



Airline mitigation of propagated delays via schedule buffers: Theory and empirics[☆]

Jan K. Brueckner^{a,*}, Achim I. Czerny^b, Alberto A. Gaggero^c

^a University of California, Irvine, United States

^b Hong Kong Polytechnic University, Hong Kong

^c University of Pavia, Italy

ARTICLE INFO

Keywords:

Flight buffer
Ground buffer
Delay propagation
Airline scheduling

ABSTRACT

This paper presents an extensive theoretical and empirical analysis of the choice of schedule buffers by airlines. With airline delays a continuing problem around the world, such an undertaking is valuable, and its lessons extend to other passenger transportation sectors. One useful lesson from the theoretical analysis of a two-flight model is that the mitigation of delay propagation is done entirely by the ground buffer and the second flight's buffer. The first flight's buffer plays no role because the ground buffer is a perfect, while nondistorting, substitute. In addition, the apportionment of mitigation responsibility between the ground buffer and the second flight's buffer is shown to depend on the relationship between the costs of ground- and flight-buffer time. The empirical results show the connection between buffer magnitudes and a host of explanatory variables, including the variability of flight times, which simulations of the model identify as an important determining factor.

1. Introduction

Flight delays are a worldwide problem, a consequence of the substantial growth in air travel over recent decades. In the US alone, the cost of delays for passengers and airlines was estimated at \$32.9 billion in 2010 by Ball et al. (2010). In response to the problem, the US Department of Transportation requires all major US airlines to provide monthly information about delays, which generates widely viewed on-time rankings of the carriers. The European Union has imposed rules for passenger compensation and assistance in the event of long flight delays.

A major source of flight delays is airport congestion, which is the subject of a large literature (see Zhang and Czerny (2012) for a survey). But whether congestion leads to flight delays is largely under the control of the airlines, since they are free to set scheduled flight durations. In other words, the congestion-related lengthening of flight times can be built into airline schedules through a practice known as “schedule padding,” whose recent growth is documented by Forbes et al. (2019). While airport congestion may make flights longer, this schedule adjustment prevents them from arriving late.

Despite this overall adjustment in response to broad trends, flight times are still influenced by many random daily factors, including weather, mechanical issues, and unanticipated congestion, which can vary by day and hour around some expected level. Airline

[☆] We thank Tiziana D'Alfonso and several referees for helpful comments and Alex Luttmann for comments and for providing us with one of his data sets. Any shortcomings in the paper are our responsibility. Czerny acknowledges research support from the Hong Kong Research Grants Council General Research Fund (No. 15504918). The data for this paper are available at <https://data.mendeley.com/datasets/s53gcynkth/1>.

* Corresponding author.

<https://doi.org/10.1016/j.tre.2021.102333>

Received 28 June 2020; Received in revised form 29 March 2021; Accepted 4 April 2021

Available online 5 May 2021

1366-5545/© 2021 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license

(<http://creativecommons.org/licenses/by/4.0/>).

scheduling decisions take account of these random influences through the choice of “schedule buffers.” One type of buffer is known in the airline industry as a “block-time buffer” (we call it a “flight buffer”), and it equals the amount of time added to the shortest feasible flight time to get the scheduled arrival time. While a longer flight buffer reduces the chance of late arrival, it also makes an early arrival more likely, and while passengers dislike delays, they also do not want flights to routinely arrive early, an outcome that leaves time gaps that must be filled. Longer scheduled flight times also raise the airline’s (planned) operating costs, such as the cost of crew time. In setting its flight schedules, an airline will take all three factors (disutilities from lateness and earliness as well as operating costs) into account. Flight buffers are typically positive, reflecting a greater concern about late as opposed to early arrivals on the part of the carrier (following passenger preferences). Note that these same three elements also affect scheduling by other transportation providers, such as passenger railroads and intercity bus lines.

Delays depend on more than just the operating time of a particular flight. If the incoming aircraft arrives late, then the outbound flight is likely to depart late, possibly leading to its late arrival even if operating time is normal. Late-arriving aircraft are in fact the major source of flight delays, as seen in Fig. 1, accounting for more delays than mechanical and crew-related delays (“air carrier delays”) or weather.¹ This type of delay is known as “propagated delay” since it propagates from one flight to another, and it is also present in the railroad and bus contexts.

Propagated delay can be addressed through a long flight buffer, which reduces the chance of a late-arriving aircraft, but another tool is the “ground buffer.” This buffer is defined as the difference between the scheduled ground time and the shortest feasible aircraft turnaround time. A long ground buffer can absorb a late arrival of the inbound aircraft, allowing the next flight to depart on time despite this disruption. The flight buffer for the outbound flight can also address delay propagation, allowing the flight to arrive on time even if it departs late. Lengthening this flight buffer is costly, however, and longer ground times are also costly since they require more gate space.²

The purpose of the present paper is to analyze an airline’s choice of flight and ground buffers, both theoretically and empirically. Our theoretical framework differs from most other models because it is stylized, rather than fully realistic, and starts from first principles, treating propagated delay in a setting with just two flights. Thus, the airline chooses two flight buffers and one ground buffer, taking passenger disutilities from lateness and earliness into account along with buffer costs. The analysis yields a number of insights. A principal result is that the flight buffer for flight 1 plays no role in mitigating delay propagation, which is instead handled by the ground buffer and flight 2’s buffer. The reason is that, while flight 1’s buffer and the ground buffer can both address delay propagation, the former distorts flight 1’s scheduled arrival time, making use of the ground buffer preferable. The apportionment of the mitigation roles between the ground buffer and flight 2’s buffer depends on buffer costs. Numerical examples show how buffer magnitudes are affected by the variances of the random factors affecting the operating times of the two flights, showing that greater flight-time variability raises the flight buffers while having a complex effect on the ground buffer. Given the parallel to other transportation sectors, the paper’s theoretical findings extend beyond the airline industry.

The empirical work relies on US Department of Transportation data showing the daily operations of thousands of commercial aircraft. These data allow computation of flight and ground buffers over an aircraft’s operating day, which are then related to exogenous explanatory factors suggested by the model. For example, one set of regressions relates the magnitude of the flight buffer to the standard deviation of operating times for a particular flight (computed across the months of the sample for that flight’s operations in the previous year). The empirical work provides confirmation of some of the hypotheses suggested by the model while providing general insight into the determinants of flight and ground buffers.

The paper is related to several strands of previous work. Earlier papers in operations research, including [Deshpande and Arikan \(2012\)](#) and [Arikan et al. \(2013\)](#), present models of the choice of flight buffers, noting the similarity to the classic “newsvendor” problem of [Whitin \(1955\)](#) (a flight’s early (late) arrival is analogous to a vendor ending up with a surplus (shortage) of newspapers). [Deshpande and Arikan \(2012\)](#) consider US airlines and use the newsvendor approach to estimate the airlines’ ratios of earliness to lateness costs. They show that flight buffer choices depend on carrier types, route market shares, and route characteristics. [Arikan et al. \(2013\)](#) develop schedule robustness measures for airline networks and use them to show how US carriers use flight and ground buffers to absorb delay propagation. [Zhang et al. \(2018\)](#) present an analysis related to those of [Deshpande and Arikan \(2012\)](#) and [Arikan et al. \(2013\)](#). They show that the historical evolution of flight durations cannot explain increases in scheduled ground and flight times in the US and conclude that these increases instead have strategic motivations. [Kafle and Zou \(2016\)](#) theoretically decompose delays into propagated and newly-formed delays, and their empirical work investigates the determinants of propagated delays.³

In the economics literature, [Wang \(2015\)](#) offers a different theoretical approach to the choice of ground buffers, relating this choice to the level of competition while providing empirical evidence of this link. In other empirical work by economists, [Forbes et al. \(2019\)](#) use a much larger dataset to confirm the earlier finding of [Shumsky \(1993\)](#) showing that US carriers have added schedule padding

¹ This figure shows the apportionment of total US flight delays to different causes.

² Our analysis ignores the possibility that ground times may depend on passenger scheduling preferences, which are not present in the model. For instance, airlines may schedule a later departure and accept longer aircraft ground time if the result is a departure closer to the passenger’s preferred departure time.

³ In addition, [AhmadBeygi et al. \(2010\)](#) use a linear programming approach to show how readjustment of ground buffers, with no increase in total ground time among a set of flights, can achieve a reduction in propagated delays. In contrast to our stylized approach, they use a detailed and realistic model and rely on simulation. [Barnhart and Cohn \(2004\)](#) highlight the importance of operations research methods for the airline industry and describe the (at the time) state-of-the-art in airline schedule optimization. They further highlight the importance of schedule robustness and delay propagation (which they call “snowballing”) in the airline industry.

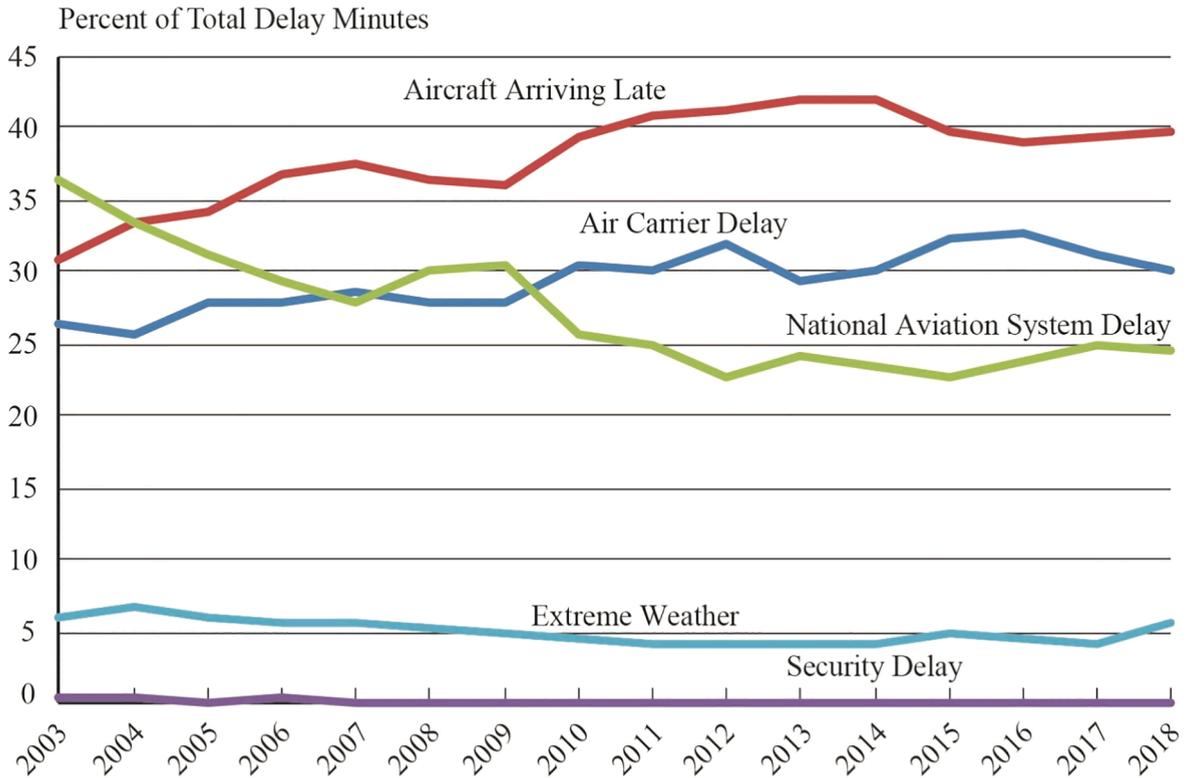


Fig. 1. Incidence of Delay Propagation.

(buffer time) over the years to improve their on-time performance. Forbes et al. (2019) show that this effect is stronger when airlines are large enough for required reporting of on-time performance, a criterion that excludes many regional carriers.⁴ Complementing these schedule-padding papers, transportation engineers have also extensively studied the determinants of scheduled block times, with contributions by Sohoni et al. (2011), Hao and Hansen (2014), Kang and Hansen (2017), Kang and Hansen, 2018 and Wang et al. (2019). Other studies by economists, which are not directly linked to our work, show the connection between market structure (mainly competition) and on-time performance (see, for example, Mazzeo (2003) and Prince and Simon (2015)).

The plan of the paper is as follows. Section 2 analyzes choice of the flight buffer in an introductory model with just one flight, where delay propagation is not an issue, and Section 3 analyzes the two-flight model. Section 4 presents numerical examples, while Section 5 offers a model extension that incorporates connecting passengers. Section 6 describes the empirical setup, and Section 7 presents the regression results. Section 8 offers conclusions.

2. Buffer choice for a single flight

The analysis starts by considering the buffer choice in a situation where the airline operates only a single flight, denoted flight 1. After this analysis is complete, the focus turns to the case where two flights are operated. While delay propagation is not an issue in the single-flight case, it is a crucial factor when two flights are operated. It should be noted that the models we analyze are highly stylized. They are designed to expose the economic incentives faced by airlines in the choice of schedule buffers, without being fully realistic.

Flight 1 departs at time 0 and has an uncertain duration. If no outside influences affect the flight’s operation, its duration is given by m_1 . Outside influences such as weather, mechanical issues, and unanticipated congestion can cause the flight duration to exceed m_1 , with the actual duration equal to $m_1 + \epsilon_1$, where ϵ_1 is a continuous, positive random variable with support $[0, \bar{\epsilon}_1]$. Note that while the congestion effect is partly random, appearing in ϵ_1 , a persistently high level of airport congestion would lead to a large value for m_1 .

In scheduling the flight’s arrival time, the presence of this random term will lead the airline to include a flight buffer, denoted b_1 , which is added to m_1 . The scheduled arrival time of the flight is thus $t_{a1} = m_1 + b_1$, whereas the actual arrival time, which captures the

⁴ For analysis of the incentives for integration of regional and mainline carriers and its impacts, see Forbes and Lederman (2009, 2010).

random effect, is $\widehat{t}_{a1} = m_1 + \epsilon_1$. The flight is late in arriving if $\widehat{t}_{a1} > t_{a1}$, or if $\epsilon_1 > b_1$, and it is early if $\epsilon_1 < b_1$. Passengers dislike being late or early, valuing a minute of late time by the amount x and valuing a minute of early time by the amount y . Therefore, if the flight is late, it generates disutility equal to $x(\epsilon_1 - b_1)$, whereas disutility is $y(b_1 - \epsilon_1)$ if the flight is early.

