Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 170 (2017) 29–55

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Incentives in landing slot problems ☆

James Schummer [a,*], Azar Abizada [b]

[a] *MEDS, Kellogg School of Management, Northwestern University, Evanston IL 60208, USA*
[b] *School of Business, ADA University, 11 Ahmadbay Aghaoglu St., Baku AZ1008, Azerbaijan*

Received 19 May 2016; final version received 30 January 2017; accepted 11 April 2017
Available online 18 April 2017

Abstract

During weather-induced airport congestion, landing slots are reassigned based on flights' feasible arrival times and cancelations. We consider the airlines' incentives to report such information and to execute cancelations, creating positive spillovers for other flights. We show that such incentives conflict with Pareto-efficiency, partially justifying the FAA's non-solicitation of delay costs. We provide mechanisms that, unlike the FAA's current mechanism, satisfy our incentive properties to the greatest extent possible given the FAA's own design constraints. Our mechanisms supplement Deferred Acceptance with a "self-optimization" step accounting for each airline's granted right to control its assigned portion of the landing schedule.

JEL classification: C78; D47; D82; R4

Keywords: Mechanism design; Landing slots; Matching; Incentives

☆ We are grateful for comments from Josh Cherry, Tayfun Sönmez, Eric Weese, and seminar participants at Boston College, Northwestern University, and the Fifth Matching in Practice Workshop (Brussels). Numerous referees and a co-editor helpfully improved the exposition of our results. We particularly thank Fuhito Kojima and Utku Ünver for extensive comments improving Subsection 5.3. Schummer thanks the ISER at Osaka University for its hospitality during the beginning of this project.
* Corresponding author.
E-mail addresses: schummer@kellogg.northwestern.edu (J. Schummer), aabizada@ada.edu.az (A. Abizada).

http://dx.doi.org/10.1016/j.jet.2017.04.003
0022-0531/

1. Introduction

Weather-caused flight delays frustrate policy makers as much as they frustrate airline passengers: the annual economic cost of such delays is measured in billions of dollars.[1] Though weather delays are unavoidable, the resulting delay costs are mitigated by rescheduling delayed flights into earlier landing slots that have been vacated by newly canceled flights. In the U.S. (as elsewhere) this rescheduling is done only *after* airlines report privately known flight information through a centralized mechanism. While this problem has yielded a significant optimization literature, there has been little analysis of airlines' incentives to report their information in the first place. We formalize this problem with mechanism design constraints appropriate for the setting, focusing on three forms of incentives pertaining to: reporting flight delays, reporting waiting costs, and making and reporting flight cancelations.

Our first set of results can be viewed as an incentives-based justification for the fact that the FAA's rescheduling mechanism is not fully Pareto-efficient. Specifically, we show that Pareto-efficiency would be incompatible with *any single one* of our three incentive conditions. Nevertheless a weaker form of efficiency—the one considered in the transportation literature on this problem—*is* simultaneously compatible with two of our incentive conditions and a weakened version of the third. We construct rules exhibiting this compatibility by supplementing the Deferred Acceptance algorithm (Gale and Shapley (1962)) with a procedure that accounts for the airlines' granted rights to rearrange their own portions of the landing schedule.

Our most significant finding is that our rules give strong incentive for airlines to execute and promptly report flight cancelations. This result is robust to dynamic specifications of the model and to the endogeneity of cancelation decisions. This is important during periods of congestion, when cancelations create positive spillovers for other airlines. Under any of our rules, in fact, *a flight cancelation is necessarily welfare improving*: each remaining flight is assigned a (weakly) better landing slot. In contrast, Schummer and Vohra (2013) show that the FAA's current mechanism can provide a strict disincentive for an airline to cancel flights even in a static model.

1.1. Ground delay programs

To justify our modeling assumptions and motivate our design constraints, we describe the relevant institutional details of a Ground Delay Program (GDP). A GDP is used to reduce the rate of air traffic at an airport when demand for landing slots is projected to exceed capacity, e.g. when landing rates are to be reduced due to inclement weather.

Hours in advance of a forecasted weather event, air traffic management declares a GDP to be in effect. First, flights destined for the affected airport are given delayed departure times while still on the ground at their origination airport. This *Ration-by-Schedule* step of a GDP simply spreads out arrivals so as not to exceed the new, reduced landing capacity. For example an airport that normally lands sixty flights per hour may be reduced to thirty flights per hour due to weather. Thirty 2-minute slots replace sixty 1-minute slots and are assigned to thirty flights based on the original schedule. We take this process as given, and it is not part of our analysis.

[1] A U.S. Senate report (Schumer and Maloney (2008)) estimates the economic cost of *all* flight delays to exceed $40 billion per year for the U.S., around half of which is direct cost to the airlines. Weather causes roughly one fifth of all delays.