Overbooking *

Jeffrey C. Ely † Daniel F. Garrett ‡ Toomas Hinnosaar §

April 28, 2016

Abstract

We consider optimal pricing policies for airlines when passengers are uncertain at the time of ticketing of their eventual willingness to pay for air travel. Auctions at the time of departure efficiently allocate space and a profit maximizing airline can capitalize on these gains by overbooking flights and repurchasing excess tickets from those passengers whose realized value is low. Nevertheless profit maximization entails distortions away from the efficient allocation. Under regularity conditions, we show that the optimal mechanism can be implemented by a modified double auction. In order to encourage early booking, passengers who purchase late are disadvantaged. In order to capture the information rents of passengers with high expected values, ticket repurchases at the time of departure are at a subsidized price, sometimes leading to unused capacity.

JEL: D42, D44, D82

Keywords: airlines, overbooking, revenue management, dynamic mechanism design

*We would like to thank Alessandro Pavan, Maher Said, and Jean Tirole for helpful comments. Thanks also to seminar participants at Aalto University, Bocconi University, Collegio Carlo Alberto, Columbia University, the CSIO/IDEI Joint Workshop on Industrial Organization, Microsoft Research, the National University of Singapore, Northwestern University, the University of Auckland, the University of Melbourne, the University of New South Wales, the University of Utah, the Midwest Economics Theory meeting at Washington University, IIOC 2013, and the Conference on Economic Design in 2015. Jeff Ely would also like to acknowledge the hospitality of Microsoft Research New England where some of this work was carried out.

†Northwestern University, jeffely@northwestern.edu.

‡Toulouse School of Economics, dfgarrett@gmail.com.

§Collegio Carlo Alberto, toomas@hinnosaar.net.

1 Introduction

Overselling limited seating space is standard practice among airlines. According to U.S. Department of Transportation reports, between April and June 2011, about one per cent of passengers ticketed on a U.S. domestic flight were denied boarding on that flight. To mitigate the potential inconvenience to passengers, airlines typically look for volunteers willing to transfer to later flights, raising the level of compensation offered until enough passengers willing to delay their travel have been found. As a result, only about one per cent of passengers who are denied boarding on oversold flights are bumped involuntarily.[1]

Airlines typically explain their practices as being motivated by the fact that a certain number of passengers can be expected not to show up (on time) for a flight and overbooking capitalizes on slack capacity, improving efficiency. From this perspective, the extent to which an airline should overbook is chiefly a statistical question, one which requires airlines to balance the value of unused capacity against the costs associated with overbooked flights, including inconvenience to passengers and possible compensation. This is also the view expressed in the extensive body of literature in operations research (early contributions include Beckmann (1958) and Rothstein (1971)).[2]

In this paper, we explore a different but complementary rationale for overbooking. This rationale is based on price discrimination among passengers who face uncertainty about the eventual value they will place on being seated on the flight. Unlike the operations-based perspective on overbooking, passenger incentives—both to purchase tickets and to give up their seats for compensation—are central to our theory. We focus on a setting where ticket sales are not needed to achieve an efficient allocation, but where they instead play a role in extracting surplus from passengers who are uncertain about their values for flying in advance of the flight. Of course, one expects that, in practice, the traditional rationale for overbooking also plays a role. Our modeling choice is motivated by simplicity and a desire to understand which properties of pricing policies used by airlines arise from the surplus-extraction motive.

We consider an airline which can offer homogeneous tickets at a single price, akin to offering only a single fare class. We show that the airline may profit from selling more tickets than the available capacity. Rather than indiscriminately canceling tickets when a flight is overbooked, i.e., resorting to involuntary bumping of passengers, the

[1]Attempts to elicit volunteers are becoming more sophisticated. For instance, Delta Airlines has recently implemented a system in which passengers indicate, at the time of check-in, their willingness to take a later flight. They do so by specifying the minimum value of the travel voucher they would be willing to receive to delay their travel.

[2]There are a few recent exceptions. For example in Gallego and Sahin (2010) the airline sells options to fly and passengers whose values turn out to be low take the refund rather than fly. However, in contrast to our assumptions, the refund in their model is independent of the demand.

airline always finds it profitable to treat the purchase of a ticket as conferring the right to a seat. The airline must then repurchase these rights. Although, in practice, airlines withhold their authority to bump even unwilling passengers, subject to rules surrounding compensation,[3] our finding does seem consistent with the status quo in the U.S., where involuntary bumping occurs only rarely.

Procuring oversold capacity via auction offers airlines the chance to improve the ex-post efficiency of seating allocations. The choice of reallocation mechanism also affects the demand for tickets. Compensation that passengers anticipate in the event of overbooking improves the value of obtaining a ticket in advance of the flight. Therefore, the design of mechanisms for determining the final seat allocation cannot be separated from the choice of ticket pricing policies. We shed light on this connection.

In our theory, the passengers who purchase tickets in advance of the flight will be those who have sufficiently favorable information about their value for flying. Among ticket purchasers, those with the least favorable information will be at the margin. Such passengers are, of course, more likely to benefit from compensation offered on the date of travel. The possibility of future compensation allows the airline to raise the price of the ticket without affecting the payoffs of the marginal ticket holders and to do so without reducing the number of tickets sold. In so doing, airlines reduce the consumer surplus of the infra-marginal ticket holders. This is simply because infra-marginal ticket holders pay the higher ticket price but are less likely to benefit from the compensation for giving up their seat. In this way, the promise of future compensation is a tool for the airline to limit consumer surplus.

While the above observations seem consistent with practice, our theory also suggests possible ways that airlines could tailor seat reallocation to improve profits. For example, the airline may profit from repurchasing seating rights even when capacity constraints do not bind. While many airlines do agree to repurchase seating rights through partial refunds, the profitability of such policies naturally depends on the distribution of passenger values (and there may well be other reasons airlines would find encouraging empty flights undesirable).

Airlines can also improve profits by reallocating seating rights to passengers who do not purchase tickets in advance, either because they anticipate a low value for flying or because they are "out of the market". A passenger being out of the market simply captures the possibility that their need for travel has yet to arise. We show that airlines profit by giving such passengers the opportunity to fly, but that doing so improves the

[3]The U.S. Department of Transportation sets out the rules for compensation in the case of involuntary bumping in 'Fly-Rights: A Consumer Guide to Air Travel'. Compensation for involuntary bumping is mandatory, and may be as high as US$1,300.

value of the option associated with *not* purchasing a ticket. It therefore reduces the demand for tickets. This consideration encourages the airline to charge higher prices to unticketed passengers, an observation which seems to help explain the substantially higher ticket prices often observed in the two or three weeks immediately preceding a flight (see McAfee and te Velde (2006) and Lazarev (2012) for these pricing patterns).

We are able to quantify an airline's incentive to allocate rationed space between ticket holders and non-ticket holders using a version of the familiar virtual surplus measure. These formulas show how allocation decisions not only affect revenues directly through departure-time transfers, but also affect revenues indirectly through their effect on passengers' willingness to pay for tickets in advance. Under appropriate regularity conditions, we show that the optimal mechanism can be implemented using a modified double auction.

Finally, our analysis offers a way to evaluate the efficiency and welfare consequences of overbooking and seat reallocation. This kind of analysis may be important for policy makers who are considering imposing possible restrictions on overbooking.[4] We argue that overbooking and subsequent reallocation is potentially efficiency enhancing (although we provide no general results about the direction of the effect).

The rest of the paper unfolds as follows. After discussing related literature in the rest of this section, we introduce a model in Section 2. Section 3 introduces the mechanisms at work when a single kind of ticket is sold before the date of travel as described above. We then introduce a stylized example to illustrate the key trade-offs. In Section 4 we consider the general model and derive the optimal transfers and allocation rules for these pricing mechanisms. We then show how to implement the optimal mechanism through a modified double auction. Section 5 discusses the possibility of multiple fare classes and Section 6 concludes the paper. There are two Appendices: Appendix A provides proofs of all results while Appendix B explains an approach to analyzing multiple fare classes and deriving the optimal mechanism without any restrictions.

Related literature

Julian Simon first suggested auction mechanisms as a response to overbooking in the airline industry in the mid-1960s (see Simon (1994)). These suggestions were implemented by the industry in the U.S. starting in 1978, with volunteers being selected on the basis of their willingness to give up their seats, rather than arbitrarily.[5] In another notable

[4]In the U.S., overbooking has long been accepted practice, but policy is still evolving in other jurisdictions. In the Philippines, the regulator recently proposed a ban on overbooking, although policy has now evolved towards permitting and regulating it ("Passenger bill of rights out soon", www.philstar.com, October 27, 2012).