The probability that flight 1 is late is equal to the probability that ϵ_1 exceeds b_1 , which is given by

$$\Pr(\widehat{t}_{a1} > t_{a1}) = \int_{b_1}^{\overline{\epsilon}_1} f_1(\epsilon_1) d\epsilon_1 = 1 - F_1(b_1), \tag{1}$$

where f_1 is the density and F_1 the cumulative distribution function of ϵ_1 . Conversely, the probability of an early arrival is $F_1(b_1)$. The passenger's expected disutility from earliness or lateness is equal to

$$\Omega_1 = x \int_{b_1}^{\overline{\epsilon}_1} (\epsilon_1 - b_1) f_1(\epsilon_1) d\epsilon_1 + y \int_0^{b_1} (b_1 - \epsilon_1) f_1(\epsilon_1) d\epsilon_1, \tag{2}$$

where the first integral captures lateness and the second captures earliness. Passenger disutility from late and early arrivals affects fare revenue, and assuming that the airline seeks to maximize profit, it takes into account this disutility.⁵

But the airline also incurs planned operating costs from scheduled flight time and ground time. Scheduled flight costs include expenditures on fuel and crew salaries, while ground costs consist mainly of gate rental costs. To facilitate comparison with the two-flight model, suppose that the aircraft is only flown for part of the day, sitting on the ground for the remaining time, so that the airline has "excess capacity" (an alternate assumption is used below in the two-flight model). With T denoting the length of the day, ground time is $T - (m_1 + b_1) > 0$, where $m_1 + b_1$ is scheduled flight time. The capital cost of the aircraft is sunk, but letting c_f denote the operating cost per minute of scheduled flight time and c_g denote the cost of ground time, total operating costs are $c_f(m_1 + b_1) + c_g(T - (m_1 + b_1))$. The goal of the profit-maximizing airline is to minimize the sum of this expression and Ω_1 by choice of b_1 , which (ignoring constant terms) means minimizing Ω_1 from (2) plus $b_1(c_f - c_g)$.

The first-order condition for this minimization problem is

$$\frac{\partial \Omega_1}{\partial b_1} = -[1 - F_1(b_1)]x + F_1(b_1)y + c_f - c_g = (x + y) \left[F_1(b_1) - \frac{x}{x + y} \right] + c_f - c_g = 0, \tag{3}$$

and the second-order condition, which requires $(x + y)f_1(b_1) > 0$, is satisfied. Rearranging (3), the optimal buffer, denoted b_1^* , satisfies

$$F_1(b_1^*) = \frac{x + c_g - c_f}{x + y}, \tag{4}$$

a formula similar to one derived by [Deshpande and Arikan \(2012\)](#). Differentiation of (4) yields

$$\frac{\partial b_1^*}{\partial x} > 0, \quad \frac{\partial b_1^*}{\partial y} < 0, \quad \frac{\partial b_1^*}{\partial c_g} > 0, \quad \frac{\partial b_1^*}{\partial c_f} < 0. \tag{5}$$

with a greater disutility from lateness raising the buffer (thus reducing the chance of lateness) and a greater disutility from earliness reducing it. Similarly, a higher c_f (c_g) reduces (increases) b_1 .

3. The two-flight model

3.1. The setup

The airline is now assumed to operate two flights using the same aircraft, so that a delay for flight 1 can cause lateness of flight 2. While the existence of two flights introduces the possibility that some passengers connect from one flight to another, we assume initially that connections are absent, showing later in the paper how they affect the analysis.

Following the single-flight assumptions, m_2 denotes the undisrupted flight time for flight 2, with the actual flight time given by $m_2 + \epsilon_2$. The positive random term ϵ_2 has density f_2 , cumulative distribution function F_2 , and support $[0, \overline{\epsilon}_2]$. The flight buffers are b_1 and b_2 , so that the scheduled arrival times of flights 1 and 2 are $t_{a1} = m_1 + b_1$ and $t_{a2} = t_{d2} + m_2 + b_2$, where t_{d2} is flight 2's scheduled departure time. The scheduled aircraft ground time is denoted t_g , and the ground buffer, given by $b_g = t_g - \bar{t}_g$, is the excess of ground

⁵ Letting G denote the benefit from travel, the net benefit from a flight, denoted V , is equal to G minus the disutility expression in (2). Letting C denote expenditure on a nontravel good, overall passenger utility equals $U = C + V = Y - P + V$, where Y is income and P is the fare for travel. The airline sets the fare at a level that makes the consumer indifferent between traveling and not traveling, which yields utility Y (in this case $V - P = 0$). The fare for a flight is then equal to V , or G minus (2), and assuming that the passenger count is normalized to 1, fare revenue equals this expression.

Table 1
Flight 2's arrival time.

Flight 2's departure	Occurs when	Flight 2 late (early)	Arrival delay	Delay propagation?
On time	$\epsilon_1 \leq b_1 + b_g$	as $\epsilon_2 > (\leq) b_2$	$\max\{0, \epsilon_2 - b_2\}$	NO
Delayed	$\epsilon_1 > b_1 + b_g$	as $\epsilon_2 > (\leq) b_1 + b_2 + b_g - \epsilon_1$	$\max\{0, \epsilon_1 + \epsilon_2 - (b_1 + b_2 + b_g)\}$	YES if late

time over the minimum feasible aircraft turnaround time, denoted \bar{t}_g . The size of the ground buffer is thus set by choice of t_g .⁶ The scheduled departure time of flight 2 is then $t_{d2} = t_{a1} + t_g = m_1 + b_1 + t_g$, and the flight's scheduled arrival time is

$$t_{a2} = t_{d2} + m_2 + b_2 = t_{a1} + t_g + m_2 + b_2 = m_1 + b_1 + t_g + m_2 + b_2. \tag{6}$$

As before, the actual arrival time of flight 1 is $\hat{t}_{a1} = m_1 + \epsilon_1$. The actual arrival time of flight 2 equals $\hat{t}_{a2} = \hat{t}_{d2} + m_2 + \epsilon_2$, where \hat{t}_{d2} is the actual departure time of flight 2. To find \hat{t}_{d2} , note that if flight 1 is late in arriving, then the ground time will be reduced below the scheduled time t_g in an attempt to prevent delay propagation via late departure (and possible late arrival) of flight 2. However, ground time cannot be reduced below the minimum feasible time, equal to \bar{t}_g . Therefore, the actual departure time of flight 2 is given by

$$\hat{t}_{d2} = \max\{m_1 + \epsilon_1 + \bar{t}_g, t_{d2}\} = \max\{m_1 + \epsilon_1 + \bar{t}_g, m_1 + b_1 + t_g\}. \tag{7}$$

Flight 2 departs on time if $\epsilon_1 + \bar{t}_g < b_1 + t_g$ or

$$\epsilon_1 < b_1 + t_g - \bar{t}_g = b_1 + b_g. \tag{8}$$

Note that satisfaction of (8) is ensured if flight 1 is early, or if $\epsilon_1 < b_1$. But the inequality can also be satisfied when flight 1 is late provided that it is not too late. Using (8), the probability of an on-time departure for flight 2 is $F_1(b_1 + b_g)$, with late departure (which reverses the inequality in (8)) having probability $1 - F_1(b_1 + b_g)$.

The formula for \hat{t}_{d2} allows derivation of conditions for late arrival of flight 2. Using (6), flight 2 is late in arriving when

$$\hat{t}_{a2} = \hat{t}_{d2} + m_2 + \epsilon_2 > t_{a2} = m_1 + b_1 + t_g + m_2 + b_2. \tag{9}$$

If flight 2 departs on time, so that $\hat{t}_{d2} = m_1 + b_1 + t_g$ from (7), then (substituting in (9)), it arrives late when

$$\epsilon_2 > b_2, \tag{10}$$

an outcome that has probability $1 - F_2(b_2)$ (early arrival occurs when (10) is reversed and has probability $F_2(b_2)$). If flight 2 departs late, so that $\hat{t}_{d2} = m_1 + \epsilon_1 + \bar{t}_g$, then (substituting in (9)), it arrives late when

$$\begin{aligned} \epsilon_1 + \epsilon_2 > b_1 + b_2 + b_g, \quad \text{or} \\ \epsilon_2 > b_1 + b_2 + b_g - \epsilon_1. \end{aligned} \tag{11}$$

With late departure, flight 2 arrives early when (11) is reversed.

Table 1 summarizes the preceding information, while showing flight 2's arrival delay. A key observation from the table is that when flight 2's departure is delayed, its arrival delay depends on the sum $b_1 + b_2 + b_g$. Therefore, the ground buffer affects delay by altering this sum, conditional on a late departure for flight 2. However, b_g also affects whether a late departure occurs via the direction of the inequality in the second column of the table. As will be seen, the overall impact of b_g on the airline's objective function operates through both these channels.

Using (11) and the reverse of inequality (8), the probability of late departure and late arrival for flight 2 is given by

$$\Pr(\hat{t}_{a2} > t_{a2} \cap \hat{t}_{d2} > t_{d2}) = \int_{\epsilon_1 = b_1 + b_g}^{\bar{\epsilon}_1} \int_{\epsilon_2 = b_1 + b_2 + b_g - \epsilon_1}^{\bar{\epsilon}_2} f_1(\epsilon_1) f_2(\epsilon_2) d\epsilon_2 d\epsilon_1. \tag{12}$$

Similarly, using the reverse of inequality (11) and the reverse of (8), the probability of late departure and early arrival for flight 2 is

$$\Pr(\hat{t}_{a2} < t_{a2} \cap \hat{t}_{d2} > t_{d2}) = \int_{\epsilon_1 = b_1 + b_g}^{\bar{\epsilon}_1} \int_{\epsilon_2 = 0}^{b_1 + b_2 + b_g - \epsilon_1} f_1(\epsilon_1) f_2(\epsilon_2) d\epsilon_2 d\epsilon_1. \tag{13}$$

Combining all this information, the probability of a late arrival for flight 2 is given by

⁶ Note that, unlike flight times, ground time does not have a stochastic component, which would increase the complexity of the analysis. Random ground time could be a consequence of weather at the airport, which could lead to a delayed departure exactly like a late inbound flight. However, in the simpler model of Brueckner et al. (2020), where ϵ_1 and ϵ_2 are discrete rather than continuous random variables, the effect of random ground times can be analyzed.

$$\begin{aligned} & \Pr(\widehat{t}_{a2} > t_{a2}) = \\ & \Pr(\widehat{t}_{a2} > t_{a2} \mid \widehat{t}_{d2} = t_{d2})\Pr(\widehat{t}_{d2} = t_{d2}) + \Pr(\widehat{t}_{a2} > t_{a2} \cap \widehat{t}_{d2} > t_{d2}) = \\ & [1 - F_2(b_2)]F_1(b_1 + b_g) + \int_{\epsilon_1=b_1+b_g}^{\bar{\epsilon}_1} \int_{\epsilon_2=b_1+b_2+b_g-\epsilon_1}^{\bar{\epsilon}_2} f_1(\epsilon_1)f_2(\epsilon_2)d\epsilon_2d\epsilon_1. \end{aligned} \tag{14}$$

Note that a conditional probability can be used in the first half of (14) because whether flight 2 is late conditional on an on-time departure (which depends only on ϵ_2) is independent of whether the departure is on-time (which depends only on ϵ_1). The absence of this kind of independence when flight 2 is late requires the different kind of expression in the last half of (14).

Similarly, the probability of early arrival for flight 2 is given by

$$\begin{aligned} & \Pr(\widehat{t}_{a2} < t_{a2}) = \\ & \Pr(\widehat{t}_{a2} < t_{a2} \mid \widehat{t}_{d2} = t_{d2})\Pr(\widehat{t}_{d2} = t_{d2}) + \Pr(\widehat{t}_{a2} < t_{a2} \cap \widehat{t}_{d2} > t_{d2}) = \\ & F_2(b_2) \left[1 - F_1(b_1 + b_g) \right] + \int_{\epsilon_1=b_1+b_g}^{\bar{\epsilon}_1} \int_{\epsilon_2=0}^{b_1+b_2+b_g-\epsilon_1} f_1(\epsilon_1)f_2(\epsilon_2)d\epsilon_2d\epsilon_1. \end{aligned} \tag{15}$$

3.2. The airline's objective function

The disutility for passengers of flight 1 is again given by in Ω_1 in (2). The expected disutility from late arrival of flight 2 is given by

$$\begin{aligned} \Omega_{2, \text{late}} &= xF_1(b_1 + b_g) \int_{b_2}^{\bar{\epsilon}_2} (\epsilon_2 - b_2)f_2(\epsilon_2)d\epsilon_2 + \\ & x \int_{\epsilon_1=b_1+b_g}^{\bar{\epsilon}_1} \int_{\epsilon_2=b_1+b_2+b_g-\epsilon_1}^{\bar{\epsilon}_2} \left[\epsilon_1 + \epsilon_2 - (b_1 + b_2 + b_g) \right] f_1(\epsilon_1)f_2(\epsilon_2)d\epsilon_2d\epsilon_1. \end{aligned} \tag{16}$$

The first half of (16) captures disutility from a late arrival when flight 2 departs on time (note that $1 - F_2$ in (14) is replaced by the integral). The rest of (16) captures late disutility when the departure is late (note that the bracketed term is added inside the integrand in (14)).

Similarly, the expected disutility from early arrival of flight 2 is

$$\begin{aligned} \Omega_{2, \text{early}} &= yF_1(b_1 + b_g) \int_{\underline{\epsilon}_2}^{b_2} (b_2 - \epsilon_2)f_2(\epsilon_2)d\epsilon_2 \\ & + y \int_{\epsilon_1=b_1+b_g}^{\bar{\epsilon}_1} \int_{\epsilon_2=0}^{b_1+b_2+b_g-\epsilon_1} \left[b_1 + b_2 + b_g - (\epsilon_1 + \epsilon_2) \right] f_1(\epsilon_1)f_2(\epsilon_2)d\epsilon_2d\epsilon_1. \end{aligned} \tag{17}$$

The passenger-disutility portion of the airline's objective function, denoted by Ω , is given by the sum of (2), (16), and (17):

$$\Omega = \Omega_1 + \Omega_{2, \text{late}} + \Omega_{2, \text{early}}. \tag{18}$$

In contrast to the single-flight case, our airline cost expression assumes the absence of excess capacity in the two-flight case, with the airline adjusting its aircraft lease to exactly cover the scheduled usage of the plane (see below for an alternate assumption). The airline's costs are then given by $c_f b_1 + c_f b_2 + c_g b_g$ plus a constant involving the m_i s, with the carrier paying for only the ground time allotted to the ground buffer and not for any time beyond the termination of flight 2. The airline minimizes the sum of Ω and this cost expression, and the derivatives of Ω with respect to b_g , b_1 and b_2 are computed in the appendix. Adding the relevant buffer cost to these derivatives and setting the resulting expressions equal to zero yields the following first-order conditions:

$$b_g : \quad \frac{\partial \Omega}{\partial b_g} + c_g = \tag{19}$$

$$(x + y) \int_{\epsilon_1=b_1+b_g}^{\bar{\epsilon}_1} \left[F_2(b_1 + b_2 + b_g - \epsilon_1) - \frac{x}{x + y} \right] f_1(\epsilon_1)d\epsilon_1 + c_g = 0$$

$$b_1 : \quad \frac{\partial \Omega}{\partial b_1} + c_f = \frac{\partial \Omega}{\partial b_g} + (x + y) \left[F_1(b_1) - \frac{x}{x + y} \right] + c_f = 0 \tag{20}$$

$$b_2 : \quad \frac{\partial \Omega}{\partial b_2} + c_f = \frac{\partial \Omega}{\partial b_g} + (x + y)F_1(b_1 + b_g) \left[F_2(b_2) - \frac{x}{x + y} \right] + c_f = 0. \tag{21}$$

Note in (20) and (21) that $\partial \Omega / \partial b_1$ and $\partial \Omega / \partial b_2$ are equal to $\partial \Omega / \partial b_g$ plus the second term in the relevant equation.

