Incentives in Landing Slot Problems*

James Schummer† Azar Abizada‡

Current version: January 5, 2017
First version: March 2012

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.

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

*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. MEDS, Kellogg School of Management, Northwestern University, Evanston IL 60208, USA. Email: schummer@kellogg.northwestern.edu.

‡School of Business, ADA University, 11 Ahmadbay Aghaoglu St., Baku AZ1008, Azerbaijan. Email: aabizada@ada.edu.az.

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] A U.S. Senate report (Sen. C. Schumer and C. 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.

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.

In fact this Ration-by-Schedule step would be the end of the process if not for certain operating constraints of the airlines. When a flight is assigned a new arrival (and hence departure) time, its airline may have to cancel it (e.g. if the crew now might exceed its legal work hours). Such flight cancelations—made *after* Ration-by-Schedule—have the interesting effect of freeing up additional units of the scarce resource: landing slots.

To utilize these newly created gaps in the landing schedule, the GDP's *Reassignment* step is executed, which is the focus of our analysis. First, airlines report their relevant flight information—cancelations and the feasible arrival time of each remaining flight. Next, a centralized mechanism[2] feasibly reschedules the remaining flights to eliminate vacancies in the landing schedule. Due to various design constraints, however, there is no single, obvious choice of mechanism to use at this step.

First, a flight cannot be assigned to an arbitrary earlier slot since this might violate its feasible arrival time constraint. Second, even though a landing schedule specifies precisely which flight occupies which landing slot, an airline may wish to swap the positions of two of its own flights. Such

[2]The FAA currently uses the Compression algorithm for Reassignment. Compression was introduced in the 1990's, in part to solve an incentives problem with the previous assignment method; see Wambsganss (1997).

swapping not only seems reasonable, but is a right explicitly granted by the FAA's operations guidelines.[3] This novel design constraint plays a significant role in Section 5, implying that even if a mechanism prescribes one landing schedule, an airline may ultimately consume a different one based on its privately known preferences for such swapping.

Third, the implementation of *Pareto-efficient* landing schedules requires knowing the airlines' full preferences over landing schedules. Significantly, *the FAA does not solicit such information* during the Reassignment step. It may be impractical for airlines to evaluate and report such complex information, unique to every GDP event. Regardless, we consider mechanisms both with and without this third (soft) design constraint. Section 4 considers Pareto-efficient mechanisms when such preference information exists, and Section 5 considers mechanisms that use only the information solicited by the FAA. The negative results of Section 4 lend support to the restriction of Section 5.

1.2 Related Literature

The paper most related to ours is by Schummer and Vohra (2013), who consider our problem using an incomplete notion of airline preferences. Without imposing all of the design constraints discussed above, they analyze the weak core and provide results on weak incentives, discussed in Subsection 5.4. Our paper also relates to both an operations-oriented literature on GDP's and a game theoretic literature on object assignment and matching.

The operations literature on GDP's emphasizes optimization; incentives are mentioned but not formalized.[4] Vossen and Ball (2006a) use a linear programming approach to minimize airline costs, yielding a generalization of the Ration-by-Schedule process. Vossen and Ball (2006b) interpret the FAA's Compression algorithm as a barter exchange process. Various papers generalize this optimization problem by modeling endogenous flight cancelations (Ball, Dahl, and Vossen (2009)), intra-flight arrival constraints (Hoffman and Ball (2007)), downstream costs from delays (Niznik (2001)), or prioritization by flight distance (Ball and Lulli (2004) and Ball, Hoffman, and Mukherjee (2010)).

[3] The rule is in Section 17–9–5 of the *Facility Operation and Administration Handbook*. The handbook is available through http://www.faa.gov/atpubs.

[4] An interesting exception is a rigorous, explicit example of manipulability provided by Wambsganss (1997), in an historical perspective on GDP's.

Our emphasis on incentives fits more closely within the assignment and matching literatures. With airlines exchanging endowed landing slots, our model naturally appears to be a generalization of the (1-sided) housing market model of Shapley and Scarf (1974).[5] However, our approach in Section 5 is to embed this problem into a version of the celebrated *two-sided* College Admissions model of Gale and Shapley (1962). We extend their Deferred Acceptance algorithm to respect the design constraints mentioned earlier in Subsection 1.1, deriving incentive conditions appropriate for this setting.

Despite some resemblance to the two-sided College Admissions model, it is our *colleges* (i.e. the airlines) that have economically meaningful preferences, not our *students* (i.e. landing slots). Thus our environment reverses the School Choice environment of Abdulkadiroğlu and Sönmez (2003) in which *students* have preferences but *colleges* do not.[6] This seemingly minor difference—making the *college* side strategic—is significant since incentive compatibility is well-known to be more elusive for the college side of such markets.[7] For instance, (student-proposing) Deferred Acceptance gives the *student* side incentive to truthfully report preferences, but no analogous result holds for the college side. *Despite this, we obtain positive incentive results for our setting in Subsection 5.3 and Subsection 5.4.*

Our model includes *endowments* of objects (slots) to airlines. Endowments appear in the house allocation model of Abdulkadiroglu and Sönmez (1999), which yields an individually rational, Pareto-efficient, strategy-proof mechanism. However the consumption of *multiple* objects again leads to negative results: Konishi et al. (2001) show the weak core is often empty; Atlamaz and Klaus (2007) show efficient rules to be manipulable by destroying, concealing, or transferring endowed objects.[8] While Theorem 3 parallels a result of theirs, neither result implies the other. More importantly we

[5] Schummer and Vohra (2013) take this approach. Balakrishnan (2007) uses the Shapley–Scarf model directly by treating flights (rather than airlines) as agents. This allows use of the Top Trading Cycle algorithm, but ignores incentives at the airline level.

[6] Relatedly, see Balinski and Sönmez (1999), Sönmez (2013), Sönmez and Switzer (2013), and Kominers and Sönmez (2016).

[7] See Roth and Sotomayor (1990) for a survey, along with Dubins and Freedman (1981), Roth (1982), Roth (1985), and extensions by Sönmez (1996), Alcalde and Barberà (1994), and Takagi and Serizawa (2010).

[8] Endowment-manipulation is also considered by Postlewaite (1979). Sertel and Ozkal-Sanver (2002) show that Deferred Acceptance is manipulable through the hiding of monetary endowments, contrasting interestingly with our positive result in Subsection 5.3.

provide a contrasting *positive* result in Subsection 5.3.

2 Model

There is a finite set of airlines, $\mathcal{A}$. Each airline $A \in \mathcal{A}$ has a finite set of flights denoted F_A; let $F = \bigcup_{A \in \mathcal{A}} F_A$. Flights are to be assigned to an ordered set of available landing slots, denoted by a set of integers $S \subset \mathbb{N} \equiv \{1, 2, \ldots\}$, with $|S| \geq |F|$. We interpret the slot labels as physical units of time, so slot 1 is the earliest slot, slot 5 is two units of time later than slot 3, etc.[9]

Each flight $f \in F$ is to be assigned a slot no earlier than its **earliest feasible arrival time** $e_f \in \mathbb{N}$. Assigning flights to slots, a **landing schedule** is a function $\Pi\colon F \to S$ that is injective ($f \neq f'$ implies $\Pi(f) \neq \Pi(f')$). Landing schedule Π is **feasible** if for all $f \in F$, $\Pi(f) \geq e_f$.

Given some *initial* landing schedule Π^0, one can infer for any flight $f \in F_A$ that slot $\Pi^0(f)$ is initially endowed to airline A; however Π^0 does not specify endowment of any initially vacant slots. As discussed in section 1.1, whichever airline vacated such a slot maintains some degree of property rights over it (footnote 3). Therefore we introduce the concept of a **slot ownership function**, a function $\Phi\colon \mathcal{A} \to 2^S$ such that $A \neq B$ implies $\Phi(A) \cap \Phi(B) = \emptyset$. If $s \in \Phi(A)$ is vacant according to Π^0, the interpretation is that A canceled a flight previously occupying s.[10]

We say that Φ is consistent with a landing schedule Π when occupation (under Π) implies ownership (under Φ): $\forall A \in \mathcal{A}, \forall f \in F_A, \quad \Pi(f) \in \Phi(A)$. Any pair (Π, Φ) satisfying this consistency condition is called an **assignment**.

Preferences We model airlines' preferences to reflect the fact that *earlier is better*: all else being constant, airline A prefers flight $f \in F_A$ to be assigned to as early a landing slot as possible (though not earlier than e_f). While this

[9]We model a "single runway, single airport" problem. Our results easily extend to "multiple runway" problems in which each time-unit contains multiple, identical landing slots. We ignore the modeling of multiple airports (like the majority of the cited operations literature in Subsection 1.2) since it would be rare to see many (distant) airports simultaneously experiencing GDP's due to unexpected weather events.

[10]Thus an initial landing schedule and slot ownership function describe the scenario in which cancelations have already been made. The strategic decision whether to cancel flights in the first place is implicitly captured by the non-manipulability conditions of Definition 6 and Definition 12, and the result of Observation 2.

assumption is what motivates our analysis, it says nothing about how an airline evaluates tradeoffs amongst moving different flights to earlier slots in the schedule. If given the choice, airline A's preference to move flight $f \in F_A$ to an earlier slot might be more intense than its preference to move $g \in F_A$.[11] Such preference intensities could vary due to differences in flights' operating costs, numbers of passengers, deadlines for crews timing out, future needs of the aircraft, etc.

While this could make the class of real-world airline preferences very rich (and complicated), it turns out that our results are quite robust to the modeling of such preferences. To begin with, we perform our analysis using a simple model of *linear delay costs*, in which an airline evaluates schedules by aggregating its flights' costs. It is then easily argued that our positive results extend to any model of preferences in which *earlier is better* as defined above. In addition, any of our negative results immediately extends to any richer class of airline preferences using standard arguments. Since we believe that any model of this problem should at least *contain* our class of preferences, our conclusions are robust to this modeling specification.

Formally each flight $f \in F_A$ has a **weight** $w_f > 0$, which can be interpreted as a (relative) delay cost per unit of time. To illustrate, suppose A has two flights, $F_A = \{f, g\}$, and consider landing schedules Π and Π' where $\Pi(f) = 5$, $\Pi(g) = 6$, $\Pi'(f) = 3$, and $\Pi'(g) = 9$. Moving from Π to Π', A gains $2 \cdot w_f$ units through f and loses $3 \cdot w_g$ units through g. If $2 \cdot w_f > 3 \cdot w_g$ then A prefers Π' to Π.

Definition 1. A list of weights $(w_f)_{f \in F}$ induces for each airline $A \in \mathcal{A}$ a (weak) **preference relation over feasible landing schedules** $\succeq_A^w$ (with strict part $\succ_A^w$) as follows. For all feasible landing schedules Π, Π':

$$\Pi \succeq_A^w \Pi' \quad \Leftrightarrow \quad \sum_{f \in F_A} w_f(\Pi'(f) - \Pi(f)) \geq 0.$$

If Π is feasible but Π' is infeasible for A, we say $\Pi \succ_A^w \Pi'$.

This linear delay cost model is common in the operations literature on optimization in GDP's (see Subsection 1.2). There the typical objective is

[11] Schummer and Vohra (2013) assume that an airline does not make such tradeoffs, and only considers a landing schedule "better" if all its flights (weakly) improve. This yields weaker incentive compatibility concepts, discussed in Subsection 5.4.

to minimize *aggregate* delay costs, which implies inter-airline comparability of weights. Our analysis does not depend on whether these weights are comparable across airlines since they merely parameterize preferences.

To summarize our model, an **Instance** of a Landing Slot Problem is a tuple $I = (S, \mathcal{A}, (F_A)_{A\in\mathcal{A}}, e, w, \Pi^0, \Phi^0)$ of slots, airlines, flights, earliest arrival times, weights, and an *initial* (feasible) assignment (Π^0, Φ^0).[12] While our analysis is motivated by the GDP slot reallocation problem described in Subsection 1.1, our model's primitives describe any "job rescheduling" problem in which agents (airlines) control multiple "jobs" (flights) that need to be queued. The model is particularly relevant in problems where each job f has its own earliest feasible processing time e_f.[13]

3 Rescheduling Rules and their Properties

A **rescheduling rule** is a function φ that maps each instance I to a landing schedule $\varphi(I)$ that is feasible for I. We denote by $\varphi_f(I)$ the slot to which flight f is assigned, so if $\Pi = \varphi(I)$ then $\Pi(f) = \varphi_f(I)$; we also write $\varphi_A(I) \equiv \bigcup_{f\in F_A} \varphi_f(I)$. Our objective is to find rescheduling rules that improve efficiency while respecting property rights and providing incentives for airlines to promptly and truthfully report private information.

The primary objective of the Reassignment step in GDP's is to utilize any vacant slot that could improve the position of some flight. We consider both this *non-wastefulness* property along with the stronger condition of *Pareto-efficiency*.