[5]Simon (1994) notes that prior to 1978, companies such as United Airlines followed a practice of bumping "old people and armed services personnel, on the assumption that they would be least likely to

contribution, Vickrey (1972) proposed the use of an efficient auction to resolve the problem of allocation of seats in the event of overbooking. Vickrey proposed extending this mechanism to include dynamic flexible pricing schemes, and speculated about many of the issues that we address here, although without a formal model.

Our paper also contributes to more recent literature on selling to consumers who learn about their valuations over time. An important question we are concerned with is the following: When should the firm allocate available capacity—before or after buyers have learned their values? Biehl (2001), Deb (2014), Dana (1998), Möller and Watanabe (2010), and Nocke, Peitz, and Rosar (2011) provide models in which sellers choose to discriminate between consumers based on their prior information, for instance through offers of "advanced purchase discounts". DeGraba (1995) and Courty (2003) study related models in which there is no role for intertemporal discrimination due to buyers initially lacking private information. The key distinction relative to our work is that capacity constraints in these papers are either respected (e.g., DeGraba (1995) and Möller and Watanabe (2010))—i.e. tickets sold do not exceed the available capacity—or completely absent (e.g., Biehl (2001), Courty (2003), Deb (2014), and Nocke, Peitz, and Rosar (2011)). We instead demonstrate how a seller (airline) may find it optimal to over-allocate capacity and then use an appropriate mechanism for ensuring that capacity constraints are respected on the date of consumption.

There is also a recent branch of literature in revenue management and marketing that considers buyers with evolving valuations and shows that advance selling with partial refunds (or selling options to fly) ensures higher revenue than selling non-refundable tickets (Gallego and Sahin, 2010; Gallego, Kou, and Phillips, 2008) and that discounts for advance selling may increase revenues (Shugan and Xie, 2000; Xie and Shugan, 2001; Gallego and Sahin, 2010; Gallego, Kou, and Phillips, 2008). However, while these papers include more institutional details than ours, they assume that refunds for cancellations are fixed, which seriously limits the choices available to airlines. In particular, in our model, the refund as well as whether the passenger will fly or not will be determined as a comparison with other passenger values at a re-allocation auction before the departure.

In work independent of our own, Fu, Gautier, and Watanabe (2012) also study a setting in which an airline can over-allocate capacity, and then either repurchase it or randomly select passengers not to fly. Unlike our work, the airline can only allocate seats to passengers who buy tickets in advance of the flight. The overselling of tickets therefore improves efficiency by helping the airline match the highest value passengers with the available capacity. In contrast, the airline in our model can achieve full efficiency without advance ticket sales.

complain."

Restricting the available fare options to a single fare class means the airline loses some of its ability to screen initial passenger information. To screen passengers, the airline may strictly profit from setting ticket prices that are unattractive for some passengers who, although they are in the market and able to make advance ticket purchases, believe they are likely to have low values for flying. While such passengers choose not to purchase tickets in advance, they still have the opportunity to participate at a later date. In sum, the airline balances trade-offs which are absent in most of the literature on dynamic mechanism design with agents who learn about their preferences over time (e.g., Baron and Besanko (1984), Courty and Li (2000), Battaglini (2005), Eso and Szentes (2007), and Pavan, Segal, and Toikka (2014)). This literature does not restrict the class of mechanisms and finds that the principal (say, the seller) finds it at least weakly optimal to contract with agents at the first opportunity. This is simply because of the revelation principle, since any outcome obtained by (temporarily) excluding some agents can be replicated by a mechanism which induces participation by all agents at the first opportunity (in particular, the principal can use a direct revelation mechanism in which the first report occurs on the date of the agent's arrival). Below we discuss the airline's optimal mechanism when it is not restricted in the mechanisms it can offer (see Appendix B for a full analysis). Here, contracting with all passengers on the date they arrive to the market is an optimal policy, and is often strictly optimal.[6]

We assume that the airline can fully commit to the ticket price and terms, which is a common assumption in the literature on dynamic mechanism design. An important exception is Deb and Said (2015). Unlike the rest of the literature, the seller in their model cannot commit to the contracts it offers at later dates. Exclusion of buyers is important as it dictates the composition of buyers available to sign later contracts, effectively providing a form of commitment to offer less generous contracts at later dates. Thus, as in our paper, the seller may strictly profit by not contracting with all available customers, although the reason for this finding is quite different (i.e., lack of commitment, rather than a restriction on the seller's ability to screen customers).

Selling tickets in our model also provides a way for the airline to distinguish timely arrivers from those who arrive to the market late. Our paper is thus related to the literature studying buyers who arrive over time but who face no uncertainty about their future valuations. Examples include Gershkov and Moldovanu (2009, 2012), Board and Skrzypacz (2015), Said (2012), and Hinnosaar (2015) (see Bergemann and Said (2011) for an overview, as well as McAfee and te Velde (2006) for a review focused on applications

[6]Contracting with an agent at the first opportunity indeed is often strictly optimal in the existing literature on dynamic mechanism design. An exception is Nocke, Peitz, and Rosar (2011). In this model, the seller finds it weakly optimal to delay contracting to some buyers (but, by the revelation principle, the same profits can be obtained by inducing immediate participation by all).

to the airline industry).[7] Further examples include work by Pai and Vohra (2013) and Mierendorff (2016), who provide models where buyers arrive over time and have private deadlines. Mierendorff shows that the optimal mechanism may feature a handicap auction that favors the early arriver, in order that the early arriver truthfully disclose a late deadline. In contrast to all of these papers, we study an environment where buyers are uncertain about their future values for the good. The treatment of a setting with both dynamic arrivals and valuation uncertainty builds on recent work by Garrett (2014, 2015).

2 Model

A risk-neutral monopolist airline is selling m seats on a flight that departs at date 1. There are n risk-neutral potential passengers who are ex-ante anonymous and symmetric, where $n \geq m$.

A passenger may arrive to the market at date 0 or only at date 1. Only passengers arriving at date 0 will have the opportunity to purchase tickets at date 0, as described below. Date-1 arrivals naturally represent passengers who only come to realize their desire to fly close to the date of travel. The possibility of such passengers is easy to accommodate in the analysis, but is also not crucial for our main qualitative results.

A passenger i arriving at date 0 receives partial information about his value for flying at date 1, v_i, which is captured by a signal θ_i. Vectors of values are denoted in bold font, i.e., $\mathbf{v} = (v_i)_{i=1}^n$, with $\mathbf{v}_{-i'} = (v_i)_{i \neq i'}$ for any individual passenger i'. A passenger's time of arrival and both the information about his value for flying and the value itself are determined independently of the other passengers' realizations and are his private information.

The signals θ_i are drawn from a distribution with a CDF F whose support is $\Theta = [0, 1] \cup \{\varnothing\}$.[8] The signal $\theta_i = \varnothing$ indicates that the passenger is out of the market and unavailable for ticket purchases at date 0. For notational convenience we will adopt the convention that $\varnothing < 0$. A signal $\theta_i \in [0, 1]$ indicates that the passenger enters the market at date 0 and we assume that F admits a density f over that range and includes the possibility of an atom at $\varnothing$. Abusing the notation slightly, $F(\theta_i|\mathbf{S})$ denotes the distribution of $\tilde{\theta}_i$, conditional on event $\mathbf{S}$.

Conditional on his signal θ_i, a passenger i's (non-negative) eventual willingness to pay $\tilde{v}_i$ is distributed according to the CDF $G(v_i|\theta_i)$, where $G(\cdot|\cdot)$ is continuously differentiable, and where the density is denoted $g(v_i|\theta_i)$. The support of the marginal distribution of $\tilde{v}_i$ is $[\underline{v}, \overline{v}]$. (The support of the distribution of $\tilde{v}_i$ conditional on a particular realization θ_i

[7] See Lazarev (2012) for a recent empirical study of dynamic pricing in the airline industry.

[8] Throughout, random variables are denoted using tildes.

of the signal may be a strict subset of $[\underline{v}, \overline{v}]$.) Abusing notation, we also let $G(\mathbf{v}|\mathbf{S})$ and $G(\mathbf{v}_{-i'}|\mathbf{S}_{-i'})$ denote the joint distributions of passenger values conditional on the events $(\theta_i)_{i=1}^n \in \mathbf{S} \subset \Theta^n$ and $(\Theta_i)_{i \neq i'} \in \mathbf{S}_{-i'} \subset \Theta^{n-1}$.