3.3. Characterizing the optimum

The immediate implication of the first-order conditions is that the flight buffer b_1 has the same value as in the single-flight model. This conclusion can be seen by using (19) to substitute $-c_g$ in place of $\partial\Omega/\partial b_g$ in (20), which yields a condition that matches (3) from the single-flight model. With flight 1's buffer set as if the flight were operating in isolation, the buffer therefore plays no role in mitigating delay propagation. Thus,

Proposition 1. *Responsibility for mitigation of delay propagation falls only on the ground buffer and the flight buffer for flight 2.*

The lack of a delay-propagation role for flight 1's buffer makes sense. Even though b_1 and the ground buffer are, in effect, perfect substitutes in addressing delay propagation, use of the ground buffer instead of the flight buffer does not distort the balance between late and early disutilities for flight 1, making it the preferred tool. However, suppose the ground buffer were somehow constrained below its optimal value, due to a shortage of airport gate capacity, for example, which could be caused by hoarding of gates by a dominant airline (see Ciliberto and Williams (2010) for empirical evidence). In this case, b_1 would help to address delay propagation along with the other buffers. Such a constraint would make $\partial\Omega/\partial b_g$ in (20) less than $-c_g$, causing b_1 to rise above its single-flight value, thus addressing delay propagation.

To see how the delay-propagation responsibility is apportioned between flight 2's buffer and the ground buffer, it is instructive to first consider the unrealistic case where $c_g = 0$, with ground time being costless. When $c_g = 0$, $F_1(b_1 + b_g) = 1$ must hold at the optimum, so that the probability of late departure for flight 2 (which requires $\epsilon_1 > b_1 + b_g$) equals zero. In effect, the lower limit of integration in (19) must be above the upper limit in order to make the integral zero. To establish this point, suppose to the contrary that $F_1(b_1 + b_g) < 1$ is satisfied along with the first-order conditions. Since $F_2(b_1 + b_2 + b_g - \epsilon_1) \leq F_2(b_2)$ holds over the range of integration in (19), which is nonempty given the maintained assumption, it follows that

$$\begin{aligned} \frac{\partial\Omega}{\partial b_g} &< (x+y) \left[F_2(b_2) - \frac{x}{x+y} \right] \int_{\epsilon_1 = b_1 + b_g}^{\bar{\epsilon}_1} f_1(\epsilon_1) d\epsilon_1 \\ &= (x+y) \left[F_2(b_2) - \frac{x}{x+y} \right] \left[1 - F_1(b_1 + b_g) \right] < 0, \end{aligned} \tag{22}$$

where the last inequality follows because $F_2(b_2) - [x/(x+y)] < 0$ by (21) (using $\partial\Omega/\partial b_g = 0$ from (19)).

The inequalities in (22) contradict the assumption that (19) equals zero, ruling out the premise that $F_1(b_1 + b_g) < 1$. Therefore, the ground buffer b_g is set at a value large enough to eliminate the chance of late departure for flight 2, so that $F_1(b_1 + b_g) = 1$.⁷

Setting $F_1(b_1 + b_g)$ in (21) equal to 1, (21) matches the optimality condition (3) for the single-flight model when $c_g = 0$, implying that the optimal b_2 equals the single-flight value. Thus, when $c_g = 0$, the flight buffer for flight 2 is set as if the flight were operating in isolation, with delay propagation not an issue. Since the same conclusion has already been established for flight 1, we can state.

Proposition 2. *When ground time is costless, the ground buffer does all the work in mitigating delay propagation, fully eliminating it, with no contribution from flight 2's buffer.*

Even though this result is a natural conclusion, its derivation is not entirely straightforward, as seen above. Now consider the realistic case where the cost of ground time is positive, with $c_g > 0$. Setting (19) equal to zero now implies that the integral must be negative, which means that the lower limit of integration cannot exceed the upper limit as before, leading to a zero value for the integral. In other words, $F(b_1 + b_g)$ must now be less than rather than equal to 1, indicating that there is a chance of late departure for flight 2.

In this case, the mitigation of delay propagation is apportioned between the ground buffer and flight 2's buffer, with the exact apportionment depending on the relationship between c_f and c_g . Suppose first that $c_f < c_g$, so that the cost of the flight buffer is less than that of the ground buffer. Letting ** denote optimal values in the two flight model, it follows from (20) that $F_2(b_2^{**}) - x/(x+y) > 0$ must hold. Since $F_1(b_1^{**} + b_g^{**}) < 1$, (21) and $F_2(b_2^{**}) - x/(x+y) > 0$ imply

$$\begin{aligned} 0 &= (x+y)F_1(b_1^{**} + b_g^{**}) \left[F_2(b_2^{**}) - \frac{x}{x+y} \right] + c_f - c_g \\ &< (x+y) \left[F_2(b_2^{**}) - \frac{x}{x+y} \right] + c_f - c_g \end{aligned} \tag{23}$$

(recall $b_1^{**} = b_1^*$). With the last line of (23) thus positive, it follows that b_2^{**} is larger than b_2^* , the single-flight value of b_2 , which makes the second line equal to zero. Since flight 2's buffer is thus larger than the value it would assume if the flight were operating in isolation, it follows that flight 2's buffer assists the ground buffer in addressing delay propagation. This conclusion is due to the relative

⁷ It should be noted that $F_1(b_1 + b_g) = 1$ does not yield a unique solution for b_g given that any b_g value satisfying $b_1 + b_g \geq \bar{\epsilon}_1$ makes the equality true. However, replacing b_1 by b_1^* , the airline might be assumed to choose the smallest b_g satisfying the equality, so that $b_1^* + b_g = \bar{\epsilon}_1$, yielding a unique solution given by $b_g^* = \bar{\epsilon}_1 - b_1^*$.

cheapness of the flight buffer, which encourages its use in addressing propagation.

In the reverse case where $c_f > c_g$, reversal of the above argument yields $b_2^{**} < b_2^*$, so that flight 2's buffer is less than its single-flight value. Now, relative cheapness of the ground buffer means that it takes extra responsibility in addressing delay propagation, causing flight 2's buffer to be reduced below its single-flight value. This adjustment means that the flight buffer now actually contributes to delay propagation, but this effect is offset by the greater role of the ground buffer. Finally, when $c_f = c_g$, flight 2's buffer equals its single-flight value ($b_2^{**} = b_2^*$), so that it neither helps nor offsets the ground buffer in mitigating delay propagation. With the buffers equally costly, the nondistorting ground buffer is thus set to do all the work in addressing delay propagation, although the chance of propagation is not reduced to zero given the costliness of the buffer. Summarizing yields.

Proposition 3. *When $c_g > 0$, mitigation of delay propagation is apportioned between the ground buffer and flight 2's buffer. When $c_f = c_g$, the ground buffer alone addresses delay propagation (while not fully eliminating it), with flight 2's buffer set at its single-flight value. When $c_f < c_g$, flight 2's buffer contributes to (partly offsets) the ground buffer's mitigation of delay propagation, taking a value above (below) its single-flight value.*

The results yield a further implication in the case where the distributions of the random flight-duration terms are equal, allowing the 1 and 2 subscripts to be removed from the F functions. With common F 's, the single-flight values of b_1 and b_2 are the same. Then, the fact that b_2 is greater than (less than) the common single-flight value as $c_f < (>)c_g$ means that flight 2's buffer is greater than (less than) flight 1's buffer, which always equals the common single-flight value, when $c_f < (>)c_g$. In other words, $b_2^{**} > (<)b_1^{**}$ holds as $c_f < (>)c_g$. The reason is that flight 1's buffer plays no role in addressing delay propagation, while flight 2's buffer contributes to (borrows from) the ground buffer's delay-mitigation effect as $c_f < (>)c_g$.

The second-order conditions for the optimization problem have not been mentioned so far. In the appendix, it is shown that, if $c_f \geq c_g$, then Ω is strictly convex in b_1 and b_2 at the optimum, so that conditional on b_g , flight buffers satisfying the first-order conditions based on (20) and (21) yield a local minimum of Ω . However, it is not possible to establish convexity of Ω in all three buffers, a condition that must be assumed to hold.

With b_1 always equal to its single-flight value, comparative-static results for this flight buffer are given by (5). Comparative statics for b_2 and b_g are generally unavailable, however, because they require total differentiation of the equation system (19)–(21), which leads to ambiguous results.

If the current model were extended to a case where the aircraft provides three or more flights per day, greater complexity would rule out derivation of analytical results, with reliance on numerical simulation required instead. Such analysis, however, would undoubtedly reproduce one of the main current results by showing that first flight's buffer equals its single-flight value. It is hard to predict, however, what the rest of the numerical analysis would show.

3.4. The effects of adding or removing excess capacity

As seen above, the single-flight model assumes that excess capacity is present, while it is absent in the two-flight model. These assumptions affect the comparison between flight buffers in the two models, making it useful to gauge the effect of alternate assumptions.

First, suppose that excess capacity is present in the two flight model, in which case the previous airline cost expression is replaced by $c_f b_1 + c_f b_2 + c_g(T - b_1 - b_2)$, where T is the total time covered by the aircraft lease. Ground time now includes the ground buffer and additional ground time following the termination of flight 2. It is easy to see that, following this change, c_g disappears from the b_g first-order condition (19), whereas c_f in (20) and (21) is replaced by $c_f - c_g$. Arguments parallel to those leading to Proposition 2 can then be applied, yielding the same conclusions as in the proposition, namely, sole usage of the ground buffer in addressing delay propagation. Intuitively, when it comes to choice of the ground buffer, excess capacity is just like a zero c_g , accounting for the identical conclusions. Note, however, that a corner solution emerges if T is not large enough to accommodate a sufficiently long ground buffer.⁸

Alternatively, suppose that excess capacity is absent in both the single- and two-flight models. Then, the single-flight cost expression becomes $c_f b_1$, and $-c_g$ vanishes from (3). This change makes the expression in (3) larger, which requires a decline in the value of b_1 in order to maintain the zero equality. With the single-flight value falling and the flight buffers in the two-flight model unchanged, $b_1^{**} > b_1^*$ now holds, while b_2^{**} could be larger or smaller than b_2^* when $c_f > c_g$. Being larger than the single flight-value, flight 1's buffer therefore now assists the ground buffer in mitigating delay propagation.

While these conclusion are not unreasonable, we believe that the results in Propositions 1–3, which are based on excess capacity in the single-flight model, are more natural and intuitive. Moreover, when only one flight is operated, the existence of excess capacity is a plausible assumption.

4. Numerical examples

Figs. 2–5 present numerical examples. It should be noted that, given the stylized nature of the model, realism in the choice of

⁸ This discussion shows that, as in the $c_g = 0$ case, unrealistic conclusions make excess capacity an undesirable assumption for the two flight-model.

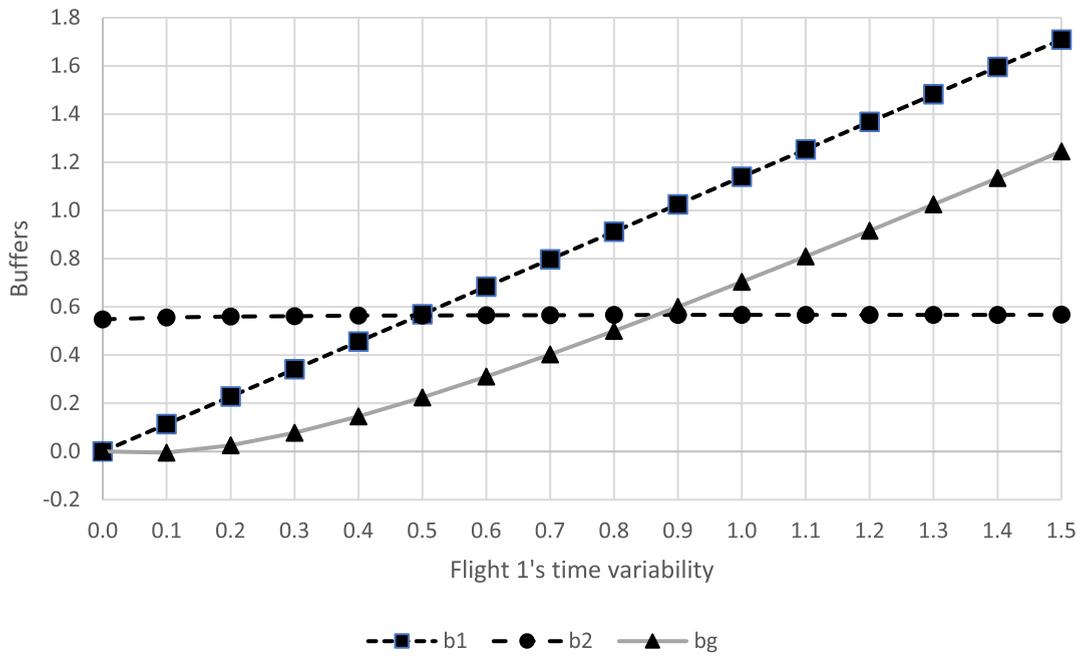


Fig. 2. Effect of flight 1's time variability on buffers.

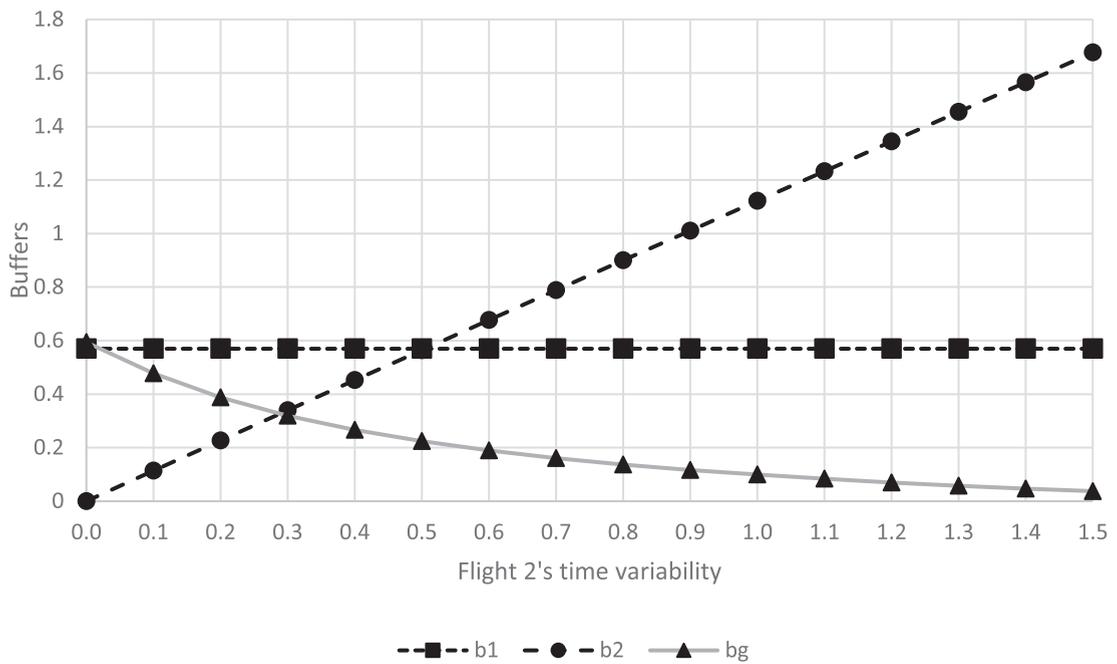


Fig. 3. Effect of flight 2's time variability on buffers.

