

Online Appendix for *Incentives in Landing Slot Problems*

James Schummer* Azar Abizada†

April 14, 2017

This document contains supplementary results and proofs for *Incentives in Landing Slot Problems*, published in *The Journal of Economic Theory*.

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 14. 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 3 (Preferences of airlines are not responsive). Consider slots $S = \{1, 2, 3, 4, 5\}$, 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.

*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.

\succ_A^w
1, 2, 3
1, 3, 4
1, 3, 5
2, 3, 4
2, 3, 5
3, 4, 5
1, 2, 4
1, 4, 5
2, 4, 5
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 $\cup_{f \in F_A} \Pi(f) \subseteq T$.¹

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 15. 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.*

¹Note that this implies $|T| \geq |F_A|$.

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.”

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 16. 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 $k = 2, \dots$: 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 .²

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

²This tie-breaking is irrelevant as in DASO.

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 . \square

While this equivalence can be intuitively attributed to the commonality of airline preferences described above, one should note that airlines do *not* have 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.³ 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

³We thank Utku Ünver for pointing this out.

airline in \ggg_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, \dots , 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.

Observation 3. Fix an instance I , and let instance $I' = I_{F \rightarrow F \setminus \{f\}}$ to be the instance obtained by deleting f from I . Fix a DASO rule (priorities \ggg), 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 \ggg .

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 \ggg_s as follows: for all airlines A, B and all $f, g \in F_A$ and $h \in F_B$, (i) $A \ggg_s B$ 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

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

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 $I' = I_{e_f \rightarrow e'_f}$. 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 (?) 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 $\varphi(I_{w \rightarrow w^\lambda})$ to $\varphi'(I_{w \rightarrow w^\lambda})$, 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 $I_{w \rightarrow w'}$, 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	<i>no</i>	<i>no</i>	<i>no</i>
weak slot destruction	<i>no</i>	<i>no</i>	Yes
slot destruction	<i>no</i>	<i>no</i>	Yes
postpone flight cancelation	<i>no</i>	<i>no</i>	Yes
selects from a weak core (S-V 2013)	<i>no</i>	Yes	<i>no</i>

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