$$

The first two conditions determine a carrier's desired flight volume (and thus the magnitude of its slot purchase or sale) when faced with the price w, while (27) says that the slot trades balance. This requirement implies $f_1 - n_1 = -(f_2 - n_2)$, which reduces to (27).

The key implication of (25) and (26) is that the expressions $p_i - \tau(f_i) - f_i\tau'(f_i),\ i = 1, 2$, are set equal to a common value. Since this same property holds at the social optimum (see (6) and (7)), slot trading is potentially efficient. The $p_i - \tau(f_i) - f_i\tau'(f_i)$ expressions, by contrast, do not assume a common value in the slot-sale equilibrium (see (13) and (14)). The reason is that those equations include airline-specific congestion impacts (the $f_i c'$ terms), a feature that accounts for the inefficient distortion of individual flight volumes. With total flights (and thus congestion) perceived as fixed under slot trading, these impacts are absent in (25) and (26).

To generate an efficient outcome under slot trading, the airport authority sets n, the number of slots distributed, equal to q^*, the socially optimal flight volume, and distributes the slots in some arbitrary fashion to the carriers. Then, it is easily seen that the solution to (25)–(27) has $f_1 = f_1^*$, $f_2 = f_2^*$ and $w = q^*c'(q^*) = \text{MCD}^*$. This conclusion follows because (25) and (26) are then identical to the conditions (6) and (7) evaluated at the social optimum. Summarizing yields

Proposition 5. *The slot-trading regime is efficient. To reach the social optimum, the airport authority distributes q^* slots to the carriers and allows them to trade.*

Observe that the equilibrium price for slots equals the atomistic congestion toll, MCD*, which was seen to be relevant in the absence of internalization. By contrast, the equilibrium value for the price z under a slot-sale regime was half as large ($\frac{1}{2}$MCD*) in the linear case. Note also that the optimality of slot trading is independent of the distribution of slots to the carriers. That distribution affects equilibrium profit levels but not the chosen flight volumes.

While much of the previous analysis generalizes naturally to a model with more than two carriers, the generalization of the slot-trading setup requires some discussion. When more than two carriers are present, bilateral trading would be replaced by a clearing house for slots, a structure similar to a slot-sale regime except that slots could be sold as well as bought. The airport authority (or other market maker) would adjust the price w until slot purchases and sales are equal. The essential behavioral difference between this setup and the slot-sale regime is the carriers' recognition that, *regardless of their choices, the total flight volume is fixed at a value equal to the total number of slots distributed*. As seen above, this recognition is the key to the efficiency of the slot-trading regime.

One could ask whether the same recognition might apply under the slot-sale regime. In other words, carriers might understand that the airport authority sets the slot price to achieve a target total flight volume, thereby viewing total flights as fixed. Even though this view might be appealing, it is not logically defensible given the rules governing slot sales. To view total flights as fixed, carrier 1 must anticipate that carrier 2 would reduce its flights on a one-for-one basis in response to its own slot purchases. While this view is correct when slots are traded, it is untenable in the context of sale regime, where carriers are free to purchase as many slots as they wish.

3.6. Slot auctions

Consider now an auction system as a means of allocating slots. Since a large number slots must be allocated, a multi-unit auction is the appropriate setup, and two variants exist. Under a "pay-as-bid" multi-unit auction (analogous to a standard first-price auction), carriers offer unit-specific bids for slots, specifying different prices for first, second, and later slots requested. The auctioneer then allocates a fixed slot supply to the highest bidders, making them pay their winning bids. Since bidders trade off the surplus gain from a lower price against a smaller chance of winning, the pay-as-bid system can elicit understatement of valuations. Under a "uniform-price" auction, which is analogous to a second-price auction, the bidders with the highest (unit-specific) bids again win slots, but they pay a common price equal to the highest bid not accepted. Since, in the event of winning at least one slot, a carrier's bids on later units may set the price that is paid (potentially being the highest unaccepted bid), an incentive again exists to understate the valuations of these later units. This latter conclusion is established by Ausubel and Cramton (2002), and the useful survey by Burguet (2000) provides further discussion.

Suppose that a uniform-price auction is used to allocate slots, and in keeping with the maintained assumption of non-manipulative behavior on the part of the carriers, imagine that strategic behavior is absent, so that carriers act on the basis of their true valuations for slots. This assumption is inconsistent with the analyses described above, but it may yield a rough approximation to the actual outcome under a uniform-price auction while facilitating a simple discussion.