Signals are ordered in the sense of first-order stochastic dominance, so, for all v_i, $G(v_i|\theta_i)$ is nonincreasing in θ_i over $[0, 1]$ (at this point, we impose no additional restrictions on $G(\cdot|\varnothing)$). To simplify some of the analysis, we further impose the weak restriction that there exists $k > 0$ such that, for all x, $G(k\theta_i + x|\theta_i)$ is nondecreasing in θ_i.[9]

The airline may sell tickets at date 0 and reallocate capacity (including to unticketed passengers) at date 1. We describe the mechanisms that the airline can use below. There is no discounting. Passengers realize utility v_i if they are allocated a seat and a utility zero otherwise, less any payments made to the airline across the two periods. So a passenger whose (possibly negative) total payment to the airline is ρ earns utility $v_i - \rho$ if he is seated and $-\rho$ if he is not. Given a fixed number of available seats m and a zero cost of seating a passenger in an available seat, the airline maximizes the expected total payment by passengers.

3 Pricing Mechanisms

We consider pricing mechanisms Ω of the following form. At date 0, the airline sets a price p for tickets, and also commits to a family of re-allocation mechanisms to be used at date 1 depending on the number of tickets sold.[10] In particular, the mechanism will be used to determine which passengers will sell back their tickets in the event of overbooking, to broker the possible transfer of seating rights from the date-0 purchasers to those unticketed passengers who are willing to pay to fly, and possibly to sell additional seats to passengers who wish to purchase tickets at date 1. Throughout, we restrict attention to mechanisms that treat passengers symmetrically.

Feasibility and incentive constraints. The mechanism Ω must satisfy familiar feasibility and incentive constraints. First, at most m passengers can be seated on the plane. Second, the mechanism must be implementable in perfect Bayesian equilibrium: whether or not a passenger purchases a ticket, his continuation play at date 1 must be optimal given updated beliefs. Our assumption is that, at date 1, absent any information

[9]This allows us to rely directly on arguments in Pavan, Segal, and Toikka (2014); see the proof of Lemma 1. In words, the condition says that there exists k such that, for any x, the probability that a passenger's value for flying exceeds his signal θ_i by more than $k\theta_i + x$, is nonincreasing in his signal. (We are grateful to Phil Reny for suggesting this interpretation.)

[10]We maintain the assumption that the airline can fully commit throughout the paper. See Deb and Said (2015) for an analysis of how dynamic rationing can provide partial commitment power for a seller who cannot otherwise commit to contracts offered at future dates.

provided by the airline, passengers act without knowing how many other tickets have been sold and the signals or values of the other passengers. However, it will turn out that, compared to Bayesian incentive-compatible mechanisms, there is no loss in revenue in using mechanisms that inform all passengers of the total number of tickets sold and that ask passengers to play dominant strategies at date 1 (see Lemma 2 below). At date 0, a passenger's decision whether or not to purchase a ticket must be optimal given the anticipated continuation play.

Our treatment of participation constraints is familiar from the dynamic mechanism design literature. Tickets are a contractual commitment for both the airline and the passengers, so we treat ticketed passengers as required to participate at date 1. However, we will show that, by designing ticket prices and transfers appropriately, the airline need not rely on the passengers' commitments.[11] For unticketed passengers, there is no contractual commitment (indeed, ticket sales are the *only* contract formed at date 0 and ticket prices are the only date-0 transfers) and so these passengers are not compelled to participate at date 1.

Additional constraints on the mechanism. Our analysis focuses on a restricted class of mechanisms. The first restriction, already implicit in the discussion above, is that the airline is permitted to offer tickets for a single fare class with a single price p. We assume that all passengers present at date 0 simultaneously decide whether to purchase a ticket at price p.

Intuitively, selling tickets is a way of dividing customers among those who are present at date 0 and have certain beliefs over their willingness to pay, and those who do not. It is thus equivalent to a restriction on the message space available at date 0. A buyer present at date 0 faces the decision either to communicate his desire to purchase a ticket (by sending the only available message, "purchase"), or not to communicate at all at date 0. When the set of possible customer date-0 beliefs is sufficiently rich, the airline could often profit by choosing a finer partition, requiring a richer message space. This could be achieved by offering multiple fare classes, each with a different set of terms, a possibility that we discuss below. However, our main focus is on the case where the airline offers a single fare class, and thus has no way to distinguish between the different passengers who request a ticket at date 0.

There are at least two reasons to focus on limited screening at date 0. First, it may be unrealistic to expect airlines to distinguish finely between passenger beliefs (in our model, the airline would find it profitable to offer a continuum of contracts with different

[11] This is because the airline has a degree of freedom to shift payments across time. By increasing the price of the ticket and reducing any date-1 payment (equivalently, increasing any date-1 subsidy) to ticket holders, the airline can ensure that ticket holders find date-1 participation optimal, even if they are not contractually bound.

terms; we derive the optimal mechanism with no restrictions on tickets in Appendix B).[12] Second, while it may be realistic to consider a small number of fare classes, restricting attention to only one simplifies the analysis. In particular, it allows us to focus on the decision of passengers *whether* to contract in advance of the flight, rather than the decision of *which kind* of contract to purchase.

Note that the restriction to a single fare class is by no means essential for our conclusion that the airline may enter more contracts at date 0 than it has seats. As is true in most of the work on dynamic mechanism design (as reviewed briefly above), were the airline not restricted in its choice of mechanism, it would optimally write contracts with all passengers arriving to the market at date 0. In other words, all passengers available at date 0 would participate in the mechanism at this date. Given our restriction on date-0 contracting, however, an airline may profitably use date-0 participation to screen passengers. As explained above, it can do this by selling tickets to some passengers but not others. Our interest then is in understanding how (and when) the airline optimally uses this screening device.

We assume that all passengers requesting a ticket at price p receive one. This comes at no cost to the airline: in particular, the airline could not gain by limiting the number of tickets sold and applying a randomized-rationing rule.[13] This is simply because the profits of any pricing mechanism which (randomly) rations ticket sales can be attained by an alternative mechanism in which (i) a ticket is sold to every passenger who requests one, and (ii) ticketed passengers are randomly assigned either the allocations and payments associated with holding a ticket in the original mechanism with rationing, or the allocations and payments associated with not holding a ticket in that mechanism. Below we will provide conditions under which the airline *strictly* profits by selling tickets to all buyers who request them, rather than limiting the number sold and rationing randomly.

Our second restriction is that having a ticket cannot decrease the probability of being seated. In particular, if the mechanism seats an unticketed passenger with value v_i at date 1, then it will also seat the same passenger if he has instead purchased a ticket at date 0 (holding fixed the actions of the other passengers, and assuming that the passenger in question behaves optimally at date 1). In other words, obtaining a ticket at date 0 does not hurt a buyer's probability of being seated at date 1.[14] The idea here is that tickets

[12]Indeed, part of our objective is to explain how an airline can profit from overbooking with simple instruments which resemble those already used in the industry.

[13]The randomized-rationing rule, which gives all buyers asking for a ticket an equal probability of receiving one, is common in the literature (see Tirole, 1988, pp 213-214). In our setting, randomized rationing follows from the assumptions that (i) passengers are ex-ante identical and anonymous, and (ii) passengers have no way of expressing their date-0 information other than expressing their demand for a ticket.

[14]Note that this is different to assuming that ticketed passengers are seated ahead of unticketed pas-

are required to play the role with which we are familiar, i.e., providing privileged access to their holders. Our analysis focuses on the extent to which the airline profitably uses tickets in this role.

The above conditions allow us to focus attention on a class of mechanisms which is easy to analyze. Assuming for simplicity that passengers purchase tickets whenever indifferent, we show that such mechanisms have the following property.

Lemma 1. *In any pricing mechanism satisfying the above conditions, there exists a threshold value θ^* such that passengers purchase tickets if and only if they arrive at date 0 and have a signal no less than θ^*.*

Excluding boundary cases in which all or no available passengers purchase tickets, θ^* must be a signal for which passengers are indifferent. Hence, we refer to this as the "marginal signal" and to the passenger as the "marginal ticket holder". Buyers with initial signals above θ^* anticipate higher values for flying in the sense of first-order stochastic dominance. Since obtaining a ticket increases a buyer's chances of being seated, buyers with signals above θ^* strictly prefer purchasing tickets at date 0.

If the marginal signal θ^* is less than one, then (given that buyers draw signals symmetrically) there is a positive probability that the airline sells tickets to all n passengers. If, in addition, $m < n$, then the airline practices overbooking. That is, there is a positive probability that more tickets are sold than available seats. Corollary 1 below will give sufficient conditions for θ^* to be less than one in the profit-maximizing pricing mechanism.

The date 1 mechanism. Before describing how the airline sets ticket prices, consider the mechanism it uses to allocate seats at date 1. We assume that the airline faces no restriction on date-1 mechanisms, and, without loss of generality, we consider direct revelation mechanisms. Payments and allocations in these mechanisms depend on which passengers hold tickets and their reported values for flying.