Definition 2. A rescheduling rule φ is **non-wasteful** if, for any instance I, there is no flight $f \in F$ and no slot $s \in S$ such that (i) $s \notin \bigcup_{g\in F} \varphi_g(I)$ (s is vacant) and (ii) $e_f \leq s < \varphi_f(I)$ (f can feasibly move up to s). It is **Pareto-efficient** if, for any instance I, there is no landing schedule Π' such

[12] The initial schedule is interpreted as the one that the FAA issues during the Ration-by-Schedule step described in Subsection 1.1.

[13] E.g. consider a manufacturing facility rescheduling jobs from various product design firms. The starting time of job (f) is constrained by how early (e_f) its firm owner can deliver that job's design specifications. A similar story can be told for research institutes submitting jobs to a shared supercomputer. Schummer and Vohra (2013) relate a similar model to the problem of organ allocation when each agent (region) controls a set of patients joining a common waiting list. Geographical fairness constraints could prevent patient f appearing earlier than position e_f in the list.

that (i) for all $A \in \mathcal{A}$, $\Pi' \succeq_A^w \varphi(I)$, and (ii) for some $A \in \mathcal{A}$, $\Pi' \succ_A^w \varphi(I)$.[14]

An even stronger definition would minimize the sum of delay costs across all airlines. Such a requirement would conflict with the incentive to truthfully report weights and with any form of property rights.

A trivial way to achieve *Pareto-efficiency* is to use serial dictatorship: allow airline A to occupy whichever slots it wants, then allow airline B to occupy its favorite of the remaining slots, etc. Aside from being unfair, such a method would disrupt operational planning for most airlines, since none (except A) could rely on keeping their initially assigned slots (following the Ration-by-Schedule step of a GDP). Thus there is an operational argument for giving airlines the option to keep their initially assigned slots. To achieve this we impose an individual rationality constraint: no airline should be worse off than they are at the initial assignment (Π^0, Φ^0).

There are two natural ways to define such a requirement in landing slot problems depending on whether one grants airlines the right to use *any* slot it initially "owns" according to Φ^0. Specifically, a weak definition of individual rationality would simply require airlines to prefer the final landing schedule to the initial one, Π^0. A stronger definition would first determine how each airline A could optimally utilize the subset of slots it intially owns, $\Phi^0(A)$, and require airlines to prefer the final landing schedule to that scenario.[15] It turns out that our negative results hold even under the weaker definition formalized as follows.

Definition 3. A rescheduling rule φ is **individually rational** if for any instance $I = (S, \mathcal{A}, (F_A)_{A\in\mathcal{A}}, e, w, \Pi^0, \Phi^0)$ and airline A we have $\varphi(I) \succeq_A^w \Pi^0$.

On the other hand our positive results would hold even under the stronger definition; see the discussion following Theorem 5.

[14] One could consider both stronger and weaker notions of efficiency. One stronger definition would minimize the *sum* of all airlines' delay costs; such a requirement conflicts with the incentive to truthfully report weights and with *individual rationality* (Definition 3). A *weak* Pareto-efficiency condition only rules out assignments from which *every* airline can *strictly* improve. This condition is less appealing in our application since there may typically be an airline whose schedule can no longer be improved at all. Nevertheless Theorem 1 and Theorem 2 hold even under *weak Pareto-efficiency* using our current proofs.

[15] In the language of Section 5, the stronger definition would first *self-optimize* the initial assignment, and then impose a standard IR condition. This definition would be justified by the regulations documented in Footnote 3.

3.1 Incentives

The various incentive properties we consider are defined by comparing the outputs of a rule before and after some change in the parameters of an instance, e.g. changes to the list of weights w, to the feasible arrival times e, etc. Therefore we use the following "replacement notation." For any instance $I = (S, \mathcal{A}, (F_A)_{A\in\mathcal{A}}, e, w, \Pi^0, \Phi^0)$, let

$$I_{w\to w'} \equiv (S, \mathcal{A}, (F_A)_{A\in\mathcal{A}}, e, w', \Pi^0, \Phi^0)$$

denote the instance identical to I except with weights w replaced with w'. Similarly $I_{e_f\to e'_f}$ denotes I with e_f replaced by e'_f, etc.

Though airline preferences are parameterized by feasible arrival times and weights, we separate the incentives to report these two types of information.[16]

Definition 4. A rescheduling rule φ is **manipulable by intentional flight delay** if there is an instance $I = (S, \mathcal{A}, (F_A)_{A\in\mathcal{A}}, e, w, \Pi^0, \Phi^0)$, airline $A \in \mathcal{A}$, flight $f \in F_A$, and $e'_f > e_f$ such that A gains from delaying f to e'_f, i.e. $\varphi(I_{e_f\to e'_f}) \succ^w_A \varphi(I)$.

History strongly motivates this concept. The FAA's previous slot allocation method was abandoned in the 1990's in part because it gave airlines a *disincentive* to report certain changes in their feasible arrival times (e_f's).[17]

Our condition has two interpretations. One is that e_f's are observable but an airline can take some private action of "sabotage" (e.g. delay the reassignment of a pilot) that commits a flight to some delay; hence the requirement $e'_f > e_f$. Another interpretation is that e_f's are private information and an airline can misreport them without detection. In this case our restriction to *downward* manipulations ($e'_f > e_f$) weakens the non-manipulability condition. This makes our conclusions stronger, however, since our results regarding this condition are negative.

Definition 5. A rescheduling rule φ is **manipulable via weights** if there is an instance $I = (S, \mathcal{A}, (F_A)_{A\in\mathcal{A}}, e, w, \Pi^0, \Phi^0)$, airline $A \in \mathcal{A}$, flight $f \in F_A$, and weight w'_f such that A gains from reporting w'_f: $\varphi(I_{w_f\to w'_f}) \succ^w_A \varphi(I)$.

[16]We consider the incentive to misreport only a *single* flight's information (e_f or w_f). Our results would continue to hold if airlines could misreport multiple flights' information.

[17]Much of the work cited in Subsection 1.2 mentions this "double penalty" problem.

The applicability of this condition depends on the degree to which delay costs are observable to the planner. For example, if one interprets weights merely as the (observable) fuel cost of keeping a particular type of aircraft waiting, then this manipulability may not be an issue. But typically, other *privately known* factors determine delay costs, such as the potential need to change exhausted flight crews, number of connecting passengers on a flight, etc., making the concept important.

Cancelations. Our third incentive property concerns the creation and reporting of vacated slots. The FAA knows that a slot has been vacated only after an airline announces it.[18] When this announcement is timely, the slot can often be given to another airline. With a sufficiently late announcement, however, the slot could go to waste (e.g. once airlines have committed to their current schedules). Even more perverse would be a situation in which an airline decides *not* to cancel a flight it otherwise would have. During times of congestion, airlines clearly should be given proper incentives to make and announce cancelations. Interestingly this concern is reflected in a 1996 US Department of Transportation memo (Oiesen (1996)), which stated

> "If an airline sits on a slot that it is not planning to use, is there any way for ATMS to detect this and to take this slot away from the airline? Should this be done?"

We consider two ways to formalize such incentives. The following, weaker one considers a scenario in which an airline gains by permanently *destroying* a slot that it initially owns.[19] A practical motivation follows the definition.

Definition 6. A rescheduling rule φ is **manipulable via slot destruction** if there is an instance $I = (S, \mathcal{A}, (F_A)_{A\in\mathcal{A}}, e, w, \Pi^0, \Phi^0)$, airline $A \in \mathcal{A}$, and slot $s \in S$ such that

(i) $s \in \Phi^0(A)$ (A owns s),

(ii) $\nexists f \in F$ such that $\Pi^0(f) = s$ (s is initially vacant), and

[18] Wambsganns (1997) explains "Airlines will send an 'open' message indicating a vacant slot is eligible [for reallocation]. This is essential, since airlines may assign different flights to their arrival slots until the slot 'times out' ... The only way for [the planner] to know if a slot is available... is by the users providing that information."

[19] Atlamaz and Klaus (2007) and Schummer and Vohra (2013) use related definitions.

(iii) $\varphi(I_{S \to S \setminus \{s\}}) \succ_A^w \varphi(I)$ (A gains when s is deleted from I).[20]

The FAA gives airlines "the capability to 'freeze' flights they don't want moved up through the submission of earliest time of arrival" (Wambsganss (1997)). By freezing flight d in slot s, airline A achieves the instance $I_{S \to S \setminus \{s\}}$) in Condition (iii). Even without this real world design constraint, one could imagine an airline A failing to announce cancelation of flight d, occupying slot s, with $e_d = s$ and w_d arbitrarily large. Seeing d as an active flight, an *individually rational* rule would keep d in slot s, effectively removing s from I, again yielding the reduced instance $I_{S \to S \setminus \{s\}}$ in which A gains in Condition (iii).

Non-manipulability via slot destruction may be too weak of a requirement since Condition (iii) removes s from the problem. Imagine a dynamic setting where rescheduling rules are applied iteratively (as done by the FAA). If airline A announces the cancelation of d after an initial run of a rescheduling rule, then what happens to slot s? When the rescheduling rule is re-run on the new instance, the FAA's rules of operation (Footnote 3) grant A first rights to use slot s for another of its flights, possibly improving A's outcome even more. We eliminate this broader form of *manipulation by postponing flight cancelations* in Definition 12, where A *hides* s as above, but then recovers s for possible consumption *after* the rule operates.

We consider both Definition 6 and Definition 12 in order to emphasize the gap between our negative and positive results. Our *DASO rules* (Subsection 5.3) satisfy the stronger non-manipulability condition. In contrast, neither *Pareto-efficient* rules (Theorem 3) nor the FAA's Compression algorithm (Schummer and Vohra (2013)) satisfy even the weaker one.

4 Efficient Rules and Manipulability

Here we show that any one of our incentive conditions is incompatible with Pareto-efficiency under the minimal constraint of *individual rationality*. These incompatibilities will motivate us to consider rules in Section 5 that solicit only the information necessary to compute *non-wasteful* assignments.

While efficiency conflicts with incentives in broader economic models, our results in this section stand out for two reasons. First, our results hold even though we have restricted ourselves to the linear-weight preference model.

[20] To be clear, $I_{S \to S \setminus \{s\}}$ is instance I but without slot s, so s is also deleted from Φ^0.

Second and more importantly, we consider the consequences of only a *single* manipulability condition at a time rather than full *strategy-proofness*. Specifically, an airline's preferences depend both on earliest arrival times (e) and flight weights (w). The analog of a full *strategy-proofness* condition in this model would allow airlines the flexibility to misreport either (or both) of these parameters. When we allow an airline to misreport only one of these variables we restrict the dimension in which an airline misreports its preferences. In this sense, the results of this section are stronger than analogous results in the matching literature. All proofs appear in the appendix.

Theorem 1. *If a rescheduling rule is* Pareto-efficient *and* individually rational*, then it is* manipulable by intentional flight delay*.*

Note that when an airline misreports feasible arrival times under Definition 4, it is required to abide by whatever landing schedule is output by the rule. A stronger definition would further allow the manipulating airline to subsequently rearrange flights amongst its assigned slots, i.e. what we call "self-optimize" in Section 5. One also could argue that such manipulation is more easily detectable (since manipulating airlines would frequently adjust their schedules in ways that initially appear infeasible or inefficient). Theorem 1 makes such arguments irrelevant since Pareto-efficient rules are manipulable *even if airlines cannot reshuffle their flights* after the mechanism has operated. A similar observation also applies to the next two theorems.

Theorem 2. *If a rescheduling rule is* Pareto-efficient *and* individually rational*, then it is* manipulable via weights*.*

The intuition behind the proof is that an efficient rule could require airline A to sacrifice a desirable slot to airline B in exchange for one *or more* of B's desirable slots elsewhere in the schedule; think of this as the "price" B pays to A. Because such a price must remain *individually rational*, A (or B) could find it beneficial to misreport weights in order to raise (or lower) this price. A related intuition proves Theorem 1.

Finally, any *Pareto-efficient* and *individually rational* rule can give an incentive to withhold slots from the system.[21]

[21] In a model of multi-object consumption with separable preferences, Atlamaz and Klaus (2007) obtain a similar conclusion. However our results are not logically related mainly for two reasons: our model's consumption constraints lead us to a different defini-

Theorem 3. *If a rescheduling rule is* Pareto-efficient *and* individually rational*, then it is* manipulable via slot destruction*.*

The proof uses an example based on a potential 3-airline trade of six slots. Airlines A and B would gain from this trade but airline C would lose. However both A and B own vacant slots that can be used to compensate C for his loss in the 3-way trade. Efficiency requires either A or B (or both) to "compensate" C to execute the trade. But by destroying its slot, A (or respectively B) can make C's compensation too high to pay (with respect to *individual rationality*), forcing the *efficient* rule to make only the other airline B (or respectively A) compensate C instead.