As explained above, a carrier reports to the auctioneer a bid function giving its marginal willingness-to-pay for incremental slots, and without strategic behavior, this function coincides with actual marginal valuations. Carrier i thus reports a bid function equal to $p_i - \tau(f_i) - f_i\tau'(f_i) - c(n)$, where n now represents the number of slots to be auctioned. Since this number is announced in advance by the auctioneer, a carrier's bid function treats total flights (and thus the level of congestion) as fixed, as in the slot-trading regime.

Based on these bid functions, the auctioneer then sets a price y such that: carriers bidding at least y for incremental slots receive them; a total of n slots is allocated. It is easy to see that the resulting equilibrium conditions are identical to those from the sloting-trading regime, being given by (25)–(27) with y in place of w. As a result, the previous argument establishing efficiency of slot trading applies as well to a slot auction under the present assumptions, yielding[14]

Proposition 6. *With non-strategic behavior, a uniform-price, multi-unit slot auction is efficient. To reach the optimum, the airport authority auctions q^* slots.*

3.7. The effect of market power

In a closely related paper, Verhoef (2008) analyzes slot-based regimes when demand is imperfectly elastic, allowing carriers to exercise market power. Uninternalized congestion again tends to make flight volumes excessive, but carriers have a new incentive to limit their flight volumes in order to raise fares. The net effect of these two distortions is ambiguous, so that flight volumes in the laissez-faire equilibrium could either be too large or too small. Since Verhoef (2008) assumes that, under the toll regime, the two distortions must be addressed using a single toll instrument, the possibility of insufficient flight volumes means that optimal tolls could be negative.[15]

In Verhoef's model, two carriers serve the same market (facing a downward-sloping, linear demand curve), and asymmetry between them is generated by a difference in costs. The $\tau(\cdot)$ function is constant in his model (yielding constant returns), the $c(\cdot)$ function is linear, and the functions differ across carriers in a way that gives one carrier lower costs at any flight volume. Given the assumed cost structure and service to a single market, the high-cost carrier does not operate at the social optimum, an outcome that is supported by levying a positive toll on that carrier. By contrast, the low-cost carrier faces a negative toll (receives a subsidy), which raises its flight volume and corrects the market-power distortion. Note that uninternalized congestion vanishes with the absence of the other carrier, making this the only distortion present.

As in the present analysis, a slot-sale regime eliminates the differentiation of tolls required to support the social optimum. One consequence is that the high-cost carrier may operate under the slot-sale regime rather than being forced out of business, an undesirable second-best outcome. This pattern matches one of the conclusions of Proposition 3 (too large a flight volume for the smaller carrier), although Verhoef is not able to state an equally clearcut result given the greater complexity of his model. Nevertheless, the two papers offer a similar verdict on slot sales by highlighting the inefficiency of a structure that imposes a common charge on asymmetric carriers.

Verhoef also analyzes slot trading and shows that the trading equilibrium efficiently removes the high-cost carrier from the market. However, unless the airport authority can compel the low-cost carrier to use all of the slots it ends up holding, the social optimum cannot be generated. In other words, the authority can issue total slots equal to the socially optimal flight volume, but the low-cost carrier, even though it ends up holding this many slots, can choose instead the smaller flight volume corresponding the best monopoly output, leaving some slots unused.

4. Conclusion

This paper has analyzed slot-based regimes for management of airport congestion, using a model where airlines are asymmetric and internalize congestion. Under these circumstances, optimal congestion tolls differ across carriers, and since a slot-sale regime (with its uniform slot price) cannot duplicate this pattern, the equilibrium it generates is inefficient. Flight volumes tend to be too low for large carriers and too high for small carriers. Under a slot-trading regime, however, the distribution of a fixed number of slots causes carriers to treat total flight volume (and thus congestion) as fixed, and this difference leads to an efficient outcome as long as the number of slots is optimally chosen. A slot auction is efficient for the same reason.

All of these conclusions presume the absence of manipulative or strategic behavior on the part of the carriers, a view that may be inaccurate. Nevertheless, the results provide a benchmark for the evaluation of slot-based regimes, which may give an approximation to actual outcomes under more-sophisticated behavior by the carriers. Other aspects of the model are also highly stylized, but since the principles underlying the results would also emerge in more detailed and realistic frameworks, the findings are likely to be highly robust.