Let $s \in \{0, \dots, n\}$ be the number of passengers holding tickets purchased at date 0. We will use subscripts j to denote passengers that hold tickets at date 1 and subscripts k to denote those who do not. Thus, we can order passengers so that $j = 1, \dots, s$ hold tickets and the remaining passengers $k = s+1, \dots, n$ do not. We let $\mathbf{S}^s$ denote this event, i.e.

$$\mathbf{S}^s = \{\theta_j \geq \theta^*, j = 1, \dots, s; \theta_k < \theta^*, k = s+1, \dots n\} \tag{1}$$

and we denote by $\mathbf{S}^s_{-i}$ the corresponding event with passenger i excluded.

At date 1, each passenger observes his realized willingness to pay as well as whether or not he is holding a ticket. A direct revelation mechanism is described by a collection

sengers with the same value. We instead derive this as an implication of our model (under appropriate regularity conditions) in Proposition 2.

of functions $(\mathbf{q}, \mathbf{t})$, such that, for each $s \in \{0, \dots, n\}$, $\mathbf{q}^s(\mathbf{v}) = (q_i^s(\mathbf{v}))_{i=1}^n$ gives the probability that each passenger i is seated while $\mathbf{t}^s(\mathbf{v}) = (t_i^s(\mathbf{v}))_{i=1}^n$ gives the payment that each passenger i makes to the airline at date 1.[15] Conditional on s tickets sold, we say passengers play the "s-mechanism". To ease notation, and without loss of optimality, payments and probabilities of allocation are defined (unless otherwise specified) for all values in $[\underline{v}, \overline{v}]$, which includes values that may be inconsistent with equilibrium play.[16]

Consider a passenger i who truthfully reports his value v_i, while other passengers report $\mathbf{v}_{-i}$, and suppose the total number of tickets sold is s. Disregarding the ticket price paid at date 0, if any, the payoff to passenger i is $U_i^s(\mathbf{v}) = q_i^s(\mathbf{v})v_i - t_i^s(\mathbf{v})$.

Let $\tilde{r}$ be the random variable representing the number of tickets purchased by passengers other than i, which thus follows a binomial distribution with parameters $(n-1, 1-G(\theta^*))$. The date 1 mechanism $(\mathbf{q}, \mathbf{t})$ is Bayesian incentive compatible if, for each ticketed passenger j, all v_j and all reports $\hat{v}_j$

$$\mathbb{E}\left[U_j^{\tilde{r}+1}(v_j, \tilde{\mathbf{v}}_{-j})\right] \geq \mathbb{E}\left[q_j^{\tilde{r}+1}(\hat{v}_j, \tilde{\mathbf{v}}_{-j})v_j - t_j^{\tilde{r}+1}(\hat{v}_j, \tilde{\mathbf{v}}_{-j})\right] \quad (2)$$

and, for each unticketed passenger k, all values v_k and all reports $\hat{v}_k$,

$$\mathbb{E}\left[U_k^{\tilde{r}}(v_k, \tilde{\mathbf{v}}_{-k})\right] \geq \mathbb{E}\left[q_k^{\tilde{r}}(\hat{v}_k, \tilde{\mathbf{v}}_{-k})v_k - t_k^{\tilde{r}}(\hat{v}_k, \tilde{\mathbf{v}}_{-k})\right]. \quad (3)$$

The mechanism is *ex-post* incentive compatible if, for all s, all i, v_i, $\hat{v}_i$ and $\mathbf{v}_{-i}$,

$$U_i^s(v_i, \mathbf{v}_{-i}) \geq q_i^s(\hat{v}_i, \mathbf{v}_{-i})v_i - t_i^s(\hat{v}_i, \mathbf{v}_{-i}). \quad (4)$$

Thus in an ex-post incentive-compatible mechanism the airline can publicly announce the number of ticket holders and truth-telling would remain ex-post optimal for all passengers.

The following lemma shows that any Bayesian incentive-compatible mechanism can be replaced by an ex-post incentive-compatible mechanism that generates the same expected profit. The proof follows from essentially the same arguments made in Gershkov, Goeree, Kushnir, Moldovanu, and Shi (2012).

Lemma 2. *There is no loss of optimality in restricting attention to ex-post incentive-compatible mechanisms, i.e., mechanisms in which the airline announces the number of tickets sold and truth-telling is ex-post optimal for all passengers.*

[15]While we allow for random allocations, there is no loss of generality in considering deterministic transfers due to the linearity of airline and passenger payoffs.

[16]In particular, we require the mechanism to be defined if *one* of the bidders claims to have a value that is inconsistent with his decision whether to purchase a ticket, a possibility which arises because the support of $G(\cdot|\theta_i)$ may be a strict subset of $[\underline{v}, \overline{v}]$ for $\theta_i \in [0, 1]$.

Ticket pricing. The airline designs the date-1 mechanisms, taking into account not only the revenues they generate, but also how they affect the passengers' willingness to pay for a ticket at date 0. To formalize the trade-offs we will need some notation.

We denote the expected utility of a ticketed passenger as a function of his date-0 signal, given the s-mechanism, by $\overline{V}^s(\cdot)$. For the unticketed passenger, it is $\underline{V}^s(\cdot)$. These expected payoffs are gross of the ticket price. Thus, for the signals of ticketed passengers θ_j and of unticketed passengers θ_k,

$$\overline{V}^s(\theta_j) = \mathbb{E}\left[U_j^s(\tilde{\mathbf{v}})|\tilde{\theta}_j = \theta_j, \tilde{\theta}_{-j} \in \mathbf{S}^s_{-j}\right] \text{ and } \underline{V}^s(\theta_k) = \mathbb{E}\left[U_k^s(\tilde{\mathbf{v}})|\tilde{\theta}_k = \theta_k, \tilde{\theta}_{-k} \in \mathbf{S}^s_{-k}\right].$$

The expected per-passenger profit earned at date 1 from ticketed and unticketed passengers in the s-mechanism is given by

$$\overline{\Pi}^s = \mathbb{E}\left[t_j^s(\tilde{\mathbf{v}})|\tilde{\theta} \in \mathbf{S}^s\right] \text{ and } \underline{\Pi}^s = \mathbb{E}\left[t_k^s(\tilde{\mathbf{v}})|\tilde{\theta} \in \mathbf{S}^s\right].$$

Note that $\overline{\Pi}^s$ and $\overline{V}^s(\theta_j)$ are defined only for $s > 0$, since, when the number of ticketed passengers is $s = 0$, the profit from and value for a ticketed passenger are both irrelevant. Similarly, $\underline{\Pi}^s$ and $\underline{V}^s(\theta_j)$ are defined only for $s < n$, since, when $s = n$, there are no unticketed passengers.

The expected payoff as of date 0 to a passenger i with signal θ_i who purchases a ticket at price p is $\mathbb{E}_{\tilde{r}}\overline{V}^{\tilde{r}+1}(\theta_i) - p$, where the number of other ticket holders, $\tilde{r}$, is determined according to the binomial distribution given above. The expected payoff to this passenger if not purchasing a ticket is $\mathbb{E}_{\tilde{r}}\underline{V}^{\tilde{r}}(\theta_i)$. In an optimal pricing mechanism, for any $\theta^* \le 1$, the ticket price p must satisfy

$$\mathbb{E}_{\tilde{r}}\overline{V}^{\tilde{r}+1}(\theta^*) - p = \mathbb{E}_{\tilde{r}}\underline{V}^{\tilde{r}}(\theta^*). \tag{5}$$

Equation (5) states that the marginal ticket holder is indifferent between purchasing a ticket and not purchasing. When $\theta^* > 0$, given that signals are continuously distributed on $[0, 1]$, this condition is necessary for passengers to optimally follow the desired ticket purchasing strategy. If $\theta^* = 0$, then a range of ticket prices are consistent with incentive compatibility, however Eq. (5) specifies the airline's uniquely profit-maximizing choice.[17]

Equation (5) allows us to express the ticket price in terms of the marginal ticket-buying

[17] A lower choice of p is consistent with $\theta^* = 0$, but the airline can then increase profits by raising p without affecting passenger incentives.

signal. For any threshold signal θ^*

$$p=\sum_{r=0}^{n-1}\binom{n-1}{r}\left(1-F\left(\theta^*\right)\right)^r F\left(\theta^*\right)^{n-1-r}\left[\overline{V}^{r+1}(\theta^*)-\underline{V}^r(\theta^*)\right]. \tag{6}$$

The airline's expected profit is therefore

$$\pi=n\left(1-F\left(\theta^*\right)\right)p+\mathbb{E}_{\tilde{s}}[\tilde{s}\overline{\Pi}^{\tilde{s}}+(n-\tilde{s})\underline{\Pi}^{\tilde{s}}], \tag{7}$$

where the number of ticket purchasers, $\tilde{s}$, has binomial distribution with parameters $(n, 1-F(\theta^*))$. After substituting Eq. (6) and rearranging, expected profit can be expressed in terms of the date-1 mechanisms and a threshold ticket-buying signal θ^*, as Eq. (8) below.