5 FAA-conforming Rules

The results of Section 4 motivate us to consider what we call *FAA-conforming* rules, which adhere to the way in which the FAA currently collects information and decentralizes certain scheduling decisions. Most significantly, while the FAA solicits cancelations and arrival constraints (e_f's), airlines do *not* directly report other preference information (i.e. weights). Yet weight information remains relevant since airlines may rearrange their own portions of the landing schedule. These observations imply the *FAA-conforming* constraints formalized in Subsection 5.1. We then provide a class of such rules that satisfy two of our incentive properties and a weakened version of the third. Our rules also necessarily reward an airline for canceling a flight, which is of obvious importance during periods of airport congestion.

5.1 FAA Conformation

Flight weights could be used to find efficiency improvements via both inter-airline and intra-airline trades. By comparing flight weight ratios across multiple airlines, one might find Pareto-improving (inter-airline) trades. Unfortunately Theorems 1–3 show that the execution of all such efficient trades would lead to three forms of manipulability. By using flight weights within a single airline, however, one may optimally rearrange that airline's flights

tion of Pareto-efficiency than theirs, and our preferences have additional structure imposed by the arrival time constraints (e_f's).

(intra-airline) *within its own portion* of the landing schedule. Indeed each airline is granted the right to reorder its flights by the FAA's *Facility Operation and Administration Handbook* (see Subsection 1.1). This implies a design constraint necessary to our analysis of incentives. An assignment is *self-optimized* if each airline uses *its own slots* in the best possible way.

Definition 7. An assignment (Π, Φ) is **self-optimized** (for instance I) if there exists no airline A and no landing schedule Π' such that both (i) $\Pi' \succ_A^w \Pi$ and (ii) $\Pi'(f) \in \Phi(A)$ for all $f \in F_A$. We also call a landing schedule Π *self-optimized* if it is part of a self-optimized assignment (Π, Φ) for some Φ. A rescheduling rule φ is *self-optimized* if it always outputs a *self-optimized* landing schedule.

Our motivation for this condition is strong: since an airline *has the procedural right* to rearrange its own part of the schedule, it is without loss of generality to restrict attention to *self-optimized* rescheduling rules. Any attempt to implement a *non-self-optimized* schedule would be thwarted by the airlines' right to subsequently reorder its flights.

The FAA does not directly solicit information from airlines analogous to weights.[22] Nevertheless such information is used when airlines *self-optimize*. In practice this means that a rule uses weight information *only* to *self-optimize* and not to determine the set of slots any one airline receives.

Definition 8. A rescheduling rule is **simple** if for any instances I and I' with weight profiles w and w', if $I' = I_{w \to w'}$ then for all $A \in \mathcal{A}$, $\varphi_A(I) = \varphi_A(I')$. Thus the *set* of slots consumed by an airline is invariant to changes in weights.

Besides the above institutional motivation for *simple* rules, a theoretical motivation comes from the results of Section 4. The economic value of weight (delay cost) information is reduced due to the incompatibility of efficiency and incentive compatibility. A final, practical motivation for *simplicity* is that such rules reduce the burden on airlines to compute delay cost information across all flights.

The above requirements motivate the following class of rules.

Definition 9. A rescheduling rule is **FAA-conforming** if it is *non-wasteful*, *simple*, and *self-optimized*.

[22] An exception is the Slot Credit Substitution procedure which does so implicitly but has been described as unwieldy by Robyn (2007).

5.2 Manipulability through arrival times

Since *simplicity* is a form of invariance with respect to the airlines' reported weights w, it is not surprising that it essentially rules out manipulability by weights. Trivially, any rule that is both *simple* and *self-optimized* is *non-manipulable by weights*. (More generally, *self-optimizing* any *simple* rule makes it *non-manipulable by weights*.)

Less obvious is whether such a rule is vulnerable to other forms of manipulation: by *intentional flight delay* or by *slot destruction*. The next result shows that no such rule can avoid the former, though we provide positive news for the latter in Subsection 5.3.

Theorem 4. *If an* FAA-conforming *rescheduling rule is* individually rational*, then it is* manipulable by intentional flight delay*.*

Remark 1. A subtle observation about our proof slightly strengthens its interpretation. The proof only uses manipulations by intentional flight delay under which the initial landing schedule still appears feasible. In other words, it is sufficient in the proof to consider only manipulations for which $e'_f \leq \Pi^0(f)$. The interpretation for such a restriction is that, while the central planner may not know each flight's true earliest arrival time e_f, in some applications it might be common knowledge that the initial landing schedule is feasible. Hence our result would hold even in such environments.

5.3 Deferred Acceptance with Self Optimization

To define our rules based on Deferred Acceptance, we define airline *choice functions* over sets of slots and slots' *priority orders* over airlines. We begin with the former, extending the concept in Roth (1984) to our environment where airlines care about how flights are assigned to slots.

5.3.1 Choice sets

To define choice functions, consider an airline $A \in \mathcal{A}$ with flights F_A and preferences $\succ_A^w$. How would A choose to assign its flights within some set of slots $T \subseteq S$? Assuming it can feasibly do so, determining A's "self-optimal" assignment of F_A to T typically requires knowing weights $(w_f)_{f \in F_A}$. Even without weight information, however, one can determine the *subset* of T that

A would choose to occupy. Clearly A would not want to assign some $f \in F_A$ to a slot t while leaving vacant some slot s with $e_f \leq s < t$. This necessary condition is sufficient to identify the unique subset of T that A would choose to occupy, which allows us to define choice functions as follows.

Definition 10. Fix an instance I, airline A, and set $T \subseteq S$ such that A's flights can feasibly be scheduled within T. Airline A's **choice function**, $C_A()$, over such sets $T \subseteq S$, is the output of the following simple algorithm.

- Order flights in F_A in increasing order of e_f (break ties arbitrarily).
- Assign flights sequentially to the earliest slot in T that each flight can feasibly use.
- Denote the set of occupied slots $C_A(T) \subseteq T$.

It is straightforward to see that if an airline A could assign its flights (self-optimally) within $T \subseteq S$, then its flights would occupy $C_A(T)$ in some order. It also can be shown that such choice functions satisfy the classic substitutability condition of Kelso and Crawford (1982) and Roth (1984).

5.3.2 DASO Rules

We define a class of rules based on the Deferred Acceptance algorithm, augmented with a self-optimization step. An example appears below. Each rule in the class is parameterized by an arbitrarily profile of priority orders over the set of $\mathcal{A}$, which we have fixed. For any slot $s \in \mathbb{N} = \{1, 2, \ldots\}$, a **priority order** $\gg_s$ is simply a linear order over $\mathcal{A}$. For any fixed set of slots $S \subset \mathbb{N}$ a **profile of priority orders** for S is a list $(\gg_s)_{s \in S}$.

Definition 11 (DASO rules)**.** Fixing priorities $(\gg_s)_{s \in \mathbb{N}}$ over the set of $\mathcal{A}$, the **Deferred Acceptance with Self-Optimization (DASO) rule with respect to** $(\gg_s)_{s \in \mathbb{N}}$ associates with every instance I the landing schedule computed by the following "DASO algorithm."

Step 0: Each slot $s \in S$ *proposes* to the airline A that owns it (i.e. $\Phi^0(A) \ni s$). For each $A \in \mathcal{A}$, let T_A^0 ($\equiv \Phi^0(A)$) denote the slots who proposed to $A \in \mathcal{A}$ and determine $C_A(T_A^0)$. We say that A *rejects* each slot $s \in T_A^0 \setminus C_A(T_A^0)$. If there are no rejected slots, proceed to the Self-optimization step.

Step $\boldsymbol{k = 1, 2, \ldots}$: Each slot s rejected in step $k - 1$ proposes to the highest-ranked airline in $\gg_s$ that has not already rejected s in some earlier step. (If no such airline exists, s is to be permanently unassigned.) Let T_A^k denote the slots who proposed to A in step k *plus* those in $C_A(T_A^{k-1})$. For each airline A, determine $C_A(T_A^k)$. We say that A *rejects* each slot $s \in T_A^k \backslash C_A(T_A^k)$. If there are no rejected slots, proceed to the Self-optimization step.

Self-optimization step: For each $A \in \mathcal{A}$ assign A's flights to the last $C_A(T_A^k)$ so that the resulting landing schedule is self-optimized. Break ties among equally-weighted flights by preserving their relative order in Π^0.

The DASO algorithm supplements the well-known algorithm of Gale and Shapley (1962) with two adjustments: an instance-specific adjustment of priorities in Step 0, and the addition of a *self-optimization* step. Step 0 plays two related roles: it ensures *individual rationality* by giving each airline a chance to keep its endowment,[23] and, when the initial landing schedule is feasible, it guarantees that each airline holds a *feasible* (if perhaps *wasteful*) set of slots at each interim step of the algorithm.

Example 1. Consider a model with three airlines and fix a DASO rule whose first eight slot priorities are defined as follows:

$$\begin{array}{llll} A \gg_1 B \gg_1 C & A \gg_2 B \gg_2 C & C \gg_3 B \gg_3 A & B \gg_4 A \gg_4 C \\ A \gg_5 B \gg_5 C & B \gg_6 C \gg_6 A & A \gg_7 B \gg_7 C & A \gg_8 B \gg_8 C \end{array}$$

We calculate the rule's outcome on the following instance.

Slot	Flight	Airline	Earliest	Weight
1	vacant	C		
2	a_2	A	1	w_{a_2}
3	b_3	B	1	w_{b_3}
4	c_4	C	2	w_{c_4}
5	vacant	A		
6	a_6	A	5	w_{a_6}
7	a_7	A	5	w_{a_7}
8	b_8	B	5	w_{b_8}

Step 0: Each slot proposes to its owner, and each airline rejects any slot not in its choice set. Slots $T_A^0 = \{2, 5, 6, 7\}$ propose to airline A who chooses

[23] Guillen and Kesten (2012) use a similar "propose to owners first" idea.

$C_A(T^0_A) = \{2,5,6\}$, rejecting slot 7. Similarly slots $\{3,8\}$ propose to airline B who rejects neither, and slots $\{1,4\}$ propose to airline C who rejects slot 1.

Step 1: Each rejected slot proposes to the highest priority airline that has not yet rejected it: slot 1 proposes to airline A and slot 7 proposes to B. From the set $T^1_A = \{1,2,5,6\}$, A chooses $C_A(T^1_A) = \{1,5,6\}$, rejecting slot 2. Similarly B rejects slot 8 from $T^1_B = \{3,7,8\}$.

Step 2: The remaining steps are similar. Slot 2 proposes to airline B, who now rejects slot 3. Slot 8 proposes to its final airline C, who rejects it.

Step 3: Slot 3 proposes to airline C, who now rejects slot 4.

Step 4: Slot 4 proposes to airline B, who rejects it.

Step 5: Slot 4 proposes to airline A, who rejects it.

Self-Optimization step: Each airline's flights are self-optimally assigned to its current slots. In this example, feasibility uniquely determines the slot of all flights except a_6 and a_7. If $w_{a_6} \geq w_{a_7}$ then a_6 is assigned to slot 5, otherwise a_7 is. Assuming $w_{a_6} < w_{a_7}$ for example, the rule outputs

Slot	Flight	Slot	Flight
1	a_2	5	a_7
2	b_3	6	a_6
3	c_4	7	b_8
4	vacant	8	vacant

as its final landing schedule.[24]

It may seem strange that we use a "slot-proposing" Deferred Acceptance algorithm. An "airline-proposing" version of the algorithm would have airlines proposing to their favorite sets of slots, and slots accepting proposals only from their highest-priority airline. Such a formulation would seem more appropriate within the matching literature since Deferred Acceptance algorithms favor the proposing side of the market. In our model, however, it turns out that *both versions of the algorithm produce precisely the same outcome.*

[24] In practice, without possessing information on relative weights (w_f's), the FAA would arbitrarily assign those two flights to slots 5 and 6. The airline would then request a swap only if preferred. This illustrates the low informational requirement of *simple* rules: when self-optimization decisions can be decentralized, the planner does not require full airline preference information, e.g. in the form of weights.

Observation 1. Fixing priorities and an instance, an airline-proposing version of the DASO algorithm would yield precisely the same outcome as the slot-proposing version described above.

This follows from an induction argument. If the algorithms yield identical outcomes on the first $s-1$ slots, then slot s is the "best" remaining slot for those airlines that still have use for s. The highest-ranked such airline in $\gg_s$ must receive s under either algorithm.[25] The slot-proposing algorithm simplifies some proofs and has the property that, throughout its execution, every *interim* allocation of slots to airlines yields a feasible schedule.

5.3.3 Results

The following properties of DASO rules are trivial to prove.

Theorem 5. *For any profile of priorities* $(\gg_s)_{s\in\mathbb{N}}$ *over* $\mathcal{A}$*, the corresponding DASO rule is* non-manipulable via weights, individually rational, *and* FAA-conforming.

In fact DASO rules satisfy the stronger version of *individual rationality* mentioned in Section 3: each airline A weakly prefers the outcome to what it could get by self-optimizing its own "endowment" $\Phi^0(A)$.

Our main, positive result is that DASO rules induce prompt, truthful reporting of flight cancelations. Going beyond *non-manipulability via slot destruction* (Definition 6), we show that an airline cannot manipulate a DASO rule by temporarily "hiding" a vacant slot then later retrieving it for consumption. To define this stronger condition we change part (iii) of Definition 6 by allowing A to further improve its outcome using the "destroyed" slot s. Below we discuss an alternative formulation that yields equivalent conclusions.