A key lesson of the analysis is that a slot-trading regime, where slots are distributed to the carriers and then traded through a clearing house, is equivalent to an efficient regime of carrier-specific congestion tolls. Since such a toll regime is bound to be controversial given that the tolls it generates are inversely related to carrier size, the analysis generates a clear presumption in favor of the equivalent slot-trading regime. This conclusion is welcome, given that trading already occurs at slot-constrained US airports, although at low volumes.[16] To foster more-active trading, the current bilateral system should be replaced by a central clearing house, and a slot purchase should confer clear property rights, replacing the current tenuous arrangement in which slots are ultimately the property of the FAA (Whalen et al. (2007) stress this point).

Although the analysis shows that a slot auction can also achieve efficiency, free distribution of slots might be preferable given the beleaguered airline industry's strenuous opposition to new, cost-increasing charges. To ensure that free distribution works well, hoarding of unused slots, meant to deny airport access to a carrier's competitors, must be prevented by "use-or-lose" requirements, which already exist. It might appear that new entrants (who hold no slots) are disadvantaged under this system, but their status is no different from that of an incumbent carrier seeking to increase its flight volume, which must purchase new slots to do so.[17]

Future work could relax some of the stylized assumptions of the model. One limitation is the absence of any distinction between peak and off-peak periods, and if this feature were added to the model, traffic reallocation between periods would occur in response to price-based interventions. While a multi-period setup like that of Brueckner (2002) involves greater complexity, such a model would presumably generate results similar to those of the present analysis, a conjecture that could be verified through additional research.

Another limitation of the analysis, mentioned in the introduction, is the assumption that the airport authority has the necessary information (on carrier congestion costs and passenger time values) for computing optimal congestion tolls or the optimal number of slots to auction or to distribute. While good estimates of the relevant parameters are available, as seen in the carefully calibrated numerical models of Daniel (1995) and Daniel and Harback (2008), it seems likely that actual implementation of price-based remedies would involve much less precision, while also relying on iterative or trial-and-error procedures that invite manipulation.[18] Such implementation problems were illustrated by the recent controversy over flight caps at the New York airports, where the carriers argued (possibly with some justification) that FAA caps were much too tight.[19] While these obstacles may make achievement of the optimum difficult, conceptual clarity regarding the effects of price-based remedies for airport congestion is nevertheless important, and this paper has attempted to provide it.

Appendix

In the linear case, the optimal total flight volume (which equals the optimal number of slots sold) is given by

$$q^* \quad = \quad \widehat{n} \quad = \quad \frac{p_1 + p_2 - 2\theta}{2\alpha + 4\beta}. \tag{$a1$}$$

Note that positivity of this expression is ensured if the intercept θ of the τ function is smaller than both of the full prices, a condition that is assumed to hold. Carrier 1's flight volumes at the optimum and under the slot-sale regime are given by

$$f_1^* \quad = \quad \frac{(p_1 - \theta)(\alpha + \beta) \; - \; \beta(p_2 - \theta)}{\alpha(2\alpha + 4\beta)} \quad > \quad \frac{(8\alpha + 10\beta)(p_1 - \theta) \; - \; 6\beta(p_2 - \theta)}{(4\alpha + 8\beta)(4\alpha + 2\beta)} \quad = \quad \widehat{f}_1, \tag{$a2$}$$

where the inequality follows from some algebra. Since $f_1^* + f_2^* = \widehat{f}_1 + \widehat{f}_2$, $(a2)$ implies $f_2^* < \widehat{f}_2$. Note that, although f_1^* and $\widehat{f}_1$ are both positive, positivity of f_2^* and $\widehat{f}_2$ (given by the expressions in $(a2)$ with p_1 and p_2 reversed) requires a sufficiently large value of α. This fact shows the need for a sufficient degree of decreasing returns in order to generate interior solutions.

Table 1

Numerical Examples

	Toll	*Slot Sale*	*Laissez Faire*
$p_1 - \theta = 9$ $p_2 - \theta = 5$			
$q,\ n$	0.707105	0.694426	0.842159
f_1	0.668742	0.550228	0.545672
f_2	0.038363	0.144198	0.296488
W	4.96840	4.82512	4.49634
$p_1 - \theta = 7$ $p_2 - \theta = 5$			
$q,\ n$	0.706804	0.708284	0.816954
f_1	0.562943	0.458862	0.478234
f_2	0.143861	0.249401	0.338720
W	3.72792	3.66090	3.46772

References

Ausubel, L.M., Cramton, P., 2002. Demand reduction and inefficiency in multi-unit auctions, Unpublished paper, University of Maryland.

Basso, L.J., Zhang, A., 2007. Congestible facility rivalry in vertical structures, *Journal of Urban Economics* 61, 218-237.