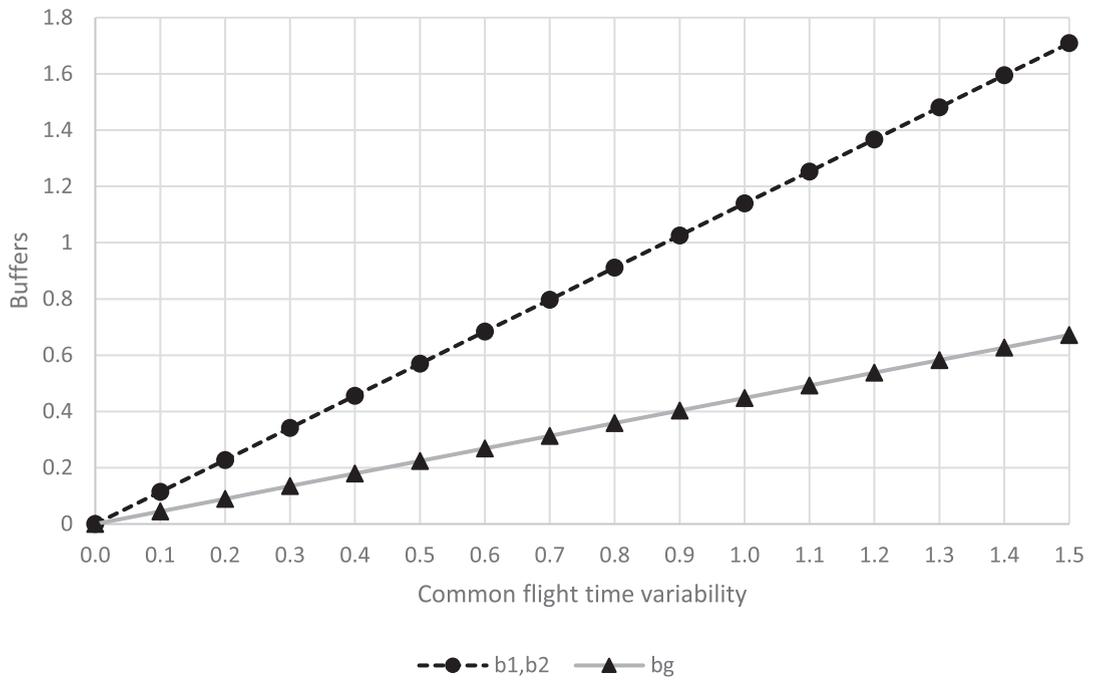


Fig. 4. Effect of common flight time variability on buffers.

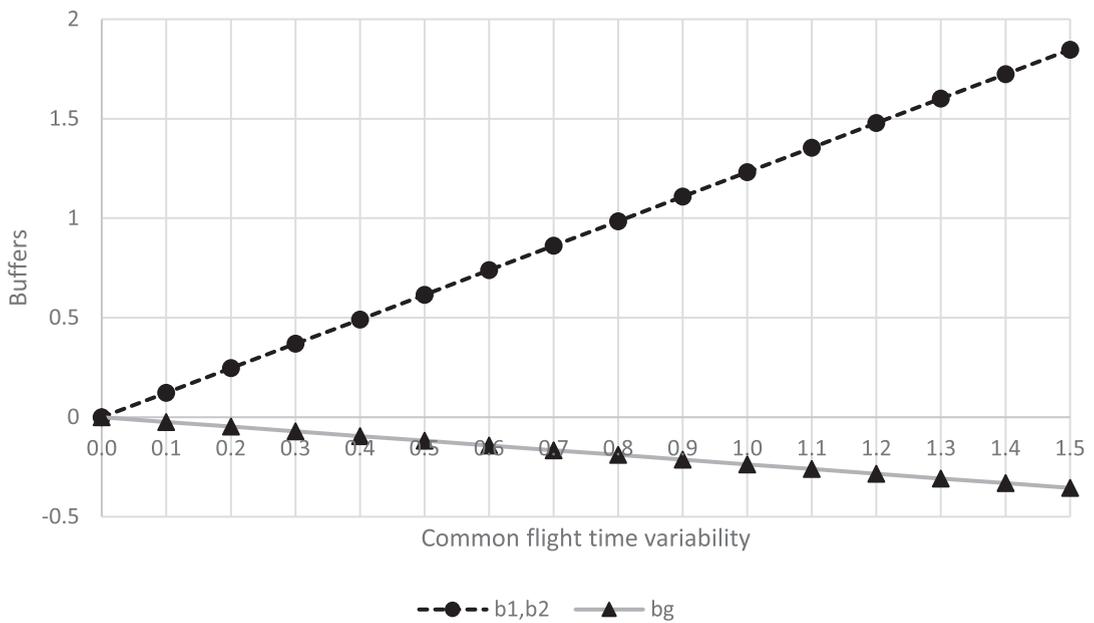


Fig. 5. Effect of common flight time variability on buffers (higher cg).

parameter values is not possible, and the qualitative (rather than quantitative) effects of parameters changes are of interest.⁹

Since some of the comparative-static effects of x and y and of c_f and c_g have been derived analytically, the analysis focuses on the effects of greater variability in the random terms ϵ_1 and ϵ_2 . To generate the results, ϵ_1 and ϵ_2 are assumed to follow independent normal distributions with a positive mean μ . Their standard deviations start out equal, satisfying $\sigma_1 = \sigma_2 = 0.0$. Then, each of the σ 's increases up to 1.5 holding the other σ fixed, allowing the effect of greater flight-time variability for flights 1 and 2 to be appraised separately. Next, the σ 's are set at a common value and increased simultaneously from 0.0 up to 1.5, allowing the effects of greater overall flight-time variability to be appraised.

Among the other parameters, the minimum turnaround time \bar{t}_g is set at 0.5 h, and the lateness and earliness disutilities are set at $x = 1.0$ and $y = 0.1$, with y realistically much smaller than x . The buffer costs are initially set at $c_f = 0.05$ and $c_g = 0.01$.

Recalling that the ϵ 's are assumed to be positive random variables, this property can be assured by choosing the normal mean μ to be sufficiently large. With σ ranging from zero up 1.5, the probability of a negative ϵ is smaller than 0.025 when the mean is at least twice the maximum σ , or at least 3.0, and its value is assumed to lie far enough above 3 to yield a negligible probability. Furthermore, inspection of the first-order conditions (19)–(21) shows that raising the value of the mean by $\Delta\mu$ (which shifts the distribution function F to the right by this amount) increases both flight buffer solutions by $\Delta\mu$ while leaving the ground buffer solution unchanged. Thus, the flight buffers move exactly in step with μ . This fact is used to facilitate presentation of the following figures, allowing flight and ground buffers to be shown on the same scale. In particular, the flight-buffer values shown are based on a μ value of zero, and they must be incremented by the amount of the actual μ (which is above 3) to get the actual solutions. Since σ 's effect on the buffers is our main focus, this need for rescaling is inconsequential.¹⁰

Fig. 2 shows the effect of increasing σ_1 from 0.0 to 1.5 with σ_2 fixed at 0.5. As can be seen, the flight buffer b_1 rises as σ_1 increases, a natural finding, while the ground buffer b_g also rises. The buffer b_2 for flight 2 appears to be constant in the figure, but it increases very slightly with σ_1 . The conclusion, therefore, is that b_1 and b_g alone do almost all the work in absorbing the greater chance of an arrival delay and subsequent delay propagation that follows from an increase in flight-time variability for flight 1.

Fig. 3 shows the effect of increasing σ_2 from 0.0 to 1.5 with σ_1 fixed at 0.5. Now b_2 rises, while b_1 is constant (note that the b_1 solution from (25) is independent of σ_2). However, the ground buffer b_g is decreasing in σ_2 , apparently because greater flight-time variability for flight 2 makes the ground buffer less effective at preventing a late arrival.¹¹

Fig. 4 shows the effect of simultaneously increasing σ_1 and σ_2 from 0.0 to 1.5. Both flight buffers naturally increase with the common σ value, and although the figure makes the buffers look equal in size, b_1 is slightly larger than b_2 , as predicted. In addition, b_g is increasing in the common σ value. However, Fig. 5 shows that the behavior of the ground buffer is reversed when $c_f = 0.05$ and $c_g = 0.03$, falling as the common σ value increases (matching the outcome in Fig. 2). Fig. 5 shows an additional point that does not arise in the other cases. In particular, b_g becomes negative as σ increases, showing that the ground time is set below the minimum feasible turnaround time \bar{t}_g . Nothing in the model prevents this outcome, which need not lead to late departure for flight 2 if flight 1's buffer is sufficiently large. In the data discussed below, however, the outcome is exceedingly rare, accounting for only 0.2% of the observations.

The implication is that the effect of flight-time variability on the ground buffer could be positive or negative depending on the magnitudes of the other parameters. If the ground buffer is sufficiently cheap compared to the flight buffers ($c_g = 0.01$ vs. $c_f = 0.05$), it is used along with the flight buffers to address the greater threat of delay propagation resulting from higher flight-time variability (Fig. 4). But when the ground-buffer cost is larger as a proportion of c_f ($c_g = 0.03$ vs. $c_f = 0.08$), then the flight buffers partly supplant the ground buffer as the threat of delay propagation rises, with b_g falling (Fig. 5).

5. Adding passenger connections

So far the analysis has suppressed the possibility that some passengers connect from flight 1 to flight 2. These passengers would travel from the origin city of flight 1 to flight 2's destination city, making a connecting trip in the absence of nonstop service between the two cities. This type of connecting travel, however, has little effect on the model. Assuming that a share α of passengers on both flights are traveling nonstop while $1 - \alpha$ are connecting, the airline's objective function would be altered in a straightforward way. In (18), Ω would be multiplied by α and then added to the term $(1 - \alpha)(\Omega_{2,late} + \Omega_{2,early})$, which represents the late and early disutilities for connecting passengers, who care only about their arrival time at flight 2's destination. Adding the two expressions, the disutility portion of the objective function reduces to (18) with α multiplying Ω_1 .

A more interesting and complex connecting scenario arises if two additional flights, again using a single aircraft, are added to the model. The turnaround city for these flights, which are denoted 1B and 2B, is the same as for the original flights, now denoted 1A and 2A. This common turnaround city can thus be viewed as a hub, which is a destination in its own right but also supports passenger

⁹ He (2019) uses such numerical exercises involving two flight legs and two aircraft to illustrate an integrated approach for overbooking and capacity planning.

¹⁰ For computational reasons, the infinite upper and lower limits of the normal distributions are replaced by 10 and -10 respectively in the calculations. With σ 's taking the values mentioned above, the probability that ϵ_1 or ϵ_2 lies outside this range is virtually zero, making the restriction inconsequential.

¹¹ Alternatively, recall that, conditional on delay propagation, the second flight's arrival delay depends on the sum of buffer times $b_1 + b_2 + b_g$ (see Table 1). If the second flight's buffer increases because of an increase in σ_1 , the result is an increase in the sum of buffer times. The ground buffer is then reduced to moderate the increase in this sum.

connections. Connecting passengers now include those traveling from flight 1A's origin (O_A) to flight 2B's destination (D_B) as well as those traveling from O_B (flight 1B's origin) to D_A (flight 1A's destination).¹²

If flight 1A arrives after the departure of flight 2B, passengers traveling from O_A to D_B miss their connection, suffering disutility V , with same conclusion applying to passengers connecting from flight 1B to flight 2A. Note that since connecting passengers using flights 1A and 2A (or 1B and 2B) do not change planes, a missed connection is not possible for them.¹³

The portion of the airline's objective function applying to nonstop trips and same-plane connecting trips is given by adding the Ω 's for the A and B flights, with the previous α modification incorporated.¹⁴ The part of the objective function that applies to the remaining connecting passengers makes use of the probability of a missed connection. For 1A-2B connections, this probability is $P_{AB} \equiv Pr(m_1 + \epsilon_{1A} > \hat{t}_{d2B})$, using (7) and adding an A subscript, while for 1B-2A connections, the probability is given by the analogous expression P_{BA} . In the previous expression, $m_1 + \epsilon_{1A}$ is the arrival time of flight 1, and a missed connection occurs when it is greater than flight 2's departure time.

Using P_{AB} and P_{BA} , the remaining (disutility) part of the objective function is given by $1 - \alpha$ times

$$P_{AB}V + (1 - P_{AB})(\Omega_{2B,late} + \Omega_{2B,early}) + P_{BA}V + (1 - P_{BA})(\Omega_{2A,late} + \Omega_{2A,early}). \quad (24)$$

The first and third terms in (24) give disutilities from missed connections, while disutilities for connecting passengers who make, rather than miss, their connection are given by the second and fourth terms. Note that these early-late disutilities are the same as for nonstop passengers on either flight 2A or 2B, being given by the terms multiplying $1 - P_{AB}$ and $1 - P_{BA}$.

These second and fourth multiplicative terms in (24) make the objective function considerably more complex than for a model with only two flights. Given the challenging nature of resulting analysis, pursuing it is beyond the scope of the paper. Intuitively, however, avoidance of missed flight connections provides an additional reason beyond mitigation of delay propagation to increase the flight buffers for flights 1A and 1B as well as the corresponding ground buffers. Despite the absence of concrete conclusions beyond this simple intuition, it is interesting nevertheless to see the logic under which flight connections can be added to the model.¹⁵

6. Empirical setup

6.1. Predictions

The empirical work aims to test some of the predictions of the theoretical model, using US data that track the flights of individual commercial aircraft for 2018. The data allow computation of the flight buffer for a particular flight, which is set equal to scheduled flight time minus the minimum observed flight time on the route, matching the model. The data also allow computation of the ground buffer, which equals scheduled ground time minus the minimum observed ground time at the airport (details on both calculations are presented below).

One of the model's predictions, which comes from the simulation analysis of the two-flight case, is that the flight buffer increases when a flight's own time variability rises, but that the buffer is unaffected by higher time variability for the other flight (Figs. 2 and 3). Also, the ground buffer is increasing (decreasing) in the time variability of the previous (subsequent) flight. To test these predictions, flight-time variability is measured by the standard deviation of actual flight times at the flight, route and month level. Later regressions replace this variability measure with proxies for variability.