Definition 12. A rescheduling rule φ is **manipulable by postponing a flight cancelation** if there is an instance $I = (S, \mathcal{A}, (F_A)_{A\in\mathcal{A}}, e, w, \Pi^0, \Phi^0)$, airline $A \in \mathcal{A}$, and slot $s \in S$ such that

[25] See the Online Appendix. This argument proves an analogous result when agents in a classic matching market have aligned preferences. Airline preferences are only partially aligned due to arrival constraints. The proof also suggests an alternative DASO algorithm suggested to us by Utku Ünver: *sequentially* assign each slot to a (remaining) flight in accordance with that slot's priorities, then self-optimize the resulting schedule.

(i) $s \in \Phi^0(A)$ (A owns s),

(ii) $\nexists f \in F$ such that $\Pi^0(f) = s$ (s is initially vacant), and

(iii) $\exists\, \Pi \succ_A^w \varphi(I)$ such that $\bigcup_{f \in F_A} \Pi(f) \subset \left(\bigcup_{f \in F_A} \varphi_f(I_{S \to S \setminus \{s\}})\right) \cup \{s\}$.

The next result shows that DASO rules are non-manipulable in this sense, and that if an airline does manipulate in this way, then *no* airline is made better off. Such a strong result is important in environments where group incentives are relevant. Even if one airline could compensate another to postpone a cancelation, this ability would be of no use.[26] By iteration, the result would hold even if airlines could postpone multiple cancelations.

Theorem 6. *For any set of airlines $\mathcal{A}$ and priorities $(\gg_s)_{s \in \mathbb{N}}$, the corresponding DASO rule is* non-manipulable by postponing a flight cancelation. *Furthermore if an airline were to postpone a cancelation, then each flight (of* any *airline) would receive a weakly later slot, making all airlines (weakly) worse off.*

The first part of the result is related to one of Crawford (1991), showing that the addition of a student to a College Admissions model benefits all colleges under (either) Deferred Acceptance algorithm.[27] This can be shown to imply that *slot destruction* (a là Definition 6) would make all airlines weakly worse off under any DASO rule. Theorem 6 strengthens that conclusion both by (i) allowing an airline to recover and consume the destroyed slot, and (ii) showing that airlines are worse off on a flight-by-flight basis.

Theorem 6 is robust in two important ways. First, the result would continue to hold in a *dynamic* setting where rescheduling rules are iteratively applied as airlines continually report updated information about cancelations, as in practice. Consider a 2-period model where a DASO rule is to be applied twice. In period one the rule creates an interim schedule based on the first-period reports. In period two airlines learn and report new cancelations, and the DASO rule is applied to period one's output. It can be shown that no airline can benefit by postponing a period one (or two) flight cancelation.

Second, Theorem 6 is robust to airlines *choosing* whether or not to cancel flights. While Definition 12 models the postponement of *exogenously determined* cancelations—s is predetermined to be vacant—the following obser-

[26] This immunity to manipulation with compensation contrasts with impossibility results in more standard revelation games (e.g. Schummer (2000)).

[27] See also Kelso and Crawford (1982) and Kojima and Manea (2010).

vation concerns the similar (but logically distinct) idea of *endogenous* cancelations. Namely, the conclusions of Theorem 6 hold whenever any *viable* flight $f \in F$ is removed from the instance, formalized as follows.[28]

Observation 2 (endogenous cancelations)**.** Under any DASO rule and for any instance I, the deletion of any flight $f \in F$ would assign each remaining flight $g \in F \setminus \{f\}$ to a (weakly) earlier slot.

This concept of flight deletion resembles, but is distinct from, the idea of *capacity manipulation* in the classic College Admissions model (Sönmez (1997)). Konishi and Ünver (2006) show that reducing a college's capacity weakly benefits the *other* colleges. A flight deletion similarly reduces an airline's "capacity" to consume slots; however *it also simultaneously changes the airline's preferences for slots* as a function of *which* flight was deleted making the two concepts logically independent.[29]

Finally, since DASO rules are not *Pareto-efficient* it is natural to ask whether any other *FAA-conforming* rule could *Pareto-improve* upon DASO while achieving the same incentive properties. This turns out to be impossible *even if we ignore incentives.* Pareto-improving upon a DASO rule would require the use of a *non-simple* rule.[30]

Proposition 1 (no Pareto-dominance)**.** *For any DASO rule, there exists no other* simple *rule that makes every airline weakly better off at every instance.*

5.4 Weak Incentives and Self-optimization

While the classic *strategy-proofness* condition makes it a dominant strategy to truthfully report *all* preference information, Definition 4 does so merely for the arrival time "dimension" of preferences. Nevertheless Theorems 1 and 4 show that even this narrow requirement is difficult to achieve. It turns out that DASO rules satisfy a milder concept which we now define: no airline can misreport arrival times in a way that would benefit *all* of its flights.

[28] We thank referees for comments leading to this robustness check of Theorem 6. Proof is in the Online Appendix.

[29] Furthermore we obtain stronger conclusions due to our model's structure, though they can be proven using the approach of Konishi and Ünver (2006). Separately, Ehlers (2010) shows that student-proposing DA satisfies a weaker capacity non-manipulation condition.

[30] We thank a referee for conjecturing this kind of result. Proof is in the Online Appendix.

Definition 13 (Schummer and Vohra (2013)). A rescheduling rule φ is **weakly non-manipulable via earliest arrival times** if for any instance $I = (S, \mathcal{A}, (F_A)_{A\in\mathcal{A}}, e, w, \Pi^0, \Phi^0)$, airline $A \in \mathcal{A}$, flight $f \in F_A$, and earliest arrival time e'_f we have

$$[\exists g \in F_A \text{ with } \varphi_g(I_{e_f \to e'_f}) < \varphi_g(I)] \Longrightarrow [\exists h \in F_A \text{ with } \varphi_h(I_{e_f \to e'_f}) > \varphi_h(I)].$$

Practically speaking, this condition implies that an airline cannot determine whether a manipulation is beneficial unless it goes to the trouble of determining its flight weight information. This weak notion[31] yields the following positive result, contrasting the negative result of Theorem 4.

Theorem 7. *For any set of airlines $\mathcal{A}$ and priorities $(\gg_s)_{s\in\mathbb{N}}$, the corresponding DASO rule is* weakly non-manipulable via earliest arrival times.

Schummer and Vohra (2013) show that two other rules satisfy this condition: the FAA's Compression Algorithm and their "TC" rule, a Top-Trading-Cycle variant. They do so in a model ignoring flight weights, so *self-optimization* is not part of their analysis. Since real world airlines *are* permitted to *self-optimize* their landing schedules, it is natural for us to ask whether Schummer and Vohra's results are preserved if we supplement Compression (or TC) with a *self-optimization* step.

Interestingly, Compression and TC remain weakly non-manipulable only if the initial schedule is self-optimized *before* the rule is applied. These rules lose the weak non-manipulability property if self-optimization is done only *after* the rule operates as we illustrate below. More precisely both the "*self-optimize-then-Compression*" and "*self-optimize-then-TC*" rules are weakly non-manipulable via earliest arrival times, but "*Compression-then-self-optimize*" and "*TC-then-self-optimize*" fail the condition.

While the formalization of these results is relegated to the Online Appendix, we provide an example to illustrate the intuition for the manipulability of a *Compression-then-self-optimize* rule.[32]

Example 2. Manipulation of a *Compression-then-self-optimize* rule.

[31] Definition 13 does not technically weaken *non-manipulability by intentional flight delay* because it allows $e'_f < e_f$. However this makes Theorem 7 a stronger result.

[32] The same example can be used to illustrate manipulability of *TC-then-self-optimize*.

Slot	Flight	Airline	Earliest	Weight
1	vacant	C		
2	a_2	A	2	1
3	b_3	B	1	5
4	a_4	A	1	4

The Compression algorithm works by "trading" each airline's unusable slot to the next flight in the schedule that can feasibly use it.[33] Here, airline C trades slot 1 for slot 3 since flight b_3 is the next earliest one that can use slot 1. Next, since C also cannot use slot 3, it is traded to a_4 (in exchange for slot 4). Following a self-optimization step, flights a_4 and a_2 end up in slots 2 and 3 respectively.

If A misreports e_{a_2} to be $e'_{a_2} = 1$, one can verify that Compression's trading steps would ultimately assign a_2 and a_4 to slots 1 and 3 respectively. Since $e_{a_2} = 2$, this would be infeasible for a_2, but a self-optimization step at this point would swap a_4 and a_2 into slots 1 and 3 respectively, which satisfies the true arrival time constraints. In particular, this misreport results in a strong manipulation: none of A's flights loses and one strictly gains.

The example illustrates how the initial order of A's flights provides an opportunity for manipulation. When slot 1 is allocated by Compression, flight a_2 is ineligible to receive it (by feasibility), and b_3 has precedence over a_4 because b_3 sits earlier in the schedule. By misreporting e_{a_2}, however, a_2 appears able to feasibly use slot 1 and has precedence over b_3, giving A the slot. While it is actually infeasible for a_2 to use that slot, a self-optimization step finally moves a_2 to slot 3, fixing the problem.

Interestingly, if the initial schedule is self-optimized *before* running Compression, the outcome changes. Flight a_4 (starting in slot 2) has precedence over b_3 and receives slot 1 immediately. In this scenario A would have no strong manipulation.[34] Thus not only is Compression sensitive to the initial ordering of flights—a fact pointed out by Vossen and Ball (2006a)—but *how* that initial ordering is determined can affect whether the mechanism is weakly non-manipulable.

[33] See Schummer and Vohra (2013) for a complete formalization which is unneeded here.

[34] This follows from our general result that *self-optimize-then-Compression* is weakly non-manipulable; see Online Appendix.

On the other hand, DASO rules are invariant to the initial ordering. The input to a DASO algorithm (Definition 11) includes the slot ownership function (Φ^0) but not the relative orderings specified in the initial schedule (Π^0). Since the outcomes of DASO rules are independent of Π^0, their weak non-manipulability has nothing to do with whether Π^0 is self-optimized. Since airlines have the right to arrange their landing schedules at any time, we argue that DASO rules satisfy Definition 13 more robustly than the other two rules.

6 Summary

Airport landing schedules—created far in advance of their actual execution—can become inefficient or even infeasible due to unforeseen weather events and subsequent flight cancelations or delays. The information needed to reschedule flights must be elicited from individual airlines, yielding the mechanism design problem analyzed here. We separately considered the airlines' incentives to report three kinds of relevant information: flights' feasible arrival times, cancelations, and relative delay costs.

The FAA does not elicit delay costs, which would be necessary to determine *Pareto-efficient* schedules. An incentive-based justification for this design choice is given by our first three results, showing that any *Pareto-efficient* rule is manipulable in *each* of the three ways mentioned above.

This motivates our study of *FAA-conforming rules* which use delay cost information only to reorder an airline's own flights among its assigned slots. Our main result provides a class of such (DASO) rules that induce airlines to both execute and promptly report flight cancelations. This contrasts with the potential disincentive to do so under the FAA's Compression algorithm (Schummer and Vohra (2013)), and emphasizes a distinction between incentives to report *cancelations* and to report *delays*—two concepts that have been intermingled in previous work. Though all FAA-conforming rules are manipulable by misreporting delays, DASO rules provide a *weak* incentive to report delays more robustly than do Compression and the TC rule of Schummer and Vohra (2013). DASO rules thus dominate the other rules that have been studied for this problem in terms of our incentive properties. Finally, while DASO rules cannot be Pareto-efficient, they cannot be

Pareto-dominated by any other *simple* rule (Proposition 1).

Though we formalize an incentive for the timely report of *exogenously* determined flight cancelations, it is also important to encourage the endogenous *decision* to cancel otherwise viable flights. This is particularly true during congestion (e.g. GDP's), when a cancelation benefits the remaining flights by reducing total delay. Hence it is significant that DASO rules always reward cancelation *decisions* by improving the slot of each remaining flight (Observation 2). Neither the Compression algorithm (Schummer and Vohra (2013)) nor any Pareto-efficient rule (Theorem 3) possesses this property.

It is natural to ask *which* DASO rules (i.e. which parameters $(\gg_s)_{s\in S}$) best accomplish any particular objective. While our results hold for any choice of these parameters, fairness considerations could justify prioritizing airlines by various factors (number of flights, total delay suffered in the GDP, etc.), or to rotate priorities by slot, to randomize priorities, etc. Depending on the distribution of primitives, certain priority structures could conceivably provide a greater average reward for airlines who decide to cancel flights;[35] however no DASO rule is Pareto-dominated by another. Since any optimization exercises within the class of DASO rules would be sensitive to empirical assumptions on the population of instances I, such future work would require historical data on GDP's, including initial schedules, feasible arrival times, and reasonable assumptions on the airlines' privately known delay costs.