Lemma 3. *The airline's expected profit can be expressed as*

$$\sum_{s=0}^{n}\phi^s(\theta^*)\left[s\left(\overline{\Pi}^s+\overline{V}^s(\theta^*)\right)+(n-s)\left(\underline{\Pi}^s-\frac{1-F\left(\theta^*\right)}{F\left(\theta^*\right)}\underline{V}^s(\theta^*)\right)\right], \tag{8}$$

where the probability of selling exactly s tickets, $\phi^s(\theta^*)$, is defined by

$$\phi^s(\theta^*)=\binom{n}{s}F\left(\theta^*\right)^{n-s}\left(1-F\left(\theta^*\right)\right)^s.$$

From Equation (8), we can directly see several implications for our analysis. First, it suggests that we may be able to solve for the optimal reallocation mechanism for each s separately. We take advantage of this separability to characterize optimal pricing mechanisms in what follows. Second, it emphasizes the relevant factors in revenue maximization. Of course, it includes revenue from selling seats to passengers on the day of departure, $\underline{\Pi}^s$, as well as (what will be nonpositive) revenue from buying back seats allocated to ticketed passengers, $\overline{\Pi}^s$. Additionally, it includes the surplus that the marginal ticketed passenger expects to receive, $\overline{V}^s$, and the value of the surplus of the marginal unticketed passenger, $\underline{V}^s$, which affects the revenue negatively, since it decreases the motivation to purchase tickets in advance. The surplus term $\underline{V}^s$ is weighted by $[1-F(\theta^*)]/F(\theta^*)$ reflecting the relative probability of a passenger being ticketed rather than unticketed; i.e., the importance of $\underline{V}^s$ in determining profits is increasing in the relative probability that a ticket is sold.[18] Notice that the surplus terms $\underline{V}^s$ and $\overline{V}^s$ are determined for marginal

[18]This is easiest to understand when $n=1$, in which case the expression in Equation (8) becomes

$$\left(1-F\left(\theta^*\right)\right)\left(\overline{\Pi}^1+\overline{V}^1\left(\theta^*\right)\right)+F\left(\theta^*\right)\left(\underline{\Pi}^0-\frac{1-F\left(\theta^*\right)}{F\left(\theta^*\right)}\underline{V}^0\left(\theta^*\right)\right)=\left(1-F\left(\theta^*\right)\right)\left(\overline{\Pi}^1+p\right)+F\left(\theta^*\right)\underline{\Pi}^0.$$

consumers, because they are the ones determining the ticket price, whereas revenue terms are averages over the infra-marginal consumers.

Illustrative Example

In this section, we will illustrate the main trade-offs faced by the airline in determining how to price tickets, how to allocate seats, and whether to overbook. We consider two passengers and one seat to be allocated. Three kinds of date-1 mechanism are then possible:

1. The "repurchase mechanism", where both passengers purchased tickets at date 0.
2. The "spot mechanism", where neither passenger purchased a ticket at date 0.
3. The "reallocation mechanism", where precisely one passenger purchased a ticket at date 0.

Departing momentarily from the distributional assumptions of the model set-up, passengers in our example either arrive at date 0 with a signal (or "type") in $\{\underline{\theta}, \overline{\theta}\}$, with $\underline{\theta} < \overline{\theta}$, or they arrive at date 1, with type $\varnothing$. These types are drawn independently with equal probability, i.e., the probability of each type is 1/3. In the second period, each passenger i realizes willingness to pay v_i, which is either $\underline{v}$ or $\overline{v}$ where $0 < \underline{v} < \overline{v}$. For illustrative purposes, the signal $\overline{\theta}$ represents a passenger who is certain he will have a high value for flying (i.e., with probability 1, his willingness to pay is $\overline{v}$). On the other hand, $\underline{\theta}$ and $\varnothing$ represent passengers who are uncertain about their eventual willingness to pay: they can have values $\underline{v}$ or $\overline{v}$ with equal probability, i.e., $g(\overline{v}|\underline{\theta}) = g(\overline{v}|\varnothing) = 1/2$.

We suppose that the airline sells tickets to passengers present at date 0 (i.e., to passengers with signals $\underline{\theta}$ and $\overline{\theta}$). As we note formally at the end of this section, this is without loss of optimality for the airline, and in fact permits the airline to achieve the optimal profit from any mechanism, not only mechanisms in our restricted class. We comment below on instances where the airline *strictly* profits from selling tickets, which will only be the case for certain values of $\underline{v}$ and $\overline{v}$.

It turns out that we can restrict attention to mechanisms whose final allocation has "no distortion at the top". That is, if there is at least one passenger with high value $\overline{v}$, then, regardless of how many tickets were purchased in the first period, the seat will not be left empty and at least one of the high-value passengers will be seated (without loss of generality, we assume that a ticketed passenger is seated ahead of an unticketed passenger). That such a policy is optimal is unsurprising, since choosing not to seat a passenger with value $\overline{v}$ does not reduce passenger rents. Thus, the problem reduces to deciding how to treat two passengers when both have low realized values ($\underline{v}$), depending on which of the two (if any) purchased tickets in the first period.

The repurchase mechanism

If two tickets are sold at date 0, the flight is overbooked. As noted above, the question to be answered is whether to allocate a seat when both passengers have low values.

Note, from Eq. (8), that the airline's repurchase mechanism is chosen to maximize

$$\overline{\Pi}^2 + \overline{V}^2(\theta^*). \tag{9}$$

Thus, the airline is maximizing a hybrid welfare function which is the sum of date-1 profits $\overline{\Pi}^2$ from the *average* ticket holder, whose signal is either $\underline{\theta}$ or $\overline{\theta}$, plus the anticipated payoff of the *marginal* ticket holder $\overline{V}^2(\theta^*)$ (here, $\theta^* = \underline{\theta}$). Intuitively, repurchasing tickets on an overbooked flight requires a compensation to ticket holders and potentially raises the value of holding a ticket, which in turn increases the price the airline can charge for tickets in the first period. On the other hand, such a compensation directly reduces the airline's profits. The fact that these two terms are conditioned on different events creates a wedge which implies that these costs and benefits are not offset one-for-one.

To illustrate, consider first the effect on the expected payoff to a given passenger with signal $\underline{\theta}$, i.e., $\overline{V}^2(\underline{\theta})$, of a one dollar cash transfer to that passenger in the event that both passengers have low values. Conditional on having signal $\underline{\theta}$, a passenger expects to have value $\underline{v}$ with probability 1/2 and expects the other ticket holder to have value $\underline{v}$ with probability $1/2 \cdot 1/2 = 1/4$. The transfer thus increases $\overline{V}^2(\underline{\theta})$ by 1/8.

On the other hand, the airline understands that each ticket holder has value $\underline{v}$ independently, with probability 1/4. Hence, the transfer reduces $\overline{\Pi}^2$ by 1/16. The net effect on the airline's objective is strictly positive. A one dollar transfer increases $\overline{\Pi}^2 + \overline{V}^2(\theta^*)$ by 1/16.

It follows that the airline has an incentive to increase the size of the subsidy as much as possible. The size of the subsidy is constrained by incentive compatibility: a buyer with value $\overline{v}$ must not prefer to claim a value $\underline{v}$.

How seats are allocated affects the incentive constraint. For example, consider an efficient mechanism where exactly one of the passengers is seated if both report values $\underline{v}$, with each passenger being seated with equal probability. In this case, the transfer τ to a given passenger cannot be larger than $\overline{v}/2$. To see why, note that, if the passenger in question has value $\overline{v}$ but reports value $\underline{v}$, he would fly with probability 1/2, receive subsidy τ, and obtain payoff $\tau + \overline{v}/2$. Incentive compatibility requires that this payoff be no larger than the payoff from reporting $\overline{v}$ truthfully, flying with probability 1 and receiving no subsidy:

$$\tau + \overline{v}/2 \leq \overline{v}. \tag{10}$$

As an alternative, the airline could use an inefficient mechanism which seats neither passenger when both ticket holders report low values. This enables the subsidy to be raised to $\overline{v}$ without violating the incentive compatibility of a high-value passenger. We can calculate the net effect on the airline's objective as follows. Increasing the subsidy from $\overline{v}/2$ to $\overline{v}$ decreases $\overline{\Pi}^2$ by $1/16 \cdot \overline{v}/2$ but increases $\overline{V}^2(\theta^*)$ by $1/8 \cdot [\overline{v}/2 - \underline{v}/2]$.