Braid, R.M., 1996. Peak-load pricing of a transportation route with an unpriced substitute. *Journal of Urban Economics* 40, 179-197.

Brueckner, J.K., 2002. Airport congestion when carriers have market power, *American Economic Review* 92, 1357-1375.

Brueckner, J.K., 2005. Internalization of airport congestion: A network analysis. *International Journal of Industrial Organization* 23, 599-614.

Brueckner, J.K., Van Dender, K., 2008. Atomistic congestion tolls at concentrated airports? Seeking a unified view in the internalization debate. *Journal of Urban Economics*, forthcoming.

Burguet, R., 2000. Auction theory: A guided tour, *Investigaciones Economicas* 24, 3-50.

Czerny, A., 2008. Managing congested airports under uncertainty. In: Czerny, A., Forsyth, P., Gillen D., Niemeier, H.-M. (Eds.), *Airport Slots: International Experiences and Options for Reform.* Ashgate, Aldershot, UK, pp. 111-126.

Gillen, D., 2008. Airport slots: A Primer. In: Czerny, A., Forsyth, P., Gillen D., Niemeier, H.-M. (Eds.), *Airport Slots: International Experiences and Options for Reform.* Ashgate, Aldershot, UK, pp. 41-59.

Daniel, J.I., 1995. Congestion pricing and capacity of large hub airports: A bottleneck model with stochastic queues, *Econometrica* 63, 327-370.

Daniel, J.I., Harback, K.T., 2008. (When) Do hub airlines internalize their self-imposed congestion delays? *Journal of Urban Economics* , 63, 583-612.

Forbes, S.J., 2006. The effect of air traffic delays on airline prices, Unpublished paper, University of California, San Diego.

Forsyth, P., Niemeier, H.-M., 2008. Setting the slot limits at congested airports. In: Czerny, A., Forsyth, P., Gillen D., Niemeier, H.-M. (Eds.), *Airport Slots: International*

Experiences and Options for Reform. Ashgate, Aldershot, UK, pp. 63-83.

Johnson, T., Savage, I., 2006. Departure delays, the pricing of congestion, and expansion proposals at Chicago's O'Hare airport. *Journal of Air Transport Management* 12, 182-190.

Kasper, D.M., 2008. Slots, property rights and secondary markets: Efficiency and slot allocation at US airports. In: Czerny, A., Forsyth, P., Gillen D., Niemeier, H.-M. (Eds.), *Airport Slots: International Experiences and Options for Reform*. Ashgate, Aldershot, UK, pp. 271-288.

Mayer, C., Sinai, T., 2003. Network effects, congestion externalities, and air traffic delays: Or why all delays are not evil, *American Economic Review* 93, 1194-1215.

Morrison, S.A., Winston, C., 2007. Another look at airport congestion pricing. *American Economic Review* 97, 1970-1977.

Oum, T.H., Zhang, Y., 1990. Airport pricing : Congestion tolls, lumpy investment, and cost recovery. *Journal of Public Economics* 43, 353-374.

Pels, E., Verhoef, E., 2004. The economics of airport congestion pricing. *Journal of Urban Economics* 55, 257-277.

Rupp, N.G., 2008. Do carriers internalize congestion costs? Empirical evidence on the internalization question. Unpublished paper, East Carolina University.

Schofield, A., 2007. Rush hour revenue: Airlines could face fees for peak-time flights if the U.S. government has its way. *Aviation Week and Space Technology*, October 29, pp. 41-42.

Small, K.A., 1992. *Urban transportation economics*, Harwood Academic Publishers, Chur, Switzerland.

Verhoef, E., 2008. Congestion pricing, slot sales and slot trading in aviation. Unpublished paper, Free University of Amsterdam.

Whalen, W.T., Carlton, D.W., Heyer, K., Richard, O., 2007. Proposal for a market-based solution to airport delays. Economic Analysis Group Discussion Paper, Antitrust Division, U.S. Department of Justice.

Zhang, A., Zhang, Y., 2006. Airport capacity and congestion when carriers have market power. *Journal of Urban Economics* 60, 229-247.

Footnotes

*I am indebted to Ami Glazer for very helpful suggestions and to Volodymyr Bilotkach, Joe Daniel, Ricardo Flores-Fillol, and especially Ken Small for additional comments. Any shortcomings in the paper, however, are my responsibility.

[1]Washington Reagan (National) Airport is also slot-constrained. See Gillen (2008) for a discussion of the history of the slot system at these airports, and see Forsyth and Niemeier (2008) for a good discussion of the economics of airport slots.