References

Abdulkadiroğlu, Atila, and Tayfun Sönmez, 1999. "House Allocation With Existing Tenants." *Journal of Economic Theory*, 88(2):233–260.

Abdulkadiroğlu, Atila, and Tayfun Sönmez, 2003. "School choice: a mechanism design approach." *American Economic Review*, 93:729–747.

Alcalde, José, and Salvador Barberà, 1994. "Top Dominance and the Possibility of Strategy-Proof Stable Solutions to Matching Problems." *Economic Theory*, 4(3):417–35.

Atlamaz, Murat, and Bettina Klaus, 2007. "Manipulation via Endowments in Exchange Markets with Indivisible Goods." *Social Choice and Welfare*, 28(1):1–18.

[35] E.g. Ergin (2002) and Kesten (2006) eliminate Deferred Acceptance's inefficiency in a unit demand model by making priorities acyclic. In our multi-unit demand model, however, DASO rules cannot be fully efficient (combining Theorem 3 and Theorem 6).

Balikrishnan, Hamsa, 2007. “Techniques for Reallocating Airport Resources during Adverse Weather.” In Proc. of the IEEE Conference on Decision and Control.

Balinski, Michel and Tayfun Sonmez, 1999. “A Tale of Two Mechanisms: Student Placement.” *Journal of Economic Theory*, 84:73–94.

Ball, Michael, Geir Dahl, and Thomas Vossen, 2009. “Matchings in Connection with Ground Delay Program Planning.” *Networks*, 53(3):293–306.

Ball, Michael, Robert Hoffman, and Avijit Mukherjee, 2010. “Ground Delay Program Planning Under Uncertainty Based on the Ration-by-Distance Principle.” *Transportation Science*, 44(1):1–14.

Ball, Michael and Guglielmo Lulli, 2004. “Ground Delay Programs: Optimizing over the Included Flight Set Based on Distance.” *Air Traffic Control Quarterly*, 12(1):1–25.

Crawford, Vincent, 1991. “Comparative statics in matching markets.” *Journal of Economic Theory*, 54(2):389–400.

Dubins, Lester E. and David A. Freedman, 1981. “Machiavelli and the Gale-Shapley Algorithm.” *American Mathematical Monthly*, 88:485–494.

Ehlers, Lars, 2010. “Manipulation via capacities revisited.” *Games and Economic Behavior*, 69:302–311.

Ergin, Haluk, 2002. “Efficient Resource Allocation on the Basis of Priorities.” *Econometrica*, 70:2489–2497.

Gale, David and Lloyd Shapley, 1962. “College admissions and the stability of marriage.” *American Mathematical Monthly*, 69:9–15.

Guillen, Pablo and Onur Kesten, 2012. “Mathcing markets with mixed ownership: The case for a real-life assignment mechanism.” *International Economic Review*, 53(3):1027–1046.

Hoffman, Robert and Michael Ball, 2007. “A Comparison of Formulations for the Single-Airport Ground-Holding Problem with Banking Constraints.” *Operations Research*, 48(4):578–590.

Kelso, Alexander and Vincent Crawford, 1982. “Job Matching, Coalition Formation and Gross Substitutes.” *Econometrica*, 50(6):1483–1504.

Kesten, Onur, 2006. “On two competing mechanisms for priority-based assignment problems.” *Journal of Economic Theory*, 127(1):155–171

Kojima, Fuhito and Mihai Manea, 2010. “Axioms for Deferred Acceptance.” *Econometrica*, 78(2):633–653.

Kominers, Scott, and Tayfun Sönmez, 2016. “Matching with Slot-Specific Priorities: Theory.” *Theoretical Economics*, 11(2):683–710.

Konishi, Hideo, Thomas Quint, and Jun Wako, 2001. "On the Shapley-Scarf Economy: The Case of Multiple Types of Indivisible Goods." *Journal of Mathematical Economics*, 35(1):1–15.

Konishi, Hideo and Utku Ünver, 2006. "Games of Capacity Manipulation in Hospital-Intern Markets." *Social Choice and Welfare*, 27:3–24.

Niznik, Tim, 2001. "Optimizing the Airline Response to Ground Delay Programs." *AGIFORS Airline Operations Conference Proceedings*, Ocho Rios, Jamaica.

Oiesen, Rick, 1996. "Issues Concerning the Compression Algorithm." Memorandum, US Department of Transportation.

Postlewaite, Andrew, 1979. "Manipulation via Endowments." *The Review of Economic Studies*, v(i):255–262.

Robyn, Dorothy, 2007. "Reforming the Air Traffic Control System to Promote Efficiency and Reduce Delays." Report to Council of Economic Advisers, The Brattle Group, Inc.,

Roth, Alvin, 1982. "The Economics of Matching: stability and incentives," *Mathematics of Operations Research*, 7(4):617–628.

Roth, Alvin, 1984. "Stability and Polarization of Interests in Job Matching," *Econometrica*, 52:47–57.

Roth, Alvin, 1985. "The College Admissions Problem is not Equivalent to the Marriage Problem," *Journal of Economic Theory*, 36:277–288.

Roth, Alvin, and Marilda Sotomayor, 1990. *Two-sided matching: A study in game-theoretic modeling and analysis*. Econometric Society Monograph 18. Cambridge, Cambridge University Press.

Schumer, Charles, Carolyn Maloney, 2008. *Your Flight Has Been Delayed Again*. Report by the Joint Economic Committee Majority Staff, U.S. Senate.

Schummer, James, 2000. "Manipulation through Bribes," *Journal of Economic Theory*, 91:180–198.

Schummer, James, and Rakesh V. Vohra, 2013. "Assignment of Arrival Slots," *American Economic Journal: Microeconomics*, 5(2):164–185.

Sertel, Murat, and Ipek Özkal-Sanver, 2002. "Manipulability of the Men-(Women-) Optimal Matching Rule via Endowments," *Mathematical Social Sciences*, 44(1):65–83.

Shapley, Lloyd and Herbert Scarf, 1974. "On Cores and Indivisibility," *Journal of Mathematical Economics*, 1:23–37.

Sönmez, Tayfun, 1996. "Strategy-Proofness in Many-to-One Matching Problems," *Review of Economic Design*, 1(1):365–380.

Sönmez, Tayfun, 1997. “Manipulation via Capacities in Two-Sided Matching Markets,” *Journal of Economic Theory*, 77:197–204.

Sönmez, Tayfun, 2013. “Bidding for Army Career Specialties: Improving the ROTC Branching Mechanism,” *Journal of Political Economy*, 121(1):186–219.

Sönnez, Tayfun, and Tobias B. Switzer, 2013. “Matching with (Branch-of-Choice) Contracts at the United States Military Academy.” *Econometrica*, 81(2):451–488.

Takagi, Shohei and Shigehiro Serizawa, 2010. “An Impossibility Theorem for Matching Problems.” *Social Choice and Welfare*, 35(2):245–266.

Vossen, Thomas, and Michael Ball, 2006a. “Optimization and mediated bartering models for ground delay programs.” *Naval Research Logistics*, 53(1):75–90.

Vossen, Thomas, and Michael Ball, 2006b. “Slot Trading Opportunities in Collaborative Ground Delay Programs.” *Transportation Science*, 40(1):29–43.

Wambsganss, Michael, 1997. “Collaborative decision making through dynamic information transfer.” *Air Traffic Control Quarterly*, 4:107–123.

A Appendix: Proofs

Proof of Theorem 1. Let φ be an *individually rational*, *Pareto-efficient* rule and consider the following instance I.

Slot	Flight	Airline	Earliest	Weight
1	b_1	B	1	2
2	a_2	A	1	6
3	a_3	A	1	4
4	b_4	B	3	3
5	a_5	A	5	3.5
6	b_6	B	5	3.5

By feasibility and *Pareto-efficiency*, φ must assign flights a_5 and b_6 to slots 5 and 6 (in some order), and assign the remaining flights to slots 1–4 (in some order); additionally b_4 must be assigned to slot 3 or 4. With the further requirement of *individual rationality*, however, airline B must receive one of the following four sets of slots: $\{1,4,6\}$, $\{2,3,6\}$, $\{2,4,5\}$, or $\{3,4,5\}$. To see this, suppose B's flights are assigned to slots $\{2,4,6\}$. Then b_1 is the only one of B's flights to receive a different (furthermore worse) slot than in the initial schedule, violating *individual rationality*. If B receives $\{1,4,5\}$, then A similarly loses out via flight a_5. On the other hand if B's flights are (efficiently) assigned to slots $\{3,4,5\}$, then b_1 is assigned slot 4 (for a loss of 6 units), b_4 is assigned slot 3 (gain of 3 units), and b_6 is assigned slot 5 (gain 3.5 units), for a net gain of $-6+3+3.5=0.5$. Similarly one can check that A also gains from receiving slots $\{1,2,6\}$. All of the remaining cases follow from similar calculations.

Next observe that two of these four *individually rational* landing schedules are Pareto-dominated. The initial landing schedule (where B receives $\{1,4,6\}$) is Pareto-dominated by the landing schedule in which B's flights are assigned to slots $\{2,4,5\}$ and A's flights are (efficiently) assigned to slots $\{1,3,6\}$. (B's payoff improves by $-2+0+3.5=1.5$ and A's improves by $6+0-3.5=2.5$.) Similarly, the schedule in which B receives slots $\{2,3,6\}$ is Pareto-dominated by the same alternative schedule. In summary we must have $\varphi_B(I) \in \{\{2,4,5\},\{3,4,5\}\}$. We show that in either case φ is manipulable by an intentional flight delay.

Case 1: $\varphi_B(I) = \{2, 4, 5\}$. In this case a_2 is assigned to slot 1, a_3 is assigned to slot 3, and a_5 is assigned to slot 6.

Let I' denote the instance that is identical to I except that A intentionally delays a_2 with $e'_{a_2} = 2$ (without altering the other e_f's). There is only one *Pareto-efficient* and *individually rational* assignment for I': a_3 is assigned to slot 1, a_2 is assigned to slot 2, and a_5 is assigned to slot 6. Airline A prefers this assignment to the one he receives under $\varphi(I)$. Therefore A is able to gain from misreporting e_{a_2}.

Case 2: $\varphi_B(I) = \{3, 4, 5\}$. In this case b_4 is assigned to slot 3, b_1 is assigned to slot 4, and b_6 is assigned to slot 5.

Let I' denote the instance that is identical to I except that B intentionally delays b_6 with $e'_{b_6} = 6$ (without altering the other e_f's). There is only one *Pareto-efficient* and *individually rational* assignment for I': b_1 is assigned to slot 2, b_4 is assigned to slot 3, and b_6 is assigned to slot 6. Airline B prefers this assignment to the one he receives under $\varphi(I)$. □

Proof of Theorem 2. Let φ be an *individually rational*, *Pareto-efficient* rule and consider the following instance I.

Slot	Flight	Airline	Earliest	Weight
1	b_1	B	1	1
2	a_2	A	1	4
3	a_3	A	2	4
4	a_4	A	4	3
5	b_5	B	4	3

One can verify that there are only two *Pareto-efficient* and *individually rational* landing schedules for I using the kind of calculations done in the previous proof. In one, B's flights are assigned to slots 2 and 4; in the other, they are assigned slots 3 and 4. Clearly B prefers the former and A prefers the latter. We show that regardless of which is selected by φ, the rule is manipulable.

Case 1: $\varphi_B(I) = \{3, 4\}$. Let I' denote the instance that is identical to I except that B reports a weight of $w'_{b_1} = 2$ (without altering the other weights). There is only one *Pareto-efficient* and *individually rational* assignment for I': B is assigned slots 2 and 4. Since $\varphi_B(I') = \{2, 4\} \succ_B^w \{3, 4\} = \varphi_B(I)$, airline B gains by misrepresenting w_1 at I.

Case 2: $\varphi_B(I) = \{2, 4\}$. Let I' denote the instance that is identical to I except that A reports a weight of $w'_{a_2} = 1.5$ (without altering the other weights). There is only one *Pareto-efficient* and *individually rational* assignment for I': B is assigned slots 3 and 4. Since $\varphi_A(I') = \{1, 2, 5\} \succ_A^w \{1, 3, 5\} = \varphi_A(I)$, airline A gains by misrepresenting w_2 at I. □

Proof of Theorem 3. Consider the instance I described in Table 1. Note that airline D plays the role of a "dummy airline," in that D's flights already occupy their most preferred slots. *Individual rationality* thus forces a rule not to move any of D's flights.

Slot	Flight	Airline	Earliest	Weight
1	a_1	A	1	3
2	b_2	B	1	4
3	b_3	B	3	3
4	c_4	C	3	2
5	c_5	C	5	3
6	a_6	A	5	4
7	a_7	A	7	0.3
8	d_8	D	8	1
9	d_9	D	9	1
10	vacant	A		
11	c_{11}	C	7	0.35
12	b_{12}	B	12	0.3
13	d_{13}	D	13	1
14	d_{14}	D	14	1
15	vacant	B		
16	c_{16}	C	16	0.35

Table 1. Main example in the proof of Theorem 3.