The latter reflects that in the eyes of the *marginal* ticket holder, the subsidy will be triggered with higher probability (i.e., with probability 1/8) than in the eyes of the average ticket holder (with probability 1/16), and in this event, his subsidy increases by $\overline{v}/2$ but his utility from flying decreases by $\underline{v}/2$ (in the efficient mechanism he would have had an equal chance of retaining his seat). Thus, the distorted mechanism with a high subsidy is preferred by the airline whenever

$$\left(\frac{1}{8} - \frac{1}{16}\right)\frac{\overline{v}}{2} - \frac{1}{8}\frac{\underline{v}}{2} > 0 \tag{11}$$

which is equivalent to $\overline{v} > 2\underline{v}$. If this condition holds, then it is profitable to leave the seat empty and increase the transfer to the low-value passenger.

The Spot Mechanism

Now consider the case in which no tickets were sold in the first period. In this case the airline may sell the seat in a spot auction to passengers arriving in the second period. According to Eq. (8), it will use a mechanism which maximizes

$$\underline{\Pi}^0 - \frac{1 - F(\theta^*)}{F(\theta^*)}\underline{V}^0(\underline{\theta}) = \underline{\Pi}^0 - 2\underline{V}^0(\underline{\theta}). \tag{12}$$

The airline trades off profit in the spot mechanism from contracting with new passengers who arrive with probability 1/3 against the implications for the rent that must be left to passengers who arrive at date 0 through adjustments to the ticket price. In particular, $\underline{V}^0(\underline{\theta})$ represents the outside option from not purchasing a ticket available to the marginal ticket holder, who has signal $\underline{\theta}$. Increases in this outside option require corresponding reductions in the ticket price.

Again, the question we need to resolve is whether to seat one of the passengers in the event that both have low values. In particular, consider seating each passenger with equal probability in this event, rather than seating nobody. This has three implications for Eq. (12):

1. It allows the airline to charge a passenger $\underline{v}/2$ when both passengers have low values, since each passenger is seated with probability $1/2$. This increases $\underline{\Pi}^0$. Conditional on no tickets sold, the event occurs with probability $1/4$.

2. It reduces the price that can be charged to a high-value passenger when the other passenger has a low value by $(\overline{v} - \underline{v})/2$ (the passenger is now charged $\overline{v} - (\overline{v} - \underline{v})/2$). This reduces $\underline{\Pi}^0$. Again, this event occurs with probability $1/4$.

3. It increases by $(\overline{v} - \underline{v})/2$ the rent a passenger can expect in the spot auction if his value is high and the other passenger's is low. A passenger with signal $\underline{\theta}$ who chooses not to purchase a ticket believes that this occurs with probability $1/4$ in the spot auction. The effect is to increase $\underline{V}^0(\underline{\theta})$.

The overall effect of seating a low-value passenger (at random) on $\underline{\Pi}^0$ is then $\underline{v}/4 - \overline{v}/8$, while the effect on $\underline{V}^0(\underline{\theta})$ is $\overline{v}/8 - \underline{v}/8$. The overall effect on Eq. (12) is $\underline{v}/2 - (3/8)\,\overline{v}$, so the airline should seat a low-value passenger if and only if $\underline{v} > (3/4)\,\overline{v}$.

It is worth comparing the airline's objective here to two alternatives. First, the mechanism that maximizes $\underline{\Pi}^0$ is simply the optimal spot auction, conditional on both passengers receiving signal $\varnothing$. As explained above, simply maximizing $\underline{\Pi}^0$ would ignore the effects of seating a low-value passenger on the ticket price (in particular, the ticket price which is acceptable to the marginal ticket-holder, who has signal $\underline{\theta}$). Second, one might compare the airline's objective to that in the case where tickets are not offered at date 0, so that passengers only participate at date 1. A further difference here is that the distribution of passenger values would be different (the probability of a high value would be 2/3 rather than 1/2).

The Reallocation Mechanism

Finally, consider the case in which a single ticket is purchased in the first period. According to Eq. (8), the reallocation mechanism should maximize

$$\overline{\Pi}^1 + \underline{\Pi}^1 + \overline{V}^1(\underline{\theta}) - \frac{1 - F(\theta^*)}{F(\theta^*)}\underline{V}^1(\underline{\theta}). \tag{13}$$

As argued above, a high-value passenger always flies ahead of a low-value passenger. Hence, a low-value ticketed passenger, for instance, is optimally induced to give up his seat in preference to a high-value unticketed one. The question is then which passenger (i.e., the ticketed or the unticketed passenger), if any, to seat if both have low values? It turns out Eq. (13) is maximized in our example by either (i) letting the ticketed passenger keep his seat in preference to a low-value unticketed passenger, or (ii) seating no low-value passenger. Flying empty is strictly better if and only if $\overline{v} > 2\underline{v}$, which is the same condition as for the repurchase mechanism.

The optimal policy described here reflects the same trade-offs as for the repurchase and spot mechanisms above. For instance, note that reallocating the seat to an unticketed low-value passenger would increase $\underline{V}^1(\underline{\theta})$, which would increase the outside option to the marginal ticket holder, type $\underline{\theta}$, of not purchasing a ticket. This would require a reduction in the ticket price.

First Period Ticket Sales

One can now check when the airline strictly profits from selling tickets at date 0 (and then overbooking the flight whenever both passengers arrive at date 0), rather than contracting with passengers only at date 1. It turns out that the airline profits in case $\overline{v} < 2\underline{v}$, since this implies that low-value passengers are seated in the optimal mechanism with positive probability. One way to understand this is that, for these parameters, tickets allow the airline to distinguish between low-value passengers arriving at date 0 and those who arrive at date 1 (recall that a ticketed passenger will always fly ahead of an unticketed low-value passenger).

More generally, we can understand the role of ticket sales in the following way. Consider first the case in which the optimal spot auction is efficient, i.e., always sells the seat to the passenger with the highest value (breaking ties at random). In such a mechanism, an early-arriving passenger with signal $\underline{\theta}$ earns rents in expectation. This is because there is a positive probability that his realized value will be $\overline{v}$ and the other passenger will have value $\underline{v}$. In that case the auction price will be $\underline{v}$ and he will earn rents equal to $\overline{v} - \underline{v}$. Consider now a two-stage mechanism in which

1. Tickets are sold at a price equal to $\underline{\theta}$'s expected rents from the spot auction.

2. Unticketed passengers are excluded.

3. The allocation and transfer rules are otherwise the same as the spot auction. In particular the value of an unticketed passenger is still used to determine the allocation and price paid by a ticketed passenger.

In this mechanism, buying a ticket is equivalent to buying admission to the spot auction. The ticket price extracts the expected rents of $\underline{\theta}$, and reduces the rents of $\overline{\theta}$ by an equal amount, but leaves the welfare from trade with ticketed passengers unchanged. These effects raise profits relative to the spot auction. Any losses come from exclusion of unticketed passengers. Thus, if the probability of $\varnothing$ is sufficiently low, then this two-stage mechanism improves upon the spot auction.

Exclusion of unticketed passengers ensures that an early arriver would obtain zero utility if he were to refuse to buy a ticket. This increases the additional surplus of the marginal passenger (i.e., the passenger with signal $\underline{\theta}$) from holding a ticket and this surplus will be extracted through the ticket price. Total exclusion of unticketed passengers is not necessary to achieve this. It is enough that the airline refuses to seat a low-value passenger unless he is holding a ticket. This guarantees that high-value unticketed passengers earn no information rents and that again, the expected utility of an unticketed passenger is zero.[19]

Unrestricted optimum

In the discussion above we focused attention on mechanisms which sell tickets to all passengers present at date 0, i.e., passengers with signals $\overline{\theta}$ and $\underline{\theta}$. Under the assumptions of the illustrative example, this turns out to be without loss of optimality. Intuitively, if the airline does not sell tickets to the passengers with low signals, then it sells tickets only to passengers who already know their valuation, and in this case it can obtain the same profits through the spot mechanism on date 1.