A second prediction is that buffer costs affect the levels of all three buffers, although the effect is unambiguous only in the case of b_1 , which is decreasing in c_f and increasing in c_g . We use airport characteristics (mainly hub status) as a proxy for c_g and aircraft size as a proxy for c_f .

A third prediction is that the buffers depend on the disutilities x and y for lateness and earliness, with unambiguous predictions only for available for b_1 , which is increasing in x and decreasing in y when c_f and c_g are close in magnitude. We use the managerial employment share of the origin city as a proxy for x on the belief that such employees are highly averse to lateness (Luttman (2019) also uses this variable).

A fourth prediction is that the position of a flight in the day's flight sequence affects the flight buffer. Recall that flight 2 has a longer buffer than flight 1 when $c_f < c_g$, with the relationship reversed when $c_f > c_g$. To test for such a sequencing effect, we include time-of-day departure variables (morning, evening, etc.) along with variables that measure a flight's exact position in the sequence of an aircraft's daily flights.

A fifth prediction is that the share of connecting passengers affects the levels of the buffers. No explicit predictions are available, but we expect the ground buffer to increase as the connecting share on the subsequent flight rises. The connecting share is computed at the

¹² Symmetry holds, with corresponding flight distances equal and with flights 1A and 1B departing at a common time.

¹³ Note that, while crew members traveling as passengers to their next flight assignment can also miss connections, with serious flight-delay consequences, analysis of this phenomenon would be complex.

¹⁴ The relevant expression is $\alpha\Omega_{1A} + \Omega_{2A,late} + \Omega_{2A,early} + \alpha\Omega_{1B} + \Omega_{2B,late} + \Omega_{2B,early}$. Note that the buffers inside the Ω expressions also acquire A and B subscripts, although their equilibrium values will be the same given symmetry.

¹⁵ A different model exploring missed connections would proceed as follows, returning to the two-flight framework but assuming the flights use different aircraft (so that delay propagation is absent). If flight 1 arrives after the departure of flight 2, then connecting passengers miss their connection, again generating disutility V . Now, there is no ground buffer per se, but the airline chooses flight 2's scheduled departure with an eye toward avoiding missed connections, assuming that a later departure imposes the same costs as a longer ground buffer.

route-airline-quarter level for the prior year using U.S. Department of Transportation (DOT) data.

6.2. Data collection and buffer measurement

The sample is obtained from the DOT and covers the US domestic airline market for the year 2018. We mainly rely on the ‘Marketing Carrier On-Time Performance’ dataset, which for each aircraft, uniquely identified by its tail number, contains information on the carrier operating the flight, the origin and destination, the departure date, the scheduled departure and arrival times, the actual departure and arrival times, and the taxi-in and taxi-out times.¹⁶ We exclude international flights, flights that are canceled or diverted, and flights from/to US Commonwealth areas and Territories.

The aim of our empirical analysis is to investigate the determinants of the two choice variables described in our theoretical model: the flight buffer and the ground buffer. The flight buffer is obtained as follows. First, for any flight i , where i defines the sequence of daily flights operated by a given aircraft during the day, we measure the actual flight time, which is the sum of the taxi-out, airborne and taxi-in times (equivalent to the so-called ‘block-time’). Using the model notation, this actual flight time is given by $\hat{m}_i = \hat{t}_{ai} - \hat{t}_{di}$, the difference between flight i ’s actual arrival and departure times. We then take the smallest value of \hat{m}_i across all aircraft of flight i ’s type flying the same route to obtain \hat{m}_{min} , the minimum flight time by route and aircraft type (the route and type subscripts are suppressed for simplicity). The flight buffer is given by the difference between flight i ’s scheduled flight time m_i and the minimum flight time of the same aircraft type flying the same route: $b_i = m_i - \hat{m}_{min}$.¹⁷

To calculate the ground buffer, we first compute the actual ground time that separates flight k and flight $k+1$ in the sequence of flights operated by the observed aircraft during the day: $\hat{t}_{kg} = \hat{t}_{d,k+1} - \hat{t}_{ak}$. Then we calculate the minimum turnaround time as the shortest actual ground time observed across all aircraft of a given type using the turnaround airport. The ground buffer for an aircraft turnaround is obtained by subtracting this minimum feasible ground time from the scheduled ground time.¹⁸ In order to have confidence in the accuracy of the observed minimum flight and ground times used in the buffer computations, we require that the number of observations used to generate them is at least equal to 30.

We initially restrict the analysis to those aircraft that only operate two flights a day, as in the model, using what we refer to as the ‘two-flight sample’. This restriction predominantly limits the focus to coast-to-coast flights or flights to/from Alaska. However, we later present results using an expanded sample that includes aircraft flying as many as 8 flights in a day, referred to as the ‘unrestricted sample.’.

6.3. Empirical approach

Our empirical analysis is based on two sets of regressions: one for flight buffers and another for ground buffers. The standard errors are clustered by route and month in order to allow the residuals of different aircraft (possibly of different carriers) flying on the same route, during the same month, to be correlated.

The general equation to be estimated is:

$$Buffer_{jcodt} = X_{jcodt}\beta + \eta_c + \phi_o + \rho_d + \lambda_t + u_{jcodt} \quad (25)$$

where *Buffer* is either the flight buffer or the ground buffer. The subscript j identifies the aircraft tail number, c the carrier, o the origin airport, d the destination airport (or the turnaround airport in the ground-buffer regression), and t the month. The η_c , ϕ_o , ρ_d , and λ_t terms are carrier, origin, destination, and month fixed effects¹⁹ and u_{jcodt} is the regression error, assumed i.i.d. with zero mean. The

¹⁶ The novelty of this dataset relative to the ‘Reporting Carrier On-Time Performance’ dataset, previously used in Forbes (2008) and other studies, is that it distinguishes whether the flight is operated by the reporting carrier or by its regional codeshare partners. Both datasets are downloadable at https://www.transtats.bts.gov/DatabaseInfo.asp?QO_VQ=EFD&Yv0x=D. However, the Marketing Carrier On-Time Performance dataset only starts with January 2018, while the Reporting Carrier On-Time Performance dates back to 1987.

¹⁷ Changing this calculation so that the minimum flight is computed by route-carrier-quarter-airline instead of by route-carrier has little effect on the results presented below. Our use of the minimum flight time follows Mayer and Sinai’s (2003) use of the minimum to compute flight delays. Other papers instead calculate delays using a low-percentile flight duration such as the 5th, 10th or 20th percentile (see Britto et al. (2012) and Zou and Hansen (2012)).

¹⁸ Since the Marketing Carrier On-Time Performance dataset does not include international flights, except those from/to US Commonwealth areas and Territories, we restrict the ground buffer to be not greater than 200 min, to exclude possible incomplete records that may occur when an aircraft flies from/to a non-domestic destination (e.g. Canada or Mexico) between two domestic flights.

¹⁹ The airport fixed effects are not only relevant in the flight-buffer regression, but also in the ground-buffer regression. Consider for example two flights that share the same turnaround airport and aircraft type, but that originate from different airports. These flights have the same minimum actual ground time, but they may have different load factors depending on the strength of demand and other forces. The potential difference in the load factor may affect the turnaround time and hence the ground buffer set by the airline. The inclusion of originating airport fixed effects (the origin airport of the previous flight) aims to control for this unobservable characteristic.

Table 2
Description and main statistics of the empirical variables.

b	b _g	Variables	Description	Two-flight	Unrestricted
				sub-sample	sample
✓		Aircraft rotation	Sequence of flights operated in a day by a given aircraft		2.99 (1.69)
✓	✓	Alaska-Southwest Competitors	Set of airline dummy variables; <i>American Airlines</i> is the omitted airline Number of competitors on the route	1.58 (1.21)	1.06 (1.11)
✓		Congestion origin	Number of landing and departing flights at the airport of origin in the same hour when the flight is scheduled to depart divided by the number of runways of the airport of origin	13.12 (7.08)	11.64 (7.64)
✓		Congestion destination	Number of landing/departing flights at the airport of destination in the same hour when the flight is scheduled to land divided by the number of runways of the airport of destination	12.23 (6.81)	11.39 (7.72)
✓	✓	Congestion turnaround	Number of landing and departing flights at the airport of turnaround in the same hour when incoming flight is scheduled to land divided by the number of runways of the airport of turnaround	12.42 (7.14)	12.42 (7.72)
✓	✓	Connecting passengers	Share of past-year connecting passengers, proxy based on DB1B Market dataset	0.31 (0.24)	0.45 (0.28)
✓		Distance	Route distance, in 100-mile units	16.94 (8.66)	8.21 (5.92)
✓	✓	February-December	Set of monthly dummy variables, <i>January</i> is the omitted month		
✓		Flight <i>i</i>	Dummy = 1 for the <i>i</i> th flight operated in a day by a given aircraft, uniquely identified by its tail number		
✓		Flight buffer (<i>b</i>)	Difference between the scheduled flight time and the minimum actual flight time, computed by route and aircraft type	39.95 (15.87)	31.73 (12.89)
✓	✓	Flight 1's variability	Flight 1's standard deviation of flight-time	14.31 (6.53)	
✓	✓	Flight 2's variability	Flight 2's standard deviation of flight-time	14.54 (7.01)	
✓	✓	Ground buffer (<i>b_g</i>)	Difference between the scheduled ground time and minimum feasible ground time, computed by turnaround airport and aircraft type	40.30 (30.46)	37.88 (25.96)
✓	✓	Heavy aircraft	Dummy variable = 1 if the aircraft has been assigned a maximum takeoff weight rating of 300,000 lb or more	0.23 (0.42)	0.03 (0.17)
✓		Hub origin	Dummy variable = 1 if airport of origin is the hub of the airline	0.51 (0.50)	0.36 (0.48)
✓		Hub destination	Dummy variable = 1 if airport of destination is the hub of the airline	0.51 (0.50)	0.36 (0.48)
✓		Hub turnaround	Dummy variable = 1 if airport of turnaround is the hub of the airline	0.45 (0.49)	0.38 (0.49)
✓		Low-cost carrier	Dummy variable = 1 if the flight is operated by a low-cost carrier (i.e. Jet Blue, Frontier Airlines, Allegiant Air, Spirit Airlines and Southwest Airlines)	0.16 (0.37)	0.29 (0.45)
✓		Managerial origin	Managerial share of the origin city's work-force	5.69 (1.30)	
✓	✓	Managerial turnaround	Managerial share of the turnaround city's work-force	5.65 (1.28)	
✓	✓	Morning-Evening	Set of departure time variables, <i>Morning</i> (9.00–11.59), <i>Afternoon</i> (12.00–15.59), <i>Late afternoon</i> (16.00–17.59) and <i>Evening</i> (18.00–23.59); the omitted category is <i>Early morning</i> (0.00–8.59)		
✓		Past-year delay	Share of past-year delayed flights, computed by carrier-route-month	0.20 (0.11)	0.17 (0.11)
✓	✓	Regional carrier	Dummy variable = 1 if the flight is operated by a regional carrier	0.07 (0.26)	0.31 (0.47)
✓		Slot controlled airport	Dummy variable = 1 if airport of turnaround is slot controlled (DCA, JFC and LGA)	0.07 (0.26)	0.05 (0.22)
✓	✓	Weekend	Dummy variable = 1 if the flight departs in the weekend	0.31 (0.46)	0.27 (0.44)

(a) The symbol ✓ denotes whether the variable is included in the flight buffer regressions (column *b*) or in the ground buffer regressions (column *b_g*).
 (b) The second-last and last columns respectively report the mean value of the variable for the two-flight sub-sample and unrestricted sample with the standard errors in parentheses. Empty cells on some variables are due to space reasons, since their inclusion would require reporting a set of multiple dummies whose mean values would not be very informative.

other independent variables are denoted by X , and they control for different aspects of the buffer decision.²⁰

Table 2 shows definitions and summary statistics for the variables in the two-flight and unrestricted samples. In addition to the buffer, time-variability, and connecting-share variables discussed above, the table shows the month and airline dummies from (25). Other variables included in the set of X variables from (25) are hub origin and hub destination dummies (replaced by a hub turnaround variable for the ground-buffer regressions), congestion measures for the origin and destination (or turnaround) airports (equal to traffic divided by runway capacity),²¹ a slot-control dummy for the turnaround airport, dummies for regional and low-cost carriers, a distance measure, time-of-day and weekend dummies, two variables indicating a flight's position in the aircraft's daily sequence, a large-aircraft dummy, and the managerial-share variable (for the origin or turnaround airport).

In the unrestricted sample, the mean flight buffer is about 32 min, indicating that carriers set scheduled flight times to be 32 min longer on average than the minimum flight time for that route and aircraft type. The mean ground buffer shows that airlines set ground times to be 38 min above minimum turnaround times. The mean connecting share is 45%,²² 36% of flights have a hub origin or destination, 31% of flights are on regional carriers, 29% are on low-cost carriers, routes have one competitor on average, average distance is a bit over 800 miles, 27% of flights are on weekends, 3% of flights use heavy aircraft, and the managerial share of the origin work force averages about 5%. Differences in means between the samples are mostly as expected (for example, mean distance more than doubles to nearly 1700 miles moving to the two-flight sample).

7. Regression results

7.1. Flight-buffer regressions with flight-time variability

The results of the flight-buffer regressions including flight-time variability are shown in Table 3. Many of the variables listed in Table 2 are omitted from these regressions (to be included later) on the grounds that they may be proxies for flight-time variability, thus being excludable given that variability is explicitly measured. In addition to other variables in Table 2, the excluded variables include the month dummies and the origin and destination fixed effects, all of which are likely determinants of flight-time variability. As mentioned above, the variability measure used in the regressions is the standard deviation of the actual flight times, computed by route, month and flight number.

In order to match the setup of our theoretical model, the regressions in Table 3 are based on a subsample of aircraft that operate only two flights a day (mainly flying coast-to-coast or to Alaska). We refer to this sample as 'two-flight sample', and it is expanded in later regressions to include aircraft operating more flights per day. In the table, column (1) shows the regression with the flight buffer of flight 1 set as dependent variable, column (2) has flight 2's flight buffer as dependent variable, and columns (3) and (4) pool the two flights, with (4) adding the dummy variable *Flight 2* to identify the second flight. Each regression includes flight variability for both flights, following the simulations.

The time-variability coefficients are positive and significant in all four regressions. Columns (1) and (2) show that an increase in either flight's time variability leads an airline to raise the flight buffer for both that flight and the other flight. This finding does not perfectly match the simulations, which showed that increasing σ for just a single flight has a small or zero effect on the other flight's buffer. Columns (3) and (4) show that when the flights are pooled, time variability for both flights affects the buffers for either flight, yielding almost the same conclusion as columns (1) and (2). Despite this imperfect correspondence, the findings confirm the overall spirit of the numerical examples, which show that greater time variability tends to raise flight buffers.