Within slots 1–6, three pairwise trades amongst airlines A, B, and C are possible. If all three are performed, A and B each gain 1 unit while C loses 1 unit. There are two ways to compensate C for this loss. One is for a_7 to move down to slot 10, giving c_{11} slot 7 (though A is not willing to move down to slot 11). Similarly B could compensate C via slot 12. We show that if A is compensating C at slot 7, then A gains by destroying slot 10. A similar argument applies to B and slot 15.

Let φ be an *individually rational*, *Pareto-efficient* rule. It can be checked that only the three landing schedules in Table 2a satisfy the conditions cor-

Slot	Π^1	Π^2	Π^3
1	b_2	b_2	b_2
2	a_1	a_1	a_1
3	c_4	c_4	c_4
4	b_3	b_3	b_3
5	a_6	a_6	a_6
6	c_5	c_5	c_5
7	a_7	c_{11}	c_{11}
8	d_8	d_8	d_8
9	d_9	d_9	d_9
10	c_{11}	a_7	a_7
11	–	–	–
12	c_{16}	b_{12}	c_{16}
13	d_{13}	d_{13}	d_{13}
14	d_{14}	d_{14}	d_{14}
15	b_{12}	c_{16}	b_{12}
16	–	–	–
A	+1	+0.1	+0.1
B	+0.1	+1	+0.1
C	+0.75	+0.75	+1.8

(a)

Slot	Π^4	Π^5
1	b_2	b_2
2	a_1	a_1
3	c_4	c_4
4	b_3	b_3
5	a_6	a_6
6	c_5	c_5
7	a_7	c_{11}
8	d_8	d_8
9	d_9	d_9
10	(destroyed)	a_7
11	c_{11}	–
12	c_{16}	b_{12}
13	d_{13}	d_{13}
14	d_{14}	d_{14}
15	b_{12}	(destroyed)
16	–	c_{16}
A	+1	+0.1
B	+0.1	+1
C	+0.4	+0.4

(b)

Table 2. Proof of Theorem 3. (a) The three possible landing schedules chosen by an *individual rational*, *Pareto-efficient* rule for I. They differ only in slots 7, 10, 12, and 15. (b) Two landing schedules following slot destruction in I.

responding to *individual rationality* and *Pareto-efficiency* for instance I, so $\varphi(I)$ must be one of them. The bottom of the table shows the relative gain in terms of weights for each airline, relative to the initial landing schedule.

We show that regardless of which of the three landing schedules is selected by φ, airline A or airline B can manipulate φ by destroying its vacant slot. First suppose that $\varphi(I) \in \{\Pi^2, \Pi^3\}$. Let $I_{S\to S\setminus\{10\}}$ be the instance obtained from I by destroying A's slot 10. There is only one landing schedule for $I_{S\to S\setminus\{10\}}$ that satisfies the conditions for *Pareto-efficiency* and *individual rationality*. It is Π^4 described in Table 2b.

Therefore $\varphi(I_{S\to S\setminus\{10\}}) = \Pi^4$. However, A's gain of one unit is greater than his gain at Π^2 or Π^3 above, i.e. $\varphi_A(I_{S\to S\setminus\{10\}}) \succ_A^w \varphi_A(I)$. Since A gains by destroying slot 10, φ is manipulable via slot destruction in this case.

Therefore we must have $\varphi(I) = \Pi^1$. Similar to the argument above, let $I_{S \to S \setminus \{15\}}$ be the instance obtained from I when airline B destroys slot 15. There is only one landing schedule for $I_{S \to S \setminus \{15\}}$, namely Π^5, that satisfies the conditions for *Pareto-efficiency* and *individual rationality*.

Therefore $\varphi(I_{S \to S \setminus \{15\}}) = \Pi^5$. However, B's gain of one unit is greater than its gain at Π^1 above, i.e. $\varphi_B(I_{S \to S \setminus \{15\}}) \succ_B^w \varphi_B(I)$. Since airline B gains by destroying its vacant slot 15, φ is manipulable via slot destruction. $\square$

Proof of Theorem 4. Consider instances I and I' defined by weights w and w' in the following table.

Slot	Flight	Airline	Earliest	Weight w_f	Weight w'_f
1	vacant	A			
2	vacant	A			
3	b_3	B	1	1	1
4	c_4	C	1	1	1
5	b_5	B	2	5	1
6	c_6	C	3	5	5
7	b_7	B	4	1	7

Let φ be a *simple*, *non-wasteful* rescheduling rule. Denote the three slots assigned to B's flights by $\varphi_B(I) = \{\varphi_{b_3}(I), \varphi_{b_5}(I), \varphi_{b_7}(I)\}$. Since φ is *non-wasteful*, it must assign all five flights to slots 1–5, so $\varphi_B(I) \subset \{1, 2, 3, 4, 5\}$. Define $\varphi_C(I)$ and $\varphi_B(I')$ similarly. Since φ is *simple*, note that $\varphi_B(I') = \varphi_B(I)$ (although $\varphi_f(I) \neq \varphi_f(I')$ is certainly possible).

There are $\binom{5}{3} = 10$ candidate subsets to consider for $\varphi_B(I)$, but B's flights cannot feasibly be assigned to $\{1, 2, 3\}$ or $\{3, 4, 5\}$. We show that for each of the remaining eight possibilities, φ must be manipulable in some way.

Case 1: $\boldsymbol{\varphi_B(I)}$ is $\boldsymbol{\{1,3,4\}}$, $\boldsymbol{\{1,3,5\}}$, or $\boldsymbol{\{1,4,5\}}$. In this case $\varphi_{b_5}(I) \geq 3$. From instance I, consider airline B delaying b_3 to $e'_{b_3} = 2$ (resulting in $\tilde{I} = I_{e_{b_3} \to e'_{b_3}}$). By *non-wastefulness*, $\varphi_{c_4}(\tilde{I}) = 1$. With *simplicity* (and *self-optimization*) this yields $\varphi_{b_5}(\tilde{I}) = 2$. This improvement for b_5 gives B a relative gain of at least 5 (weight) units. It is simple to check that B can lose at most 4 units combined through changes in the assignment of b_3 and b_7. Since B would gain from such a manipulation, this rules out Case 1. In all remaining cases, B must receive slot 2.

Case 2: $\varphi_B(I)$ is $\{2,3,4\}$ or $\{2,3,5\}$. In this case $\varphi_C(I)$ is $\{1,5\}$ or $\{1,4\}$, so $\varphi_{c_4}(I) = 1$ and $\varphi_{c_6}(I) \geq 4$. From instance I, consider airline C delaying c_4 to $e'_{c_4} = 3$ (resulting in $\tilde{I} = I_{e_{c_4} \to e'_{c_4}}$). By *non-wastefulness* B receives the first two slots, and either $\varphi_{c_4}(\tilde{I}) = 3$ or $\varphi_{c_6}(\tilde{I}) = 3$, with *self-optimization* implying the latter. This improvement for c_6 gives C a relative gain of at least 5 (weight) units, while C can lose at most 4 units through the change in the assignment of c_6. Since C would gain from such a manipulation, this rules out Case 2. In all remaining cases, C must receive slot 3.

Case 3: $\varphi_B(I) = \{2,4,5\}$. In this case $\varphi_{b_5}(I) = 2$, while b_3 and b_7 go to slots 4 and 5 (the order being irrelevant). From instance I, consider airline B delaying b_3 to $e'_{b_3} = 3$ (resulting in $\tilde{I} = I_{e_{b_3} \to e'_{b_3}}$). By *non-wastefulness*, $\varphi_{c_4}(\tilde{I}) = 1$ and $\varphi_{b_5}(\tilde{I}) = 2$.

We also show $\varphi_{b_3}(\tilde{I}) = 3$. Suppose not, so $\varphi_{b_3}(\tilde{I}) > 3$. Consider instance I'', obtained from $\tilde{I}$ by giving sufficiently high weight w''_{b_3} to b_3. By *simplicity*, B's flights would be allocated the same set of slots as in $\tilde{I}$, and in particular B would not be assigned slot 3 at I''. Since this implies $\varphi_{b_3}(I'') > 3$, $\varphi(I'')$ would violate *individual rationality* for B. Therefore $\varphi_{b_3}(\tilde{I}) = 3$.

But then at worst $\varphi_B(\tilde{I}) = \{2,3,5\}$, i.e. $\varphi_{b_5}(I) = \varphi_{b_5}(\tilde{I})$, while B gains via the other two flights. Since B would gain from such a manipulation, this rules out Case 3. In the two remaining cases, B must receive slot 1.

Case 4: $\varphi_B(I) = \{1,2,5\}$. Recall that *simplicity* then implies $\varphi_B(I') = \{1,2,5\}$, so $\varphi_{b_7}(I') = 5$. From instance I', consider airline B delaying b_5 to $e'_{b_5} = 4$ (resulting in $\tilde{I} = I'_{e_{b_5} \to e'_{b_5}}$). By *non-wastefulness*, flights c_4 and b_3 are assigned the first two slots (in some order) and $\varphi_{c_6}(\tilde{I}) = 3$. As B receives slots 4 and 5, *self-optimization* results in $\varphi_{b_7}(\tilde{I}) = 4$. This improvement for b_7 gives B a relative gain of 7 units, while B must lose strictly fewer than that through the change in the assignments of b_3 and b_5, ruling out Case 4.

Case 5: $\varphi_B(I) = \{1,2,4\}$ (so $\varphi_C(I) = \{3,5\}$).

By self-optimization, $\varphi_{c_6}(I) = 3$. From instance I, consider airline C delaying c_4 to $e'_{c_4} = 4$ (resulting in $\tilde{I} = I_{e_{c_4} \to e'_{c_4}}$). By *non-wastefulness*, $\varphi_{b_3}(\tilde{I}) = 1$, $\varphi_{b_5}(\tilde{I}) = 2$, and $\varphi_{c_6}(\tilde{I}) = 3$.

We also show $\varphi_{c_4}(\tilde{I}) = 4$. Suppose not, so $\varphi_{c_4}(\tilde{I}) = 5$. Consider instance I''' obtained from $\tilde{I}$ by giving sufficiently high weight w'''_{c_4} to c_4. By *simplicity*, C's flights would be allocated the same set of slots as in $\tilde{I}$, and in particular C would not be assigned slot 4. Since this implies $\varphi_{c_4}(I''') = 5$, $\varphi(I''')$ would violate *individual rationality* for C. Therefore $\varphi_{c_4}(\tilde{I}) = 4$.

Since $\varphi_{c_6}(\tilde{I}) = \varphi_{c_6}(I)$ and $e_{c_4} \leq \varphi_{c_4}(\tilde{I}) < \varphi_{c_4}(I)$, φ is manipulable. □

Proof of Theorem 5. The algorithm makes no use of weight information until the Self-optimizing step, by which time the set of slots to be received by any airline $A \in \mathcal{A}$ has been fully determined. Subject to the constraint that A receives precisely that set of slots, the Self-optimization step uses weights only to give A its most-preferred landing schedule. Therefore an airline has no incentive to misreport this information.

In Step 0 of the algorithm, each airline chooses its favorite set of slots from amongst those which it owns (under the implicit assumption that it may assign flights to the set in a self-optimized way). At each subsequent step, an airline either keeps its current set or selects a better one. Hence the rule is *individually rational*. It is obvious that the rule is *non-wasteful*, *simple*, and *self-optimizing*. □

Lemma 1 states that airline A never regrets rejecting a slot s at some step of the algorithm, in that if the algorithm later assigns A's flights to some set of slots T, then A could not benefit by replacing some $s' \in T$ with s.[36] This is similar to Proposition 3 in Roth (1984) ("Rejections are final").

Lemma 1. *Fix an instance I and a DASO rule φ. Suppose airline $A \in \mathcal{A}$ and slot $s \in S$ satisfy $s \in T_A^k \setminus C_A(T_A^k)$ (s was rejected by A) in some step k of the DASO algorithm. Then for all steps $\ell > k$ of the algorithm, $s \notin C_A(T_A^\ell \cup \{s\})$. In particular, for all $f \in F_A$, $e_f \leq s$ implies $\varphi_f(I) < s$.*

Proof. Denote the flights of A that could feasibly use s as $G_A = \{f \in F_A : e_f \leq s\}$. If A rejects s in step k ($s \notin C_A(T_A^k)$), then T_A^k contains at least $|G_A|$ slots strictly earlier than s (see Definition 10).

This implies that $C_A(T_A^k)$ contains $|G_A|$ slots strictly earlier than s. At every step of the algorithm $\ell \geq k$, if $C_A(T_A^\ell)$ contains at least $|G_A|$ such slots, then so do $T_A^{\ell+1}$ and $C_A(T_A^{\ell+1})$. Furthermore, at the conclusion of the algorithm, each $f \in G_A$ satisfies $\varphi_f(I) < s$. □

Proof of Theorem 6. Suppose by contradiction that there is an instance I, airline $B \in \mathcal{A}$, and slot $h \in \Phi(B)$ such that some flight is assigned an earlier

[36] An airline could regret *accepting* a slot, in that if it had not, then it would have received a better set of slots later.

arrival time when B "hides" h. Let $I' = I_{S \to S \setminus \{h\}}$ be the instance obtained from I by deleting slot h. Let Π and Π' denote the landing schedules output by the DASO rule for I and I' respectively. Finally, let Π'' be the landing schedule obtained from Π' by "unhiding h," i.e. by self-optimizing the flights of airline B over the slots $C_B(\Pi'(F_B) \cup \{h\})$.