Lemma 4. *The mechanism described above achieves the highest profit attainable by* any *mechanism.*

4 General Analysis

General properties of the optimal mechanism

Before describing how we approach the problem of characterizing the optimal allocation rule, we propose a system of date-1 transfers which maximizes profit for the airline given any implementable allocation rule $\mathbf{q}$. By the envelope theorem, the date-1 mechanism defined above must satisfy, for any s, any passenger i, and any $\mathbf{v}$,

$$U_i^s(\mathbf{v}) = U_i^s(\overline{v}, \mathbf{v}_{-i}) - \int_{v_i}^{\overline{v}} q_i^s(y, \mathbf{v}_{-i})dy = U_i^s(\underline{v}, \mathbf{v}_{-i}) + \int_{\underline{v}}^{v_i} q_i^s(y, \mathbf{v}_{-i})dy. \tag{14}$$

Note that the share of the surplus a passenger i expects to obtain at date 1, for each s and $\mathbf{v}_{-i}$, is determined up to a constant by his value v_i and the allocation rule $\mathbf{q}$. For

[19] Alternatively, the airline may choose to seat unticketed low-value passengers, but only when there are no passengers already holding tickets. As explained above, favoring ticketed passengers in this way again reduces the rents available from not purchasing a ticket.

ticketed passengers, adjusting this constant does not affect expected payoffs at date 0 provided the ticket price is correspondingly adjusted according to Eq. (5). It is therefore without loss of optimality to focus on date-1 mechanisms such that, for all s, all ticketed passengers j, and all $\mathbf{v}_{-j}$, $U_j^s(\overline{v}, \mathbf{v}_{-j}) = \overline{v}$. On the other hand, adjusting period-1 payoffs for unticketed passengers by a constant *does* affect ex-ante payoffs. In particular, if the payoff earned by the passenger with the minimum value $U_k^s(\underline{v}, \mathbf{v}_{-k})$ is greater than zero, the airline can profit by increasing transfers by unticketed passengers by choosing an appropriate constant. This not only reduces the rent left to unticketed passengers but also allows the airline to increase ticket prices according to Eq. (5). We can thus focus on mechanisms satisfying $U_k^s(\underline{v}, \mathbf{v}_{-k}) = 0$ for all s, all unticketed passengers k and all $\mathbf{v}_{-k}$. The corresponding transfers are given by

$$
\begin{aligned}
t_j^s(v_j, \mathbf{v}_{-j}) &= v_j q_j^s(v_j, \mathbf{v}_{-j}) + \int_{v_j}^{\overline{v}} q_j^s(y, \mathbf{v}_{-j}) dy - \overline{v}, \\
t_k^s(v_k, \mathbf{v}_{-k}) &= v_k q_k^s(v_k, \mathbf{v}_{-k}) - \int_{\underline{v}}^{v_k} q_k^s(y, \mathbf{v}_{-k}) dy.
\end{aligned}
\tag{15}
$$

The airline's profit now depends only on the allocation rule. Note also from Eq. (14) that passenger payoffs are not affected by adjusting the allocation to ensure passengers with values $\overline{v}$ fly wherever possible (i.e., setting $q_i^s(\overline{v}, \mathbf{v}_{-i}) = 1$ wherever possible and adjusting transfers accordingly). Moreover, since a buyer i's allocation remains monotonic in his valuation,[20] the allocation remains implementable. In other words, it is possible and costless to provide the passenger with the option to fly by reporting $\overline{v}$ whenever the number of passengers reporting $\overline{v}$ is less than the available seats m. This observation, together with Eq. (15), implies the following result.

Proposition 1. *It is optimal to structure the pricing mechanism so that*

1. *A ticket is an option to fly and passengers are never bumped involuntarily.*[21]
2. *Ticket holders who are seated make (and receive) no payments.*
3. *Any ticket holder who does not fly receives a payment which at least matches his value in compensation.*

Part 1 of the proposition states that, at least in the above framework, airlines cannot profit from using involuntary bumping. In particular, there is no loss to airlines in allowing

[20] This monotonicity is both necessary and sufficient for the ex-post implementability of the allocations by the date-1 mechanism.

[21] Note that passengers will not be bumped involuntarily, even in case more than m passengers have the highest value $\overline{v}$. In this case, all of these buyers will be indifferent between keeping their seat and instead taking the compensation $\overline{v}$ on offer.

those with the highest value of flying to do so. Indeed, the arrangement can be optimally structured so that ticketed passengers do not pay anything to keep their seats (Part 2 of the proposition) but are compensated for giving up a seat. Under this arrangement, enticing passengers with high values to give up their seats (as required, if the flight is overbooked and the passengers with tickets have the high values), while still seating those with the *very high* valuations, requires a high compensation to the unseated ticketed passengers. While the compensation provided by the airline in such cases may be large, the surplus provided to ticketed passengers can be recouped through the date-0 ticket price.

Note that the optimality of a mechanism without "involuntary bumping" follows from the use of sufficiently sophisticated auction mechanisms for resolving which passengers will fly. For instance, if the flight is overbooked by passengers with high values for being seated, the level of refund offered responds to their reported values and is correspondingly high. A possible disadvantage of using such mechanisms in practice is that both the available refund and the seat price for unticketed passengers is completely determined only after passengers report their values on date 1. Airlines may favor simpler mechanisms where prices are specified before passengers have learned their final values for flying. If these mechanisms cannot respond to excess demand (e.g., when there are more ticket-holders than there are seats that realize high values for flying), then involuntary bumping would be necessary to reconcile the demand with the number of seats.[22]

Optimal allocations

We now derive the optimal allocation. From Eq. (8), for each s, the airline's profits are proportional to

$$s\left(\overline{\Pi}^s + \overline{V}^s(\theta^*)\right) + (n-s)\left(\underline{\Pi}^s - \frac{1-F(\theta^*)}{F(\theta^*)}\underline{V}^s(\theta^*)\right). \tag{16}$$

From Eq. (14), the revenues earned from a ticketed passenger j and non-ticketed passenger k through the s-mechanism are

$$\overline{\Pi}^s = \int_{\mathbf{v}} t_j^s(\mathbf{v})dG(\mathbf{v}|\mathbf{S}^s) = \int_{\mathbf{v}} \left[v_j q_j^s(\mathbf{v}) - U_j^s(\mathbf{v})\right] dG(\mathbf{v}|\mathbf{S}^s)$$

[22]We believe the analysis below nonetheless sheds some light on optimal pricing under less flexible or responsive mechanisms. For instance, note that in the limit, as the number of available seats and passengers becomes large, the variation in the compensation made available to ticket holders diminishes; i.e., compensation must also be approximately predictable in the class of mechanisms that we consider in this paper.

and

$$\underline{\Pi}^s = \int_{\mathbf{v}} \left[v_k q_k^s(\mathbf{v}) - U_k^s(\mathbf{v})\right] dG(\mathbf{v}|\mathbf{S}^s)$$

and the welfare of the passenger with the marginal signal is

$$\overline{V}^s(\theta^*) = \int_{v_j} \int_{\mathbf{v}_{-j}} U_j^s(\mathbf{v}) dG(\mathbf{v}_{-j}|\mathbf{S}_{-j}^s) dG(v_j|\theta^*) \tag{17}$$

for a ticketed passenger j and

$$\underline{V}^s(\theta^*) = \int_{v_k} \int_{\mathbf{v}_{-k}} U_k^s(\mathbf{v}) dG(\mathbf{v}_{-k}|\mathbf{S}_{-k}^s) dG(v_k|\theta^*) \tag{18}$$

for a passenger k without a ticket.

Focusing now on the expressions for ticketed passengers, using Eq. (14) and integration by parts yields

$$\begin{aligned}\overline{\Pi}^s = & -\int_{\mathbf{v}_{-j}} U_j^s(\overline{v}, \mathbf{v}_{-j}) dG(\mathbf{v}_{-j}|\mathbf{S}_{-j}^s) \\ & + \int_{\mathbf{v}} q_j^s(\mathbf{v}) v_j dG(\mathbf{v}|\mathbf{S}^s) + \int_{\mathbf{v}_{-j}} \left[\int_{v_j} G(v_j|\tilde{\theta}_j \geq \theta^*) q_j^s(\mathbf{v}) dv_j\right] dG(\mathbf{v}_{-j}|\mathbf{S}_{-j}^s)\end{aligned}$$

and

$$\overline{V}^s(\theta^*) = \int_{\mathbf{v}_{-j}} U_j^s(\overline{v}, \mathbf{v}_{-j}) dG(\mathbf{v}_{-j}|\mathbf{S}_{-j}^s) - \int_{\mathbf{v}_{-j}} \left[\int_{v_j} G(v_j|\theta^*) q_j^s(\mathbf{v}) dv_j\right] dG(\mathbf{v}_{-j}|\mathbf{S}_{-j}^s).$$