The share of connecting passengers (measured for flight 2) has a significantly negative effect on the flight buffers in all four regressions, an unexpected result. A more intuitive finding emerges, however, in the ground-buffer regressions presented next. The dummies indicating a hub origin or destination are mainly meant to capture the cost of ground time, with gate rentals likely to be more costly at hub airports given their large sizes. The hub coefficients are uniformly positive and significant, indicating that flights in or out of hubs have longer flight buffers, an effect that may also reflect aversion to possible disruption of flight connections.

The coefficients of the carrier dummies show long flight buffers for Delta and Virgin America (relative to American, the default carrier) and short ones for Allegiant, Frontier, Hawaiian, Jet Blue, Southwest, Spirit. While the buffers of some of the latter carriers are shorter than American's by more than 10 min, United's flight buffers are also seen to be shorter, but by less than two minutes. Regional airlines in the dataset have lower flight buffers (by about 7 min) relative to other carriers, possibly indicating that these airlines are

²⁰ Note that the regression in (25) bears some resemblance to scheduled-block-time regressions like those estimated by Kang and Hansen (2017), Kang and Hansen (2018) and others. With the flight buffer equal to scheduled-block-time minus minimum flight time, (25) would become a scheduled-block-time regression if minimum flight time were moved to the RHS and given a coefficient. Kang and Hansen's regressions, however, differ from this modified regression by not including minimum flight time as a covariate.

²¹ The data on number of runways for the sample airports are from the FAA's National Flight Data Center, available at https://www.faa.gov/air_traffic/flight_info/aeronav/aero_data/Airport_Data/. A better capacity measure would be maximum feasible flight operations per hour, but this measure is not available for all of the airports in our sample (see https://www.faa.gov/airports/planning_capacity/profiles/).

²² The connecting shares are computed using ticket data from the DOT's Airline Origin and Destination Survey (database DB1B), which comes from a quarterly 10% sample of all tickets. For each route-airline-quarter combination in 2017, the share of connecting passengers on the route can be deduced from the DB1B ticket information. This share is then matched to the route-airline-quarter observations in our data, thus indicating the average share of connecting passengers present for each of our observations. While some routes carry zero connecting passengers in actuality, other routes may show zero connections because the sampling process misses small numbers of passengers.

Table 3
Flight-buffer regressions with flight-time variability.

Dependent variable	(1)	(2)	(3)	(4)
	b_1	b_2	b_1, b_2	b_1, b_2
Flight 1's variability	0.556***	0.200***	0.377***	0.375***
Flight 2's variability	0.104***	0.385***	0.240***	0.243***
Connecting pax	-17.000***	-17.814***	-17.698***	-17.677***
Hub origin	2.198***	2.207***	2.237***	2.113***
Hub destination	2.097***	3.630***	2.892***	2.959***
Alaska Airlines	-0.849	1.372**	0.085	0.119
Allegiant Air	-14.970***	-11.735***	-13.615***	-13.481***
Delta Airlines	1.781***	2.080***	1.899***	1.916***
Frontier Airlines	-10.570***	-10.563***	-10.534***	-10.565***
Hawaiian Airlines	-0.802	-4.120***	-2.099***	-2.224***
Jet Blue	-3.490***	-4.081***	-3.902***	-3.920***
Southwest Airlines	-6.038***	-6.016***	-6.014***	-5.945***
Spirit Airlines	-16.030***	-13.608***	-14.927***	-14.960***
United Airlines	-0.792	-1.570***	-1.131***	-1.161***
Virgin America	2.885	6.003***	4.144**	4.230**
Regional carrier	-8.215***	-6.834***	-7.358***	-7.371***
Competitors	1.108***	1.312***	1.216***	1.217***
Managerial origin	1.370***	1.386***	1.367***	1.402***
Heavy aircraft	4.038***	2.696***	3.399***	3.392***
Morning			1.031***	
Afternoon			-0.762***	
Late Afternoon			0.742**	
Evening			-0.167	
Flight 2				-0.829***
Constant	24.907***	24.357***	24.810***	25.126***
R ²	0.306	0.354	0.319	0.318
Observations	113,873	104,502	218,375	218,375

(a) The estimated coefficients marked with ***, ** and * are statistically significant at, respectively the 1%, 5% and 10% level.
 (b) The standard errors, not reported to save space, are clustered by route-month.

Table 4
Ground-buffer regressions with flight-time variability.

Dependent variable	(1)	(2)	(3)	(4)
	b_g	b_g	b_g	b_g
Flight 1's variability	-0.034	-0.028	0.242***	0.209***
Flight 2's variability	-0.091***	-0.118***	-0.178***	-0.302***
Connecting pax	10.709***	9.520***	11.600***	9.728***
Hub turnaround	13.201***	12.698***	13.106***	12.642***
Alaska Airlines	8.258**		8.162***	
Allegiant Air	-13.489***		-12.382***	
Delta Airlines	2.112***		1.869***	
Frontier Airlines	4.403**		4.764**	
Hawaiian Airlines	25.918***		25.718***	
Jet Blue	-3.248***		-3.878***	
Southwest Airlines	-1.929**		-1.288	
Spirit Airlines	-2.438***		-2.412***	
United Airlines	9.995***		9.900***	
Virgin America	10.426***		10.452***	
Regional carrier	-1.070**		-0.539	
Low-cost carrier		-8.064***		-8.485***
Managerial turnaround	0.538***	0.900***	0.519***	0.909***
Heavy aircraft	-0.593	0.453	-1.010	0.058
Morning	4.546**	4.112***	4.127***	3.640***
Afternoon	11.897***	13.096***	11.335***	12.520***
Late Afternoon	12.587***	13.328***	12.166***	12.938***
Evening	14.591***	15.493***	13.997***	14.862***
Constant	15.275***	18.104***	13.027***	18.250***
R ²	0.135	0.109	0.136	0.111
Observations	124,131	124,131	124,131	124,131

(a) The estimated coefficients marked with ***, ** and * are statistically significant at, respectively the 1%, 5% and 10% level.
 (b) The standard errors, not reported to save space, are clustered by route-month.
 (c) Flight-time variability calculated by route, month and flight i in columns (1) and (2), by route and flight i in columns (3) and (4),

keen to maximize aircraft utilization. In addition, the uniformly positive and significant coefficients for the competitors variable suggest that more competition may heighten an airline's attention to on-time performance, leading to an increase in its flight buffers. An extra competitor increases flight buffers by over a minute. Routes with a high managerial share at the origin have longer buffers, possibly reflecting business passengers' aversion to late arrivals. Finally, flight buffers are unexpectedly longer for large aircraft, which have higher values of c_f . Column (4) of the table shows a slightly shorter buffer for flight 2 relative to flight 1, whereas column (3) shows a mixed pattern of time-of-day coefficients.²³

7.2. Ground-buffer regressions with flight-time variability

Table 4 presents the ground-buffer results with flight-time variability, again using the two-flight sample. Since the variability coefficients for the ground buffer are somewhat sensitive to how variability is computed, we use two approaches. In columns (1) and (2), variability is the 2017 standard deviation of flight times computed by route, month and flight number, as before, whereas variability is computed only by route and flight number in columns (3) and (4), not controlling for month.

Greater time variability for flight 2 reduces the ground buffer in all four regressions, matching the results in Fig. 3. The effect of flight 1's time variability is insignificant in the first two columns but significantly positive in columns (3) and (4), matching the results in Fig. 2. Thus, Table 4 shows that greater time variability for flight 2 (flight 1) reduces (raises or leaves constant) the ground buffer, in remarkably close correspondence to the simulation results.

Turning to the other covariates, the coefficients of the share of connecting passengers are uniformly positive and significant across the regressions, as expected. The point estimates indicate that a 0.1 increase in the connecting share on flight 2 raises the ground buffer by about 1 min. The hub turnaround variable also has a positive and significant coefficient, and since the connecting share is held constant, this variable should capture the cost of ground time, with a negative coefficient expected. The positive estimate thus suggests that the hub turnaround variable may also help to capture the presence of connecting passengers. Note that a hub turnaround raises the ground buffer by about 13 min.

Among the airline dummies, notable results are the long ground buffers of Virgin America, Hawaiian, United and Alaska, and the short buffers of Allegiant. The Southwest coefficient, while significantly negative, shows a buffer only 2 min shorter than the reference carrier (American). But when all the low-cost carriers are represented by a common dummy, the variable's significant coefficient shows ground buffers about 8 min shorter than non-LCCs. Thus, recalling the short LCC flight buffers from Tables 3 and 6, LCCs keep their buffers short both in the air and on the ground. Columns (1) and (3) show that regional carriers have slightly shorter ground buffers than non-regional airlines.

The coefficients of the heavy-aircraft variable are insignificant, a finding that is reasonable given that ground buffers are already type-specific. The managerial-share coefficient is now significantly positive, indicating that ground buffers are longer in turnaround cities with a high managerial share, in an attempt to secure an on-time departure for flight 2.

The time-of-day coefficients show that the ground buffer monotonically increases as the turnaround between the two flights occurs later in the day. This effect is not captured in the model, but it could indicate that, when the two-flight sequence starts later in the day, the effects of time-varying congestion or other factors make a longer ground buffer optimal.

7.3. The determinants of flight time variability

The next step in the empirical analysis is to run regressions where flight-time variability is replaced by a host of variables that help to determine variability. The regressions thus include the previous variables along with proxies for variability but not variability itself, providing a different and perhaps more revealing picture of the factors affecting flight and ground buffers. In this exercise, we present results for both the two-flight sample and a larger "unrestricted" sample described below.

Before showing these regressions, it is useful to explore the relationship between flight-time variability and the set of proxy variables. These variables are the month dummies, which control partly for weather conditions, the origin and destination congestion variables (which presumably influence variability), the distance variable, the weekend dummy, and the origin and destination dummies, whose coefficients are not reported.

The results are shown in Table 5 for the two-flight sample, using a separate regression for each flight. With t statistics five to ten times those of other covariates, distance has a very precisely measured positive impact, showing that longer flights naturally have greater time variability. The congestion coefficients are also mostly positive and significant, as expected, while weekend flights show low time variability. The month dummies somewhat surprisingly indicate that the good-weather months of July and August have higher flight-time variability than the Winter months. This outcome may be due to high summer travel volumes, which can generate delays from various sources.²⁴

²³ Note that inclusion of the time-of-day variables is inappropriate in the regressions of columns 1 and 2 because, for example, flight 2 cannot generally be a morning flight.

²⁴ It is worth pointing out that if flight-time variability is computed only by route and flight number, then the month dummies become statistically insignificant. This result is mainly technical: because the dependent variable is then invariant throughout the year, the month dummies do not explain much more than the constant of the regression would explain, making their coefficients insignificant.

Table 5
The determinants of flight-time variability.

Dependent variable	(1)	(2)
	Flight 1's variability	Flight 2's variability
February	-0.558**	-0.872***
March	-1.080***	-1.205***
April	0.604**	-0.503*
May	-0.297	0.183
June	-0.380	0.128
July	1.134***	0.921***
August	1.104***	1.633***
September	-0.681**	-0.115
October	-0.968***	-0.913***
November	-0.420	0.180
December	0.107	0.023
Congestion origin	0.025**	0.020
Congestion destination	0.036**	0.064***
Distance	0.135***	0.152***
Weekend	-0.109**	-0.119***
Constant	5.957*	-8.461*
R ²	0.283	0.281
Observations	114,011	113,195

(a) The estimated coefficients marked with ***, ** and * are statistically significant at, respectively the 1%, 5% and 10% level.

(b) The standard errors, not reported to save space, are clustered by route-month.

(c) All estimates include airport of origin and airport of destination fixed effects.

7.4. Flight-buffer regressions with variability proxies

The results of the flight-buffer regressions with variability proxies are presented in Table 6. The regressions in columns (1)–(4) are based on the two-flight sample and have the same structure as those in Table 3. The regressions in columns (5)–(7) remove the restriction to aircraft operating just two flights per day. The corresponding sample, which we refer to as the ‘unrestricted sample’, consists of aircraft that operate at most 8 flights in a day.²⁵ Other differences between the regression specifications are discussed below.

Where statistically significant, the monthly dummy coefficients have negative signs, indicating that compared to January, the reference month, flight buffers tend to be shorter in other months. These results point to a role for weather in influencing the flight-buffer decision. The estimated coefficients suggest that the size of the effect varies with expectations of bad weather: the flight buffer falls monotonically from March/April until September, which is a period of the year generally characterized by good weather conditions and therefore by fewer weather-related delays. The flight buffer then starts increasing monotonically, while still remaining below the January value, until the end of the year. The statistically insignificant coefficients on the February dummy and sometimes the March dummy mean that there is no difference in flight-buffer choices across the months of January, February and (partially) March. This result makes sense because January, February and March are winter months that bear the same bad-weather expectations and therefore the same delay concerns.

In the two-flight sample, the effect of weather (in columns (3) and (4)) yields as much as 6 min of reduced buffer length during the Spring and summer months. In the unrestricted sample, this effect is still present, but with a slightly smaller impact, at about 5 min maximum. It should be noted that these month effects cannot be attributed to flight-time variability given the results of Table 6. If weather effects on flight buffers operated only through flight-time variability, then with a positive buffer-variability link, the month coefficients in Table 6 would follow the pattern seen in Table 5. The contrary results in Table 6 thus appear to show that the effect of weather may not operate entirely through flight-time variability. Instead, bad weather may slow all flights (requiring longer buffers) without making flight times more variable.

As in Table 3, the effect of a hub origin or destination is positive, although the origin effect is insignificant in the unrestricted sample. Note that the coefficients of the connecting share, rather than being significantly (and inexplicably) negative as before, are now uniformly insignificant. The origin and destination congestion measures, which were not present in Table 3, have uniformly positive and significant coefficients, indicating that congestion at either endpoint raises flight buffers. Recalling from Table 5 that time variability is positively related to congestion, this positive congestion effect may thus operate through variability.²⁶

The carrier coefficients are qualitatively similar to those in Table 3, although the absolute magnitudes tend to be smaller. The

²⁵ The maximum number of flights operated by a single aircraft observed in the ‘Marketing Carrier On-Time Performance’ database is 16. However, since operation of 9 or more flights is seldom observed, representing less than 0.40% of the initial database, these cases are not likely to meet our criterion of at least 30 observations in the calculation of \bar{n}_{min} and are therefore excluded from the unrestricted sample.

²⁶ Recall that our measure of flight time is block time, which includes the taxi time along with the airborne time.

Table 6
Flight-buffer regressions with variability proxies.

Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	b_1	b_2	b_1, b_2	b_1, b_2	b_1, \dots, b_8	b_1, \dots, b_8	b_1, \dots, b_8
	Two-flight sub-sample				Unrestricted sample		
February	0.740	0.283	0.548	0.547	0.161	0.164	0.164
March	-0.762	-0.645	-0.717	-0.706	-1.679***	-1.663***	-1.662***
April	-4.153***	-0.900*	-2.563***	-2.563***	-2.930***	-2.915***	-2.914***
May	-6.551***	-0.248	-3.564***	-3.560***	-3.756***	-3.742***	-3.742***
June	-8.256***	-0.413	-4.516***	-4.513***	-4.369***	-4.347***	-4.346***
July	-8.743***	-0.727	-4.861***	-4.865***	-4.628***	-4.608***	-4.607***
August	-8.977***	-1.279**	-5.192***	-5.202***	-4.589***	-4.572***	-4.571***
September	-8.970***	-3.034***	-5.914***	-5.943***	-4.646***	-4.638***	-4.637***
October	-8.501***	-3.735***	-6.041***	-6.062***	-4.521***	-4.509***	-4.508***
November	-5.134***	-3.523***	-4.178***	-4.198***	-2.427***	-2.417***	-2.416***
December	-2.819***	-2.874***	-2.830***	-2.827***	-1.338***	-1.332***	-1.332***
Connecting passengers	-0.636	0.300	-0.050	0.026	0.182	0.194	0.199
Hub origin	1.540***	-0.034	0.037	0.108	0.039	0.012	0.034
Hub destination	1.946***	2.691***	2.673***	2.691***	1.026***	1.041***	1.024***
Congestion origin	0.155***	0.175***	0.130***	0.127***	0.195***	0.188***	0.188***
Congestion destination	0.085***	0.131***	0.136***	0.134***	0.141***	0.133***	0.132***
Alaska Airlines	1.404**	1.045	2.415***	2.490***	0.063	0.076	0.075
Delta Airlines	-0.286	-0.866	-0.599	-0.583	0.483***	0.473***	0.477***
Frontier Airlines	-1.635	-1.768	-1.604*	-1.260	-0.582***	-0.574***	-0.554***
Hawaiian Airlines	-1.707*	-6.627***	-4.348***	-4.371***	-7.235***	-7.328***	-7.310***
Jet Blue	-4.847***	-4.617***	-5.532***	-5.319***	-3.807***	-3.789***	-3.781***
Southwest Airlines	-1.391**	-2.118***	-1.797***	-1.697***	0.846**	0.859***	0.862***
Spirit Airlines	-4.294***	-4.178***	-4.318***	-4.102***	-4.104***	-4.105***	-4.088***
United Airlines	-3.932***	-5.364***	-4.702***	-4.666***	-2.307***	-2.301***	-2.305***
Virgin America	-9.026***	-5.956***	-7.990***	-7.889***	-5.518***	-5.506***	-5.502***
Regional carrier	-0.735*	-0.070	-0.108	-0.143	2.208***	2.186***	2.194***
Competitors	-0.173	0.709***	0.284*	0.276*	0.211***	0.207***	0.205***
Distance	1.150***	1.202***	1.189***	1.186***	1.457***	1.457***	1.456***
Heavy aircraft	-2.055***	-1.681***	-1.790***	-1.829***	-4.852***	-4.837***	-4.833***
Past-year delay	11.388***	5.230***	9.171***	9.218***	3.978***	4.011***	4.011***
Weekend	-0.080	-0.110	-0.090	-0.093*	0.066***	0.044***	0.044***
Morning			0.186		-0.140***		
Afternoon			0.845***		-0.349***		
Late Afternoon			1.278***		0.557***		
Evening			0.873***		0.152***		
Flight 2				0.883***		0.178***	
Flight 3						-0.039	
Flight 4						0.227***	
Flight 5						0.253***	
Flight 6						0.125**	
Flight 7						-0.610***	
Flight 8						-1.052***	
Aircraft rotation							0.238***
Aircraft rotation ²							-0.035***
Constant	29.483***	16.645***	19.836***	20.015***	14.388***	14.436***	14.177***
R ²	0.676	0.680	0.658	0.658	0.654	0.654	0.654
Observations	101,920	100,916	202,836	202,836	4,126,741	4,126,741	4,126,741

(a) The estimated coefficients marked with ***, ** and * are statistically significant at, respectively the 1%, 5% and 10% level.

(b) The standard errors, not reported to save space, are clustered by route-month.

(c) All estimates include airport of origin and airport of destination fixed effects.

regional carrier dummy now has positive coefficients in the unrestricted sample, reversing the previous negative sign (now possibly reflecting feeder carriers' concerns about missed connections at hubs). While the competition effect remains positive, the large-aircraft dummy coefficients all flip sign to negative and significant, a welcome shift that confirms expectations. The new distance variable has a uniformly positive effect on flight buffers, and since distance and time variability are strongly related, this effect presumably operates through variability. The effect is about one extra minute of buffer per hundred miles flown. As for weekend flights, the weekend dummy coefficients in the two-flight sample are negative, although the unrestricted sample shows a significantly positive effect. Note that the managerial-share variable does not appear in Table 6 given its collinearity with the airport dummies.

The time-of-day dummy coefficients in columns (3) and (5) show a clearer pattern than in Table 3, with the buffer magnitudes higher late in the day than in the default early-morning period. Matching this pattern, the flight-2 dummy coefficient is now significantly positive rather than negative.

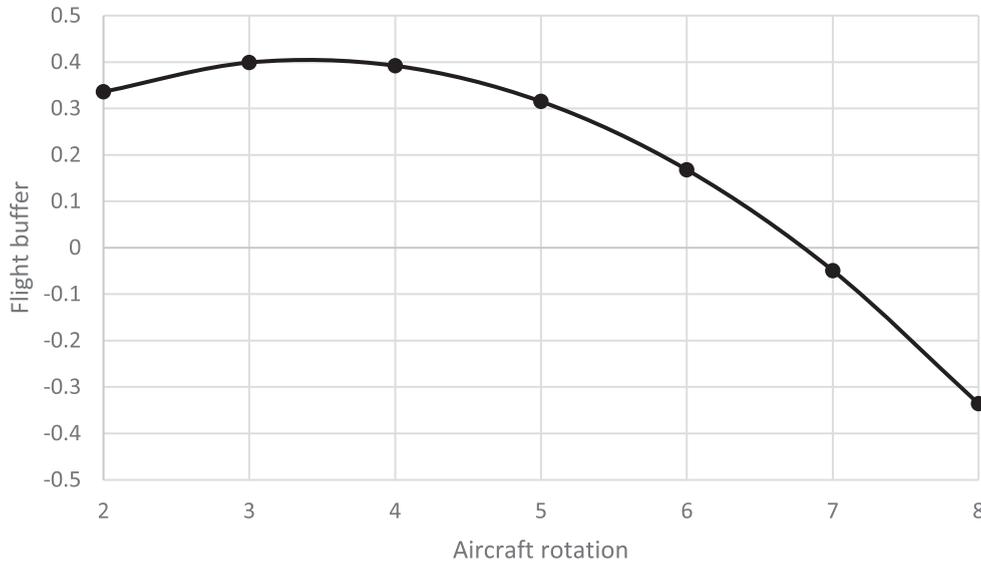


Fig. 6. Flight buffer as a function of aircraft rotation.

To explore the time-of-day effect in a different way, we investigate the effect on its buffer of a flight's position in the day's flight sequence, using either flight-sequence dummies or the aircraft-rotation variable (see Table 2).²⁷ Column (6) of Table 6 shows results using flight-sequence dummies in place of the time-of-day dummies for the unrestricted sample. With the exception of the insignificant flight-3 coefficient, the pattern is for flights 2 through 6 to have longer buffers than flight 1 (the default), while flights 7 and 8 have shorter buffers than flight 1. Because of its focus on just two flights, the theoretical model does not generate predictions for the cases covered by the unrestricted sample. But the empirical pattern appears to show a somewhat different logic than in the two-flight situation. In particular, delay propagation is apparently addressed through longer buffers for *earlier*, rather than later flights. The shortest buffers are for last few flights of the day, given that the aircraft will soon terminate its daily operations, removing any concern about delay propagation.

The regression in column (7) shows that this buffer pattern also emerges when the flight dummy variables are replaced by the continuous aircraft-rotation variable, which appears in quadratic form. The positive sign of the aircraft rotation coefficient together with the negative coefficient on the quadratic term point towards an inverse U-shape relationship between the flight buffer and the rotation variable, as shown in Fig. 6. The non-linear effect of aircraft rotations on the flight buffer appears to reconcile two opposite forces: the first force pushes towards longer flight buffers to lessen the risk of delay propagation; the second force pushes towards shorter flight buffers to maximize aircraft utilization.

Finally, the regressions in Table 6 include a variable measuring on-time performance in the previous year, on the belief that carriers would change their buffers to remedy past late arrivals. Using the DOT's 2017 'Reporting Carrier On-Time Performance (1987-present)' dataset, we calculated by carrier-route-month the proportion of delayed flights.²⁸ The coefficient on this past-year delay variable is positive and statistically significant, confirming the expectation that carriers lengthen flight buffers if delays on the route in the past year were more frequent.

7.5. Ground-buffer regressions with variability proxies

Table 7 shows ground-buffer regressions with variability proxies replacing the variability measure, doing so for both the two-flight and unrestricted samples. It is worth noting that the ground buffer is calculated at a flight's departing airport. In the two-flight sample, the observed airport corresponds to the departing airport of flight 2, whereas in the unrestricted sample, the observed airport is the

²⁷ The time-of-day dummies are not used in conjunction with either of the other two variables because they are almost functionally related. For example, while inclusion of both the time-of-day dummies and flight-sequence dummies would imply that independent variation is possible for these covariates, the fact that, say, the eighth flight of the day could never be a morning flight means that such independence is not present. For the same reason, we do not include the time-of-day dummies in the regressions in columns (1) and (2), as explained earlier.

²⁸ The 'Reporting Carrier On-Time Performance (1987-present)' dataset is essentially the same as the 'Marketing Carrier On-Time Performance (Beginning January 2018)', which we use in our empirical analysis, but it does not report the actual operating carrier. Thus, it does not distinguish between major airlines and affiliated regional carriers. For this reason and also because an airline could stop serving the route from one year to the next, some 2018 observations are not matched with the 2017 data.

Table 7
Ground-buffer regressions with variability proxies.

Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)
	b_g	b_g	b_g	b_g	b_g	b_g
	Two-flight sub-sample			Unrestricted sample		
February	0.246	0.271	0.056	-0.184	-0.176	-0.316
March	-0.773	-0.791	-1.284	-1.082***	-1.092***	-1.455***
April	-0.472	-0.716	-1.305	-0.862***	-0.871***	-1.210***
May	1.036	0.778	0.529	-0.159	-0.170	-0.486**
June	-0.629	-0.912	-1.224	-1.224***	-1.234***	-1.767***
July	0.475	0.139	-0.048	-0.906***	-0.909***	-1.467***
August	1.317	0.993	0.754	-0.563***	-0.578***	-1.170***
September	3.224***	2.935***	2.773***	1.274***	1.315***	0.829***
October	2.814***	2.543***	2.102**	0.893***	0.868***	0.360*
November	1.865**	1.705**	1.198	0.917***	0.921***	0.457**
December	0.382	0.205	0.015	0.433**	0.404**	0.063
Connecting passengers	8.245***	9.167***	7.904***	4.499***	2.172***	1.229***
Hub turnaround	14.250***	13.050***	10.085***	14.013***	13.477***	14.025***
Congestion turnaround	-0.375***	-0.368***	0.072**	-0.388***	-0.396***	0.153***
Slot controlled airport			-2.714***			-0.335*
Alaska Airlines	1.578*		5.114***	6.646***		8.595***
Allegiant Air	-11.395***		-21.571***	-5.952***		-10.330***
Delta Airlines	-3.080***		-0.664	-3.060***		-1.160***
Frontier Airlines	-4.015**		-1.468	-1.167***		-0.390
Hawaiian Airlines	-4.920***		7.824***	-4.987***		8.641***
Jet Blue	-3.958***		-6.371***	0.515		0.089
Southwest Airlines	-6.813***		-7.743***	-7.012***		-5.465***
Spirit Airlines	-6.037***		-6.573***	1.925***		3.299***
United Airlines	4.644***		7.802***	4.924***		9.248***
Virgin America	7.170***		6.824***	9.104***		11.603***
Regional carrier	-1.190*		-6.839***	-1.740***		-4.471***
Low-cost carrier		-5.649***			-4.996***	
Heavy aircraft	-0.694	-1.121*	1.503**	-2.692***	-3.721***	-1.027***
Morning	7.225***	7.033***	5.976***	4.205***	4.240***	1.979***
Afternoon	14.282***	14.122***	14.192***	4.453***	4.396***	2.816***
Late Afternoon	19.239***	19.341***	16.752***	5.840***	5.915***	3.907***
Evening	19.951***	20.043***	19.082***	9.379***	9.575***	8.089***
Weekend	1.932***	1.906***	2.647***	1.867***	1.918***	2.361***
Constant	27.496***	24.564***	37.866***	23.269***	22.581***	32.451***
R ²	0.206	0.201	0.157	0.237	0.227	0.202
Observations	133,178	133,178	133,178	4,318,387	4,318,387	4,318,387

- (a) The estimated coefficients marked with ***, ** and * are statistically significant at, respectively the 1%, 5% and 10% level.
- (b) The standard errors, not reported to save space, are clustered by route-month.
- (c) All estimates but columns (3) and (6) include airport of turnaround and airport of origin fixed effects; columns (3) and (6) include only airport of origin fixed effects.

departing airport of Flight i , with $i = 2, \dots, 8$. Thus, the flight-specific variables in the regression, such as departure time, refer to the following flight. In this way, the first flight of the day is not included in the regressions for the unrestricted sample.

The effect of weather on ground buffers, as captured by the month dummies, is not as clear as in the flight-buffer regressions. The results show that in both samples, ground buffers are longest (relative to January) in the Fall months of September, October and November, a pattern that does not have a clear weather-based interpretation. However, in the unrestricted sample, the buffers over the March-August period are significantly shorter than in January, which is partly consistent with short ground buffers being scheduled in months with better weather.

The coefficients on the connecting-share and hub-turnaround variables are uniformly positive and significant across the regressions, as in Table 4. A congested turnaround airport has longer taxi times, which may prompt the airline to cut ground time, as shown by the negative coefficients on the congestion-turnaround variable. But adding the slot-control dummy reverses this effect in both samples, leading to a significantly positive congestion turnaround coefficient in columns (3) and (6). The significantly negative slot-control effect in the two-flight subsample is large: ground buffers are almost 3 min shorter on average at such airports, which tend to be congested.²⁹

The pattern of the airline-dummy coefficients in the two-flight sample is somewhat different than in Table 4, with the coefficients of Delta and Frontier switching to significantly negative (Southwest's and Spirit's coefficients also become more negative). The

²⁹ In order to obtain these estimates, we remove the turnaround-airport fixed effects, which are perfectly collinear with the *Slot-controlled airport* variable.

coefficient pattern changes somewhat in the unrestricted sample, as seen in columns (4) and (6). The coefficients of the low-cost and regional-carrier dummy variables, however, follow the same pattern as in Table 4, being negative across regressions. The heavy-aircraft coefficient, previously insignificant, is now unstable in the two flight sample, but significantly negative in the unrestricted sample.