[2]Many European airports have slot systems, but while slot trading occurs in the UK, it is illegal at other EU airports (see Gillen (2008)).

[3]Erik Verhoef and I began discussing slot-based remedies for congestion several years ago, and he produced the main elements of a single-authored paper based on his preferred set of assumptions, a paper recently finalized as Verhoef (2008). Meanwhile, I recognized that the different approach of Brueckner and Van Dender (2008) could be applied fruitfully to the problem, leading to the present paper. The paper owes a debt to the original discussions with Verhoef, and it borrows directly from his treatment of slot trading, as explained further below.

[4]See Pels and Verhoef (2004), Zhang and Zhang (2006), Basso and Zhang (2007), and Morrison and Winston (2007). For an earlier contribution predating the analysis of internalization, see Oum and Zhang (1990). For representative studies in the road-pricing literature, see Small (1992) and Braid (1996).

[5]Brueckner (2002) and Mayer and Sinai (2003) provide evidence in favor of internalizing behavior, while Daniel (1995) and Daniel and Harback (2008) provide contrary evidence against it. A recent paper by Rupp (2008) provides additional negative evidence using an extension of Mayer and Sinai's methodology.

[6]Uncertainty regarding demand or costs is not present in the analysis. For an analysis of the impact of uncertainty in an airport-congestion model, see Czerny (2008).

[7]See Forbes (2006) for empirical evidence that airport congestion reduces fares.

[8]A sufficient condition for positivity of the Hessian determinant is $c'' \equiv 0$. However, when τ is a constant, the second-order condition fails, indicating that the optimum involves a corner solution with $f_2 = 0$.

[9]Since $p_1 > p_2$ holds and $\tau(f) + f\tau'(f) + fc'(q)$ is increasing in f holding q fixed, the f_1 value satisfying (8) must be larger than the f_2 value satisfying (9).

[10]This conclusion follows by continuity from the symmetric case, where the carriers' common optimal flight volumes are smaller than the common equilibrium level.

[11]To generate the optimum, the airport authority sets $\widehat{n}$, the number of slots sold, equal to the socially optimal flight volume, $q^* = 2f^*$, where f^* gives the carriers' common optimal flight volume (the value that satisfies (6) and (7) when $p_1 = p_2$). The equilibrium slot price z is then equal to the common optimal congestion toll, $\frac{1}{2}\text{MCD}^* = q^*c'(q^*)/2$. These conclusions follow from verifying that, when $\widehat{n} = q^*$, the solution to (13)–(15) yields $f_1 = f_2 = q^*/2$ along with $z = q^*c'(q^*)/2$.

[12]Brueckner and Van Dender (2008) show how the presumed non-internalizing behavior can occur using a simple model like the present one, but where carrier 2 is viewed as a collection of fringe airlines. However, to sustain the fringe's competitive behavior, their model requires constant returns to scale, with the $\tau(\cdot)$ function taking a constant value. Although decreasing returns in the present model precludes the use of Brueckner and Van Dender's approach, the effect of internalization's absence can still be investigated by imposing such behavior in ad hoc fashion, as described above.

[13]With the terms containing c' in the slot-sale first-order conditions (13) and (14) vanishing, it is easily seen that the solution to (13)-(15) coincides with the social optimum when the airport authority sets $\widehat{n} = q^*$. The equilibrium slot price is equal to the optimal (uniform) congestion toll, with $\widehat{z} = \text{MCD}^*$. Note that Brueckner and Van Dender (2008) refer to this toll as "atomistic" given that it would normally apply in a case where carriers are small enough to ignore their influence on congestion. This smallness is not satisfied in the present context; hence the ad hoc nature of the approach.

[14]Efficiency also obtains under a pay-as-bid auction.

[15]Another approach, suggested by Brueckner (2005), is to correct the market-power distortion through subidies paid at the level of the city-pair market (which can vary according to the extent of market competition) combined with congestion tolls levied at the airport level.

[16]Kasper (2008) stresses the overall virtues of a slot-trading regime (a secondary market) relative to congestion pricing.

[17]The analysis suggests that slot auctions are unnecessary if a slot-trading regime is already in place. The FAA's recent proposal for the New York airports does not conform to this view,

and the rationale is that auctioning a portion of the available slots each year will stimulate the existing secondary market, which is viewed as underdeveloped.

[18] Johnson and Savage (2006) provide "back-of-the-envelope" toll calculations applied to O'Hare Airport, showing that even simple methods can be useful.

[19] See Schofield (2007).