Putting these together and collecting terms we obtain an expression for the virtual surplus of ticketed passengers

$$s\left(\overline{\Pi}^s + \overline{V}^s(\theta^*)\right) = \mathbb{E}_{\tilde{\mathbf{v}}}\left[\sum_{j=1}^{s} q_j^s(\tilde{\mathbf{v}}) \overline{\mathrm{VS}}(\tilde{v}_j) \middle| \mathbf{S}^s\right], \tag{19}$$

where

$$\overline{\mathrm{VS}}(v_j) = v_j - \frac{G(v_j|\theta^*) - G(v_j|\tilde{\theta} j \geq \theta^*)}{g(v_j|\tilde{\theta}_j \geq \theta^*)}. \tag{20}$$

As anticipated above, all terms involving $U_j^s(\overline{v}, \mathbf{v}_{-j})$ are canceled: these constants are free variables in the airline's maximization.[23]

Turning now to the terms in Eq. (16) involving unticketed passengers, and using the

[23]As described above, this is intuitive because any constant added to the ticket holder's ex-post utility at date 1 can be recovered via an equal increase in the ticket price at date 0.

fact that unticketed passengers with the lowest values earn zero surplus, we have

$$(n-s)\left(\underline{\Pi}^s - \frac{1-F(\theta^*)}{F(\theta^*)}\underline{V}^s(\theta^*)\right) = \mathbb{E}_{\tilde{\mathbf{v}}}\left[\sum_{k=s+1}^{n} q_k^s(\tilde{\mathbf{v}})\underline{\mathrm{VS}}(\tilde{v}_k)\middle|\mathbf{S}^s\right], \tag{21}$$

where

$$\underline{\mathrm{VS}}(v_k) = v_k - \frac{1-G(v_k|\tilde{\theta}_k<\theta^*) + \frac{1-F(\theta^*)}{F(\theta^*)}\left(1-G(v_k|\theta^*)\right)}{g(v_k|\tilde{\theta}_k<\theta^*)}. \tag{22}$$

To summarize the above, we have derived a convenient expression for the airline's profit in an optimal mechanism implementing an allocation rule $\mathbf{q}$, which satisfies the requirements set out in Section 3.

Lemma 5. *Let s be a given number of tickets sold and consider the date-1 mechanism implementing the allocation $\mathbf{q}^s(\cdot)$. Suppose that the unticketed passenger with the lowest value earns zero surplus in this mechanism. Then the expression in Eq. (16) is equal to*

$$\mathbb{E}\left[\sum_{j=1}^{s} q_j^s(\tilde{\mathbf{v}})\overline{\mathrm{VS}}(\tilde{v}_j) + \sum_{k=s+1}^{n} q_k^s(\tilde{\mathbf{v}})\underline{\mathrm{VS}}(\tilde{v}_k)\middle|\mathbf{S}^s\right]. \tag{23}$$

The transformed expected revenue of Eq. (23) is a familiar expected virtual surplus. As usual, the first pass at a solution is to consider the allocation rule which, at every realized valuation profile, allocates seats to the m passengers with the highest (non-negative) virtual surpluses. One then verifies that the proposed allocation indeed gives precedence to ticketed passengers over unticketed passengers in the sense described in Section 3, and that the ticket price defined by Eq. (5) and the transfers defined by Eq. (15) implement the allocation. The following conditions on the virtual surpluses are sufficient to guarantee this.

Proposition 2. *Fix $\theta^* \in [0,1]$. Suppose that (i) $\max\left\{\overline{\mathrm{VS}}(v_j), 0\right\}$ and $\max\left\{\underline{\mathrm{VS}}(v_k), 0\right\}$ are non-decreasing functions, and (ii) $\overline{\mathrm{VS}}(v_i) \geq \underline{\mathrm{VS}}(v_i)$ for all $v_i \in \mathrm{Supp}\left[G\left(\cdot|\tilde{\theta}\geq\theta^*\right)\right] \cap \mathrm{Supp}\left[G\left(\cdot|\tilde{\theta}<\theta^*\right)\right]$, a nonempty set.*[24] *The allocation which maximizes virtual surplus (conditional on the marginal signal θ^*) is implementable and maximizes the airline's profit.*

Condition (i) ensures that an extended allocation can be found such that each passenger's probability of being seated is non-decreasing in his own report (whether or not that report is made in equilibrium). It thus guarantees the existence of an ex-post incentive-compatible period-1 mechanism which implements the allocation. Condition (ii) ensures

[24] The requirement that $\mathrm{Supp}\left[G\left(\cdot|\tilde{\theta}\geq\theta^*\right)\right] \cap \mathrm{Supp}\left[G\left(\cdot|\tilde{\theta}<\theta^*\right)\right]$ be nonempty is a minor technical requirement, which is satisfied in most of the cases of interest. For instance, it is enough to suppose that type $\varnothing$ has positive probability and that $G\left(\cdot|\varnothing\right)$ has full support on $[\underline{v}, \overline{v}]$.

that an allocation $\mathbf{q}$ can be found that gives precedence to ticketed passengers over unticketed passengers in the sense described in Section 3. Given that the ticket price satisfies Eq. (5), this guarantees that passengers will purchase tickets if and only if their signal exceeds θ^*.

Note that, given the conditions of Proposition 2, the optimal pricing mechanism has the property that, if there are two passengers with the same value and the one without a ticket is seated, then the passenger with a ticket must be seated as well. This preferential treatment of ticketed passengers relative to unticketed passengers is different from our assumption above, where we assumed that having a ticket cannot lower the chances of being seated for any individual passenger (we make the same point in footnote 14 above).

Understanding the trade-offs

Ticketed Passengers

We now consider the allocations described in Proposition 2 and how these reflect the airline's incentives. The virtual surplus of ticket holders is given by

$$\overline{\text{VS}}(v_j) = v_j - \frac{G(v_j|\theta^*) - G(v_j|\tilde{\theta}_j \geq \theta^*)}{g(v_j|\tilde{\theta}_j \geq \theta^*)}. \tag{24}$$

To interpret this, first consider that the decision not to seat a ticketed passenger j, i.e., by setting $q_j^s(\mathbf{v}) = 0$ for some s and realized values $\mathbf{v}$, is an advance commitment to a ticket purchaser j that his ticket will be repurchased in the event that s passengers hold tickets and the profile of all passengers' realized valuations is $\mathbf{v}$. Such a commitment impacts the airline's profits via two effects. First, it raises the payment to a ticketed passenger announcing value v_j and thus, by incentive compatibility, requires the airline to raise by an equal amount the utility of all types lower than v_j. This directly reduces the airline's profits within the date 1 mechanism by $G(v_j|\tilde{\theta}_j \geq \theta^*)$ and correspondingly lowers the airline's willingness to repurchase the ticket from passenger j. Note that $G(v_j|\tilde{\theta}_j \geq \theta^*)$ measures the probability that j's value will fall below v_j, conditional on j's purchase of a ticket at date 0.

The second effect operates indirectly via revenues from date 0 ticket sales. Indeed, some of the additional utility provided by the airline through date 1 repurchases can be recouped via increased ticket prices. At the time of ticket purchase, the marginal type θ^* assesses a probability $G(v_j|\theta^*)$ that his value will fall below v_j and that he will benefit from the increased utility resulting from the airline's commitment to repurchase a ticket from type v_j. That leads to an increased willingness to pay for a ticket that can be extracted dollar-for-dollar by increasing the price of a ticket. This indirect effect raises the airline's profits by $G(v_j|\theta^*)$ where, crucially, this measures the probability that j's value will fall below v_j, conditional on i being the *marginal* ticket purchaser, θ^*.

Putting all of this together, the airline chooses to seat a passenger based on the extent to which the surplus created v_j exceeds a measure that accounts for the net effect on the airline's profits. That measure is proportional to the difference $G(v_j|\theta^*) - G(v_j|\tilde{\theta}_j \geq \theta^*)$ in the repurchase probability assessed by the marginal and average ticket purchasers (θ^* and $\theta_j \geq \theta^*$ respectively).

In light of this wedge between marginal and average ticket holders, overbooking and repurchasing can be seen as an instrument to refine the screening of passengers by willingness to pay at date 0. For illustrative purposes, consider an airline that practices no overbooking and seats all ticketed passengers. One way the airline could seek to increase profits is to raise the ticket price, thus reducing the surplus earned by infra-marginal ticket holders. However, doing so affects the decision of the marginal ticket holders to purchase tickets, so this improvement in screening comes at the cost of reduced ticket sales. Overbooking enables the airline to capture consumer surplus without sacrificing ticket sales. The airline raises ticket prices and effectively strikes a deal with the marginal type θ^* that the price increase will be returned in expectation through repurchases in the event that his valuation turns out to be low. Since this is calculated to be a one-for-one intertemporal transfer for the marginal type θ^*, it is less favorable for passengers with signals above θ^* because they assess a strictly lower probability of repurchase. Thus, the higher price coupled with potential repurchase maintains the indifference of the marginal type and strictly lowers the consumer surplus of infra