As in Table 4, the estimated coefficients of the time-of-day dummies show that airlines keep lengthening ground buffers across the day. This effect is strictly monotonic in each regression, rising from 2-7 min extra buffer in the morning (relative to early morning) to 8–20 min extra buffer in the evening. With the model portraying operation of only two flights, it cannot predict this pattern. But the pattern suggests that, when an aircraft operates many flights per day, ground buffers may play a more prominent role than flight buffers in mitigating delay propagation late in the day.

Airlines may operate with spare capacity during weekends, since most business travel is mid-week, thus explaining why they set slightly longer ground buffers during weekends, about two minutes longer on average. Our data show that flights scheduled during weekends are more punctual than non-weekend flights, and these longer ground buffers may be part of the reason.³⁰

8. Conclusion

This paper has presented an extensive theoretical and empirical analysis of the choice of schedule buffers by airlines. With airline delays a continuing problem around the world, such an undertaking is valuable, and its lessons extend to other transport sectors such as rail and intercity bus service. One useful lesson from the theoretical analysis of a two-flight model is that the mitigation of delay propagation is done entirely by the ground buffer and the second flight’s buffer. The first flight’s buffer plays no role because the ground buffer is a perfect, while nondistorting, substitute. In addition, the apportionment of mitigation responsibility between the ground buffer and the flight buffer of flight 2 is shown to depend on the relationship between the costs of ground- and flight-buffer time.

The empirical results show the connection between buffer magnitudes and a host of variables, including the month of operation and distance of a flight, whether the flight operates early or late in the day, and congestion measures at the endpoints. In addition, the initial regressions relate buffer magnitudes to the variability of flight times, which simulations of the model identify as an important determining factor (the results show that high variability lengthens flight buffers).

Fruitful extensions to this work would most likely lie in the theoretical area. The model could be extended to include additional sequential flights, and the sketch of connecting traffic provided in the paper could be expanded into a full analysis. In addition, passenger scheduling preferences could be introduced, with a buffer-related extension of scheduled flight times possibly becoming less desirable if it leads to a divergence between a passenger’s preferred and actual arrival times. The resulting models would be complex, but additional insights could be generated, with relevance extending beyond the airline industry.

Another avenue of exploration could be in the area of market structure. For example, if a profit-maximizing airport sets the gate rental cost at an excessive level, the resulting decrease in ground buffers will impair airline mitigation of delay propagation, with negative effects on passengers. Alternatively, the entry barrier of gate shortages resulting from airport dominance by a large carrier (analyzed by Ciliberto and Williams (2010)) would also affect the ability of smaller carriers to address delay propagation via adequate ground time. Exploration of these issues could be illuminating.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A

A.1. The derivatives of Ω

This appendix section computes the derivatives of the objective function with respect to b_g, b_1 , and b_2 . Since (2) does not involve b_g , the objective function’s derivative with respect to b_g is found by differentiating the sum of (16) and (17). The derivative of the first line of (16) with respect to b_g equals

$$x f_1(b_1 + b_g) \int_{b_2}^{\bar{\epsilon}_2} (\epsilon_2 - b_2) f_2(\epsilon_2) d\epsilon_2. \tag{a1}$$

The second line of (16) can be written as

$$x \int_{\epsilon_1 = b_1 + b_g}^{\bar{\epsilon}_1} \left[\int_{\epsilon_2 = b_1 + b_2 + b_g - \epsilon_1}^{\bar{\epsilon}_2} \left[\epsilon_1 + \epsilon_2 - (b_1 + b_2 + b_g) \right] f_2(\epsilon_2) d\epsilon_2 \right] f_1(\epsilon_1) d\epsilon_1$$

³⁰ While the average arrival delay of a weekend flight is 6 min in the unrestricted sample and 0 min in the two-flight sample, the delays increase to 10 and 3 min, respectively, for non-weekend flights.

$$= x \int_{\epsilon_1=b_1+b_g}^{\bar{\epsilon}_1} Q(\epsilon_1, b_g) f_1(\epsilon_1) d\epsilon_1, \tag{a2}$$

where $Q(\epsilon_1, b_g)$ denotes the term in brackets in the first line of (a2). The b_g -derivative of (a2) is

$$x \int_{\epsilon_1=b_1+b_g}^{\bar{\epsilon}_1} \frac{\partial Q(\epsilon_1, b_g)}{\partial b_g} f_1(\epsilon_1) d\epsilon_1 - x Q(b_1 + b_g, b_g) f_1(b_1 + b_g). \tag{a3}$$

Substituting $\epsilon_1 = b_1 + b_g$ in the bracketed term in (a2) to evaluate $Q(b_1 + b_g, b_g)$ in (a3), the second term in (a3) equals the negative of (a1), so that these terms cancel. The b_g -derivative of (16) is then equal to the first term in (a3).

$\partial Q/\partial b_g$ consists of two components, the first of which comes from differentiating the bracketed expression in (a2) with respect to the limit of integration, a derivative that equals zero upon substituting the limit into the integrand. The second component comes from differentiating with respect to b_g under the integral, which yields the bracketed expression in (a2) with the integrand replaced by $-f_2(\epsilon_2)$. Therefore, the first-term in (a3) reduces to

$$-x \int_{\epsilon_1=b_1+b_g}^{\bar{\epsilon}_1} \left[\int_{\epsilon_2=b_1+b_2+b_g-\epsilon_1}^{\bar{\epsilon}_2} f_2(\epsilon_2) d\epsilon_2 \right] f_1(\epsilon_1) d\epsilon_1. \tag{a4}$$

Applying the same steps to (17), the b_g -derivative of that expression equals

$$y \int_{\epsilon_1=b_1+b_g}^{\bar{\epsilon}_1} \left[\int_{\epsilon_2=0}^{b_1+b_2+b_g-\epsilon_1} f_2(\epsilon_2) d\epsilon_2 \right] f_1(\epsilon_1) d\epsilon_1. \tag{a5}$$

Replacing the bracketed terms in (a4) and (a5) with $1 - F_2(b_1 + b_2 + b_g - \epsilon_1)$ and $F_2(b_1 + b_2 + b_g - \epsilon_1)$, respectively, the sum of (a4) and (a5) can be written as

$$\begin{aligned} &= \int_{\epsilon_1=b_1+b_g}^{\bar{\epsilon}_1} \left[-x(1 - F_2(b_1 + b_2 + b_g - \epsilon_1)) + yF_2(b_1 + b_2 + b_g - \epsilon_1) \right] f_1(\epsilon_1) d\epsilon_1 \\ &= (x + y) \int_{\epsilon_1=b_1+b_g}^{\bar{\epsilon}_1} \left[F_2(b_1 + b_2 + b_g - \epsilon_1) - \frac{x}{x + y} \right] f_1(\epsilon_1) d\epsilon_1 = \frac{\partial \Omega}{\partial b_g}. \end{aligned} \tag{a6}$$

The derivative of the objective function with respect to b_1 builds on the previous results. The b_1 -derivatives of (16) and (17) are identical to the b_g derivatives, given by (a4) and (a5), with their sum equal to (a6). Since (2) also depends on b_1 , the b_1 -derivative of the objective function is then the expression in (3) (slightly rearranged) plus (a6):

$$\frac{\partial \Omega}{\partial b_1} = (x + y) \left[F_1(b_1) - \frac{x}{x + y} \right] + \frac{\partial \Omega}{\partial b_g}. \tag{a7}$$

Using similar steps,

$$\frac{\partial \Omega}{\partial b_2} = (x + y) F_1(b_1 + b_g) \left[F_2(b_2) - \frac{x}{x + y} \right] + \frac{\partial \Omega}{\partial b_g}. \tag{a8}$$

A.2. Convexity of Ω

This appendix section first shows that the objective function Ω is strictly convex in b_1 and b_2 , with b_g fixed at the optimal value, doing so for both cases of zero and positive buffer costs. The first step is to compute the second derivatives of Ω with respect to b_1 and b_2 , assuming zero buffer costs. Using the shorthand Ω_{ij} for $\partial^2 \Omega / \partial b_i \partial b_j$, differentiation of (20) and (21) yields (after suppressing the multiplicative factor $x + y$)

$$\begin{aligned} \Omega_{11} = & f_1(b_1) - f_1(b_1 + b_g) \left[F_2(b_2) - \frac{x}{x + y} \right] \\ & + \int_{\epsilon_1=b_1+b_g}^{\bar{\epsilon}_1} f_1(\epsilon_1) f_2(b_1 + b_2 + b_g - \epsilon_1) d\epsilon_1 \end{aligned} \tag{a9}$$

$$\Omega_{22} = F_1(b_1 + b_g) f_2(b_2) + \int_{\epsilon_1=b_1+b_g}^{\bar{\epsilon}_1} f_1(\epsilon_1) f_2(b_1 + b_2 + b_g - \epsilon_1) d\epsilon_1 \tag{a10}$$

$$\Omega_{12} = \int_{\epsilon_1=b_1+b_g}^{\bar{\epsilon}_1} f_1(\epsilon_1)f_2(b_1+b_2+b_g-\epsilon_1)d\epsilon_1 < \Omega_{11}, \Omega_{22}. \quad (\text{a11})$$

Observe that the same expressions apply when buffer costs are positive since they vanish in computing the second derivatives.

Ω_{22} and Ω_{12} are positive, while the sign of Ω_{11} in (a9) is unclear. However, at the optimum in the zero-cost buffer case, the bracketed term in (a9) is zero, making Ω_{11} positive and also ensuring satisfaction of the inequality in (a11). The same conclusion holds in the positive-cost buffer case if $c_f > c_g$, in which case the bracketed term in (a9) is negative at the optimum. With $\Omega_{11}, \Omega_{22} > 0$ and $H \equiv \Omega_{11}\Omega_{22} - \Omega_{12}^2 > 0$ (a consequence of the inequalities in (a11)), Ω is thus strictly convex in b_1 and b_2 at the optimum. As a result, b_1 and b_2 values satisfying the first-order conditions yield a local minimum for Ω , holding b_g fixed at its optimal value. Convexity of Ω in all three buffers cannot be established analytically and must be assumed.

References

- AhmadBeygi, S., Cohn, A., Lapp, M., 2010. Decreasing airline delay propagation by re-allocating scheduled slack. *IIE Trans.* 42, 478–489.
- Arikan, M., Deshpande, V., Sohoni, M., 2013. Building reliable air-travel infrastructure using empirical data and stochastic models of airline networks. *Oper. Res.* 61, 45–64.
- Ball, M., Barnhart, C., Dresner, M., Hansen, M., Neels, K., Odoni, A., Peterson, E., Sherry, E., Trani, A., Zou, B., 2010. Total Delay Impact Study: A Comprehensive Assessment of the Costs and Impacts of Flight Delay in the United States. Final Report, National Center of Excellence for Aviation Operations Research (NEXTOR).
- Barnhart, C., Cohn, A., 2004. Airline schedule planning: Accomplishments and opportunities. *Manuf. Service Oper. Manage.* 6, 3–22.
- Britto, R., Dresner, M., Voltes, A., 2012. The impact of flight delays on passenger demand and societal welfare. *Transp. Res. Part E* 48, 460–469.
- Brueckner, J.K., Czerny, A.I., Gaggero, A.A., 2020. Airline schedule buffers and flight delays: A discrete model. Unpublished paper, UC Irvine.
- Ciliberto, F., Williams, J.W., 2010. Limited access to airport facilities and market power in the airline industry. *J. Law Econ.* 53, 467–495.
- Deshpande, V., Arikan, M., 2012. The impact of airline flight schedules on flight delays. *Manuf. Service Oper. Manage.* 14, 423–440.
- Forbes, S.J., 2008. The effect of air traffic delays on airline prices. *Int. J. Ind. Organ.* 26, 1218–1232.
- Forbes, S.J., Lederman, M., 2010. The performance implications of vertical integration: Evidence from the airline industry. *RAND J. Econ.* 41, 765–790.
- Forbes, S.J., Lederman, M., 2009. Adaptation and vertical integration in the airline industry. *Am. Econ. Rev.* 99, 1831–1849.
- Forbes, S.J., Lederman, M., Wither, M.J., 2019. Quality disclosure when firms set their own quality targets. *Int. J. Ind. Organ.* 62, 228–250.
- Forbes, S.J., Lederman, M., Yuan, Z., 2019. Do airlines pad their schedules? *Rev. Ind. Organ.* 54, 61–82.
- Hao, L., Hansen, M., 2014. Block time reliability and scheduled block time setting. *Transp. Res. Part B* 69, 98–111.
- He, W., 2019. Integrating overbooking with capacity planning: Static model and application to airlines. *Prod. Oper. Manage.* 28, 1972–1989.
- Kafle, N., Zou, B., 2016. Modeling flight delay propagation: A new analytical-econometric approach. *Transp. Res. Part B* 93, 42–52.
- Kang, L., Hansen, M., 2017. Behavioral analysis of airline scheduled block time adjustment. *Transp. Res. Part E* 103, 56–68.
- Kang, L., Hansen, M., 2018. Assessing the impact of tactical airport surface operations on airline schedule block time setting. *Transp. Res. Part C* 89, 133–147.
- Luttmann, A., 2019. Evidence of directional price discrimination in the U.S. airline industry. *Int. J. Ind. Organ.* 62, 291–329.
- Mayer, C., Sinai, T., 2003. Why do airlines systematically schedule their flights to arrive late? Unpublished paper, Wharton School of Business, University of Pennsylvania.
- Mazzeo, M., 2003. Competition and service quality in the U.S. airline industry. *Rev. Ind. Organ.* 22, 275–296.
- Prince, J.T., Simon, D.H., 2015. Do incumbents improve service quality in response to entry? Evidence from airlines' on-time performance. *Manage. Sci.* 61, 372–390.
- Shumsky, R.A., 1993. Response of U.S. air carriers to on-time disclosure rule. *Transp. Res. Rec.* 1379, 9–16.
- Sohoni, M., Lee, Y.-C., Klabjan, D., 2011. Robust airline scheduling under block-time uncertainty. *Transp. Sci.* 45, 451–464.
- Whitin, T.M., 1955. Inventory control and price theory. *Manage. Sci.* 2, 61–68.
- Wang, X., 2015. Service competition in the airline industry: Schedule robustness and market structure. Unpublished paper.
- Wang, Y., Zhou, Y., Hansen, M., Chin, C., 2019. Scheduled block time setting and on-time performance of U.S. and Chinese airlines: A comparative analysis. *Transp. Res. Part A* 130, 825–843.
- Zhang, A., Czerny, A.I., 2012. Airports and airlines economics: An interpretive review of recent research. *Econ. Transp.* 1, 15–34.
- Zhang, D.J., Salant, Y., Van Mieghem, J.A., 2018. Where did the time go? On the increase in airline schedule padding over 21 years. Unpublished paper.
- Zou, B., Hansen, M., 2012. Impact of operational performance on air carrier cost structure: Evidence from US airlines. *Transp. Res. Part E* 48, 1032–1048.