To simplify exposition, assume that the initial owner of any slot s has the highest priority for that slot under $\gg_s$. (This is without loss of generality due to the first step of DASO, in which slots propose to their owners.)

Let $f \in F_A$ be the *earliest* flight of any airline that receives an earlier slot under Π'' than under Π, i.e. let $f \in F_A$ be such that both $\Pi''(f) < \Pi(f)$ and

$$\forall f' \in F \setminus \{f\}, \qquad \Pi''(f') < \Pi(f') \Longrightarrow \Pi''(f') > \Pi''(f). \tag{1}$$

Denote $s' = \Pi''(f)$. Since DASO rules are *non-wasteful*, s' must be occupied by some $f' \neq f$ under Π (otherwise it could be given to f); i.e. there exist $A' \in \mathcal{A}$ and $f' \in F_{A'}$ such that $\Pi(f') = s'$. Each of two cases leads to a contradiction.

Case 1: $A' = A$. Both flights $f, f' \in F_A$ can feasibly use s', but $s' = \Pi(f') < \Pi(f)$. Therefore the *Self-optimization* step of DASO (applied to I) implies that either $w_{f'} > w_f$, or both $w_{f'} = w_f$ and $\Pi^0(f') < \Pi^0(f)$ (the tie-breaking condition). In either case, since $\Pi''(f) = s'$, this implies $\Pi''(f') < \Pi''(f)$ (because Π'' is self-optimized). This contradicts Equation 1, the assumption that f was the earliest flight to move up in Π''.

Case 2: $A' \neq A$. Observe that at Π, A would strictly gain by receiving s'. That is, A does not receive s' under Π, but desires s' in that both $e_f \leq s'$ and $\Pi(f) > s'$. Lemma 1 thus implies that s' never proposed to A at any step of the DASO algorithm under I. Since A' receives s' under I (hence s' *did* propose to A'), s' puts higher priority on A' than on A: $A' \gg_{s'} A$.

If $s' = h$ then we have $A = B \gg_{s'} A'$ (since B owns s'), which is a contradiction. Thus $s' \neq h$, implying that A receives s' under Π'.

For A to receive s' when DASO is applied to I', s' must propose to A, so A' must have rejected a proposal from s' ($A' \gg_{s'} A$) at some step of the algorithm. Lemma 1 thus implies that A' cannot strictly gain at Π'' from receiving s', hence $\Pi''(f') < s' = \Pi(f')$. This contradicts Equation 1. □

Proof of Theorem 7. Fix priorities $(\gg_s)_{s \in \mathbb{N}}$ and the corresponding DASO rule. Fix an instance I, airline $B \in \mathcal{A}$, and flight $g \in F_B$, and let Π be the outcome of the DASO rule for I.

Let $I' = I_{e_g \to e'_g}$ be the instance where B misreports e_g to be e'_g, and let Π' be the outcome of the DASO rule for I'. Suppose by contradiction that for all $f \in F_B$, $e_f \leq \Pi'(f) \leq \Pi(f)$, with $\Pi'(f) < \Pi(f)$ for at least one $f \in F_B$.

At least one slot is assigned to different airlines under Π and Π'; denote by $s \in S$ the earliest such slot. Namely, $s = \Pi(f) = \Pi'(f')$ for some $f \in F_A$ and $f' \in F_{A'}$, where $A \neq A'$, and s is the earliest such slot.

Consider the number of slots A obtains at time slot s or earlier. By our choice of s, A receives fewer such slots at Π' than at Π. Therefore, $B \neq A$.

Case 1: $A' \neq B$. Note that since $A \neq B \neq A'$, both A and A' report the same arrival times at both I and I'.

Recall that under Π and Π', the airlines receive the same sets of slots strictly earlier than s. Since A receives s under Π, A could be made better off at Π' by being offered s. Therefore (Lemma 1) s never proposed to A in the DASO algorithm under I', hence $A' \gg_s A$.

Symmetrically, since A' receives s under Π', A' could be made better off at Π by being offered s. By Lemma 1 $A \gg_s A'$, which is a contradiction.

Case 2: $A' = B$. As above, $B \neq A$ implies $A' \gg_s A$ (so $B \gg_s A$).

Since A receives s at Π, $A' = B$ must have rejected s at some stage of the DASO algorithm applied to I. Lemma 1 thus implies that B cannot gain by receiving s at Π: all of B's flights that can feasibly use s must be assigned to even earlier slots under Π. Since B receives the same number of such earlier slots at Π', B cannot feasibly use s according to the arrival times e reported at I. Therefore, B must have reported an infeasible e'_g causing it to receive a slot s that it cannot feasibly use under its true arrival times. This contradicts the assumption that airline $B = A'$ gains from such a misreport. $\square$

B Online Appendix for *Incentives in Landing Slot Problems* (not for publication)

B.1 Airline preferences: substitutable, not responsive

Preferences in our paper are defined only over sets of a fixed cardinality. However, we show that we cannot imbed such preferences into "responsive preferences" over sets of any size, as defined in the college admissions literature.

Definition: A relation P defined over *all* subsets of slots is **responsive** when, for each $s, s' \in S$,

- for each $S' \subseteq S \setminus \{s\}$, we have $S' \cup \{s\} \succeq_A^w S'$ if and only if $s \; P \; \emptyset$; and
- for each $S'' \subseteq S \setminus \{s, s'\}$, we have $S'' \cup \{s\} \succeq_A^w S'' \cup s'$ if and only if $s \; P \; s'$.

The following example shows that some weight-based preferences over subsets of size $|F_A|$ are not consistent with any responsive relation over all subsets of S.

Example (Preferences of airlines are not responsive): Consider $S = \{1, 2, 3, 4, 5\}$, and let airline A have $F_A = \{f, f', f''\}$ with $e_f = 1$, $e_{f'} = 2$ and $e_{f''} = 3$, and with $w_f = 1.5$, $w_{f'} = 1$ and $w_{f''} = 8$. This induces the following preference ordering $\succ_A^w$ over subsets of size 3.

$\succ_A^w$
$1, 2, 3$
$1, 3, 4$
$\mathbf{1, 3, 5}$
$2, 3, 4$
$2, 3, 5$
$\mathbf{3, 4, 5}$
$1, 2, 4$
$1, 4, 5$
$\mathbf{2, 4, 5}$
$\mathbf{1, 2, 5}$

Let P be a preference relation over all subsets of S that coincides with $\succ_A^w$ on the above subsets. If P is responsive, then $\{1, 3, 5\} \succ_A^w \{3, 4, 5\}$ would imply $\{1\} \; P \; \{4\}$ (by letting $S'' = \{3, 5\}$ in the definition of responsiveness).

Similarly, $\{2,4,5\} \succ_A^w \{1,2,5\}$ would imply $\{4\}$ P $\{1\}$ (by letting $S'' = \{2,5\}$). Since these conclusions are contradictory, P cannot be responsive.

Denote A's flights that can feasibly use s as

$$F_A^s \equiv \{f \in F_A : e_f \leq s\}.$$

For each airline A and each set of slots $T \subseteq S$, we say that T **is feasible for** A if there exists a (feasible) landing schedule Π such that $\bigcup_{f \in F_A} \Pi(f) \subseteq T$.[37]

The following requirement reflects the notion that if a slot is chosen from a large set $T' \subseteq S$, then it should still be chosen from within subsets of T' that contain it.

Definition 14. (e.g. see Roth (1984)) Preferences of an airline A, yielding choice function $C_A()$, satisfy **substitutability** when for each $T \subset T' \subseteq S$, with T feasible for A, we have $[T \cap C_A(T')] \subseteq C_A(T)$.

The following result holds not only on our domain of linear-weight preferences, but would hold on any airline preference domain in which "earlier is better," i.e. any domain in which an airline prefers to feasibly move one of its flights earlier, with no further restriction on preferences.

Proposition 2. *Preferences of airlines satisfy substitutability.*

Proof. Let $A \in \mathcal{A}$ and let $T \subset T' \subseteq S$ where T is feasible for A. Suppose that $s \in T \setminus C_A(T)$. We show $s \notin C_A(T')$ concluding the proof.

Since $s \notin C_A(T)$, the flights F_A^s (defined above) all can be assigned to slots within T that are earlier than s. This implies that $F_A^s = F_A^{s-1}$ and $|\{\bar{s} \in T : \bar{s} < s\}| \geq |F_A^s| = |F_A^{s-1}|$.

Since $T \subset T'$ these inequalities imply $|\{\bar{s} \in T' : \bar{s} < s\}| \geq |F_A^s| = |F_A^{s-1}|$. That is, the flights F_A^s can be assigned to slots within T' that are earlier than s. Therefore $s \notin C_A(T')$. $\square$

The only property assumed on choice functions $C_A()$ are that, if $C_A(T)$ does not contain some $s \in T$, then it must contain enough earlier slots to feasibly hold all of A's flights that could have used s. This property would hold on any preference domain in which "earlier is better."

[37]Note that this implies $|T| \geq |F_A|$.

B.2 Slot-propose and Airline-propose Deferred Acceptance coincide

On our domain of problems, both our *slot-proposing* and an *airline-proposing* version (below) of Deferred Acceptance yield the same outcome. In other words, the slot-optimal and airline-optimal stable matches coincide on our domain of landing slot problems. This equivalence is straightforward in standard models whenever one side of the market has a common preference ranking of the agents on the other side of the market. While this common ranking does not hold in our model (due to the e_f's), there is "enough" commonality in their rankings for the result to hold. Indeed, any airline that can utilize slot 1 agrees that it is, in a sense, a "best" slot (though not necessarily "the" best slot since a highly weighted flight f with $e_f > 1$ cannot use it). Therefore, stability requires slot 1 to go to its highest ranked airline that can feasibly use it. *Conditional on this*, a similar argument requires slot 2 to go to its highest-ranked airline that can feasibly use it, and so on.

Formalizing this requires us to define an airline-proposing version of Deferred Acceptance that respects the initial landing schedule in the same way that DASO rules do in Step 0. Effectively, Step 0 is equivalent to modifying the priority orders $\gg$ so that each slot ranks its owner (under the initial landing schedule) highest. Indeed DASO rules could equivalently be defined this way. Here we define A-DASO rules using this convention. The algorithm is basically three parts: modifying the priorities, classic Deferred Acceptance, and self-optimization as in DASO.

Definition 15. For any profile of priorities $(\gg_s)$ on $\mathcal{A}$, the **A-DASO rule with respect to** $\gg$ associates with every instance I the landing schedule computed with the following "A-DASO algorithm."

> **Step 0:** (Owner has top priority.) For each slot s, let $\gg'_s$ be the priority order over airlines that satisfies (i) $s \in \Phi^0(A)$ implies that A is ranked first in $\gg'_s$, (ii) $s \notin \Phi^0(B) \cup \Phi^0(C)$ implies $[B \gg'_s C \Leftrightarrow B \gg_s C]$.
>
> **Step $k = 1$:** Each airline proposes to its favorite set of slots. Each slot s tentatively accepts the offer of its highest ranked proposer under $\gg'_s$, and rejects the other proposing airlines.

Step $k = 2, \ldots$: If there were no rejections in the previous round, proceed to the Self-optimization step. Otherwise, *each* airline A proposes to its favorite set of slots from among those slots that have not already rejected A. (Note that by substitutability, A will re-propose to all of the slots that *accepted* its offer in the previous round.) Each slot s tentatively accepts the offer of its highest ranked proposer under $\gg'_s$, and rejects the other proposing airlines.

Self-optimization step: For each airline A, assign A's flights to the slots who accepted its proposal in the previous step so that the resulting landing schedule is self-optimized. Break ties among equally-weighted flights by preserving their relative order in Π^0.[38]

Theorem 8. *For any priorities $\gg$ and any instance I, the outcomes of the DASO rule $\varphi^{\gg}(I)$ and the A-DASO rule associated with $\gg$ coincide.*

Proof. Fix priorities $\gg$, and suppose by contradiction that there is I such that $\Pi \equiv \varphi^{\gg}(I) \neq \varphi^{A-DASO,\gg}(I) \equiv \Pi'$. Let s be the earliest slot for which the rules differ: $s = \Pi(f)$ implies $\Pi(f) \neq \Pi'(f)$, and $\Pi(f) < s$ implies $\Pi(f) = \Pi'(f)$.

Let $\mathcal{A}_s$ be the set of airlines A that can both (i) feasibly assign some flight $f \in F_A$ to s and (ii) assign other flights in F_A to each slot $t < s$ that A receives under Π. It is obvious by feasibility that both DASO and A-DASO must assign to s a flight from an airline in $\mathcal{A}_s$. By Lemma 1, DASO gives s to the highest ranked airline in $\mathcal{A}_s$ under $\gg$.

Denote this highest-ranked airline as A and suppose A-DASO yields the set of slots $\Pi'(A)$ to A. By definition, it is clear that $s \in C_A(T \ \cup \{s\})$, i.e. A would choose to take s in exchange for some other slot assigned by A-DASO. But this means that under an airline-proposing version of DA, A would propose first to s before ultimately proposing to one of the other slots in $t > s$ that it ends up receiving. This means that s rejects A for one of the other flights in $\mathcal{A}_s$, contradicting the fact that A is highest-ranked in $\gg'$ among $\mathcal{A}_s$. □

While this equivalence can be intuitively attributed to the commonality of airline preferences described above, one should note that airlines do *not* have

[38] This tie-breaking is irrelevant as in DASO.

common preferences over *sets* of slots. Consider an airline with two flights f and g, evaluating two (feasible) sets of slots: $X = \{1, 3\}$ and $Y = \{2, 4\}$. Depending on the flights' parameters, it is obvious the airline could prefer X to Y (e.g. whenever $e_f = e_g = 1$). But it also could prefer Y to X, e.g. when $e_f = 1$, $e_g = 2$, and w_g/w_f is sufficiently large.

B.3 Alternate Algorithm

The proof of Theorem 8 suggests another algorithmic description of DASO rules, exploiting the additional structure that our model adds to the classic college admissions model.[39] With its "greedy" structure, this algorithm may yield a more efficient implementation of DASO rules in practice. To describe it concisely, assume that the initial owner of any slot s is ranked highest in $\gg_s$ and that $S = \mathbb{N}$.

> **Step 1:** Temporarily assign slot 1 to a (feasible) flight $f \in F_A$ such that A is the highest-ranked airline in $\gg_1$ that can feasibly use slot 1. Remove f from the list of flights. (If no such flight exists, slot 1 remains vacant.)
>
> **Step 2:** Temporarily assign slot 2 to a (feasible) flight $g \in F_B$ such that, *subject to the removal of* f, B is the highest-ranked airline in $\gg_2$ that can feasibly use slot 2. Remove g from the list of flights. (If no such flight exists, slot 2 remains vacant.)
>
> **Step k:** Continue similarly with slots $3, 4, \ldots$, until all flights are temporarily assigned.
>
> **Final Step:** Self-optimize the temporary landing schedule to achieve the final schedule.

We leave it to the reader to verify that such an algorithm yields the same outcome as Definition 11.

B.4 Endogenous flight cancelations

Observation 2 from Subsection 5.3 is more formally stated as follows.

[39] We thank Utku Ünver for pointing this out.

Observation 3. Fix an instance I, and let instance $I' = I_{F \to F \setminus \{f\}}$ to be the instance obtained by deleting f from I. Fix a DASO rule (priorities $\gg$), and let Π and Π' be the landing schedules output by the rule for I and I', respectively. Then $\forall g \in F \setminus \{f\}$ we have $\Pi'(g) \leq \Pi(g)$.

As we discuss below, this proof is essentially the same as the proof of Konishi and Ünver's (2006) (logically unrelated) Capacity Lemma.

Proof. Fix an instance I and a DASO rule with priorities $\gg$.

Step 1 uses the known idea of transforming a college admissions market to a marriage market (e.g. see Roth and Sotomayor's book) by giving the student side of the market preferences over individual college "seats." Rather than breaking up a college (airline) into *arbitrary* seats, however, we order the flights by weight, which turns out to handle the *self-optimization* step of DASO. Formally, give each flight f preferences over slots, so that e_f is preferred to e_f+1 is preferred to e_f+2, etc. Give each slot s strict preferences over *individual flights*, constructed from the priority ordering $\gg_S$ as follows: for all airlines A, B and all $f, g \in F_A$ and $h \in F_B$, (i) $A \gg_s B$ implies f is preferred to h, and (ii) $w_f > w_g$ implies f is preferred to g (break ties according to flights' relative order in Π^0, as in the DASO algorithm).

Step 2 is to observe that a standard slot-proposing DA applied to this marriage market yields the DASO rule's outcome for I. This is straightforward to show, e.g. using the idea of the Alternate Algorithm we discuss in Subsection B.3. Specifically, the highest-weight flight of the highest-priority airline in $\gg_1$ will be the first flight to get (and keep!) a DA-proposal from slot 1. Given this, the highest-weight flight of the highest-priority airline in $\gg_2$ *other than the previously assigned flight* will ultimately receive (and keep) a proposal from slot 2. Continuing the argument shows that the outcome coincides with the DASO rule.

Step 3 is to apply the well known Gale-Sotomayor result that the removal of a man weakly benefits all other men under deferred acceptance in marriage markets. Hence all other flights gain in this artificial marriage market when flight f departs, meaning they receive earlier slots in the DASO outcome. $\square$

The idea of deleting a flight is reminiscent of *capacity manipulation* in the literature on college admissions problems. Consider the **Capacity Lemma** of Konishi and Ünver, stating (under responsive preferences) that when a college reduces its capacity, all other colleges benefit under DA. Indeed, the deletion of a flight f reduces airline A's demand for slots by one unit, which is effectively a capacity reduction.

There is a subtle difference here, however, in that when a flight is deleted, the airlines preferences also change *as a function of which flight is deleted*. For example, consider an airline with flights f, g, h such that

$$\begin{aligned} e_f &= 1 \quad w_f = 1 \\ e_g &= 1 \quad w_g = 1 \\ e_h &= 3 \quad w_h = 3 \end{aligned}$$

Suppose the airline deletes a flight, and ask what its resulting preferences are over, say, the two sets of slots $\{4, 5\}$ and $\{3, 6\}$. If the airline had deleted h, it would be indifferent among these two sets. On the other hand, if the airline deletes f (or identically, g), then it would have a strict preference for $\{3, 6\}$, where it is improving flight h (3 weight units) at the cost of flight g (1 unit). In contrast, the idea of capacity manipulation (Sönmez, and Konishi-Ünver) is to cap the number of students with whom a college can match, which *does not change* the underlying preference that the college initially had for sets of students strictly smaller than its true capacity.

B.5 Weak Incentives

Schummer and Vohra (2013) show that two simple rules—the FAA's Compression algorithm and the TC rule—satisfy weak non-manipulability via arrival times. Since their paper considers only simple rules and weak incentives, they need not model the part of airline preferences represented here by weights w_f. Consequently they need not consider whether any landing schedule is self-optimized (since this is irrelevant when speaking of weak incentives). Here we show that their incentive results are robust if we assume that the airlines (or the rule) first self-optimize the initial landing schedule.

Proposition 3. *Consider the rule that first self-optimizes the initial landing schedule and then applies the Compression algorithm. This rule is weakly non-manipulable via earliest arrival times.*

The same conclusion holds for the rule that applies the TC rule of Schummer and Vohra (2013) to a self-optimized initial schedule.

Proof. Let φ denote the rule that first self-optimizes the initial landing schedule and then applies the Compression algorithm. Fix an instance I, airline A, and flight $f \in F_A$. Suppose A misreports e_f to be $e'_f \neq e_f$. Let $I' = I_{e_f \to e'_f}$. Denote $\Pi = \varphi(I)$ and $\Pi' = \varphi(I')$.

Let Π_1 be the landing schedule that results from self-optimizing the initial landing schedule Π^0 using the parameters in I. Let Π'_1 be the landing schedule that results from self-optimizing Π^0 using the parameters in I'.

Suppose $\Pi_1 = \Pi'_1$, i.e. that A's misreport has no effect on the self-optimization of Π^0. Then the Compression algorithm is applied to two (optimized) instances that differ only in e_f (and not in initial schedules). The result of Schummer and Vohra (2013) thus implies the result (since they take an arbitrary initial landing schedule as fixed and allow for arbitrary misreports).

If $\Pi_1 \neq \Pi'_1$ then the change of e_f to e'_f affects the self-optimization exercise, so it must be that $\Pi_1(f) \neq \Pi'_1(f)$. We show that f ends up either with an infeasible slot or a later slot than it would without a misreport.

Case 1: $e'_f < e_f$.

Since Π'_1 is self-optimal for I' but not for I, it must be infeasible for I, i.e. $\Pi'_1(f) < e_f$. Since Compression never moves a flight to a later slot, $\Pi'(f) \leq \Pi'_1(f) < e_f$, i.e. f receives an infeasible slot. Therefore A does not benefit from this manipulation.

Case 2: $e_f < e'_f$.

Since Π_1 is self-optimal for I but not for I', it must be infeasible for I', i.e. $\Pi_1(f) < e'_f \leq \Pi'(f)$. Since Compression moves no flight to a later slot, $\Pi(f) \leq \Pi_1(f) < \Pi'(f)$, i.e. f gets a strictly later slot after the misreport.

In both cases, the misreport cannot improve the outcome of each of A's individual flights.

The proof is identical for TC. □

More generally, any rule that is weakly non-manipulable by arrival times remains so if the rule is augmented by first self-optimizing the initial landing schedule.

On the other hand when a rule *does not* self-optimize the initial schedule, but performs self-optimization only after the rule operates, it may be strongly manipulable. Example 2 illustrates this for the Compression rule. The same manipulation illustrated in that example would benefit A if we apply a self-optimization step only after using the TC rule of Schummer and Vohra (2013). That rule prescribes the same outcome for that example as Compression does. However, the manipulation by A would assign flights a_4 and a_2 to slots 1 and 2 respectively. That is, A again has a strong manipulation under the *TC-then-self-optimize* rule.

B.6 No Pareto-dominance

The following result implies Proposition 1. We are grateful to a referee for suggesting a non-Pareto-dominance result, leading to this theorem.

A rule φ *Pareto-dominates* a rule $\varphi' \neq \varphi$ if at every instance, every airline weakly prefers its outcome under φ to its outcome under φ', with a strict preference for some airline at some instance.

Theorem 9. *No* FAA-conforming *rule is Pareto-dominated by a* simple *rule.*

Proof. Suppose a *simple* rule φ' Pareto-dominates an *FAA-conforming* rule φ. Note that it is without loss of generality to assume that φ' is also *self-optimized*, since otherwise the rule φ'' that is the "self-optimization of φ'" Pareto-improves φ' and hence also Pareto-dominates φ. We also assume that when φ and φ' self-optimize flights, ties are broken in the same way by both rules (e.g. if two equal-weight flights can use the same two slots, the rules preserve the relative order those two flights had in the initial landing schedule). This is also without loss of generality since swapping equal-weight flights is a Pareto-indifferent operation.

Let I be an instance where at least one airline strictly prefers $\varphi'(I)$ to $\varphi(I)$. Let s be the earliest slot to which the rules make different assignments. Since the rules coincide on slots earlier than s, and since φ is *non-wasteful*, if φ leaves slot s vacant, then so must φ' by feasibility. Therefore φ assigns some flight f of some airline A to slot s. By our choice of s, φ' must assign f to a slot later than s. Furthermore, since φ' is *self-optimized* (and by our tie-breaking assumption), φ' does not assign another of A's flights to s.

Denote those of A's flights that φ assigns to slot s or earlier by $F' = \{g \in F_A : \varphi_g(I) \leq s\}$. By our choice of s, for each flight $g \in F' \setminus \{f\}$ we have $\varphi_g(I) = \varphi'_g(I)$. Consider a new weight profile w^λ in which we scale up the weights of flights in F' by a factor of $\lambda > 1$, and leave all other flights' weights unchanged. By *simplicity* and *self-optimization*, φ and φ' continue to assign flights $F' \setminus \{f\}$ to exactly the same slots as before. For the same reason, φ continues to assign f to s, and φ' assigns f to some slot strictly later than s. For sufficiently large λ, A would strictly prefer $\varphi(I_{w \to w^\lambda})$ to $\varphi'(I_{w \to w^\lambda})$, regardless of how φ and φ' assign A's remaining (low-weight) flights $F_A \setminus F'$. Therefore φ' does not Pareto-dominate φ. $\square$

The proof technically shows a stronger fact than the Theorem states: If the outcome of a simple rule φ' differs from that of an FAA-conforming rule φ at any instance I, then there exists an airline A and another weight profile w' such that, at instance $I_{w \to w'}$, A strictly prefers the outcome under φ. Hence this non-Pareto-comparability holds even on every (small) subdomain in which we fix *all* parameters other than weights. To the extent that real world airline preferences (e.g. weights) are private information, this yields a fairly strong non-comparability result from the perspective of an uninformed planner.

B.7 Summary of properties for three simple rules

Non-manipulable by...	Compression	TC	DASO
weak flight delay	Yes*	Yes*	Yes
flight delay	*no*	*no*	*no*
weak slot destruction	*no*	*no*	Yes
slot destruction	*no*	*no*	Yes
postpone flight cancelation	*no*	*no*	Yes
selects from a weak core (S-V 2013)	*no*	Yes	*no*

Yes*: Yes *except* when self-optimization is performed only *after* the rule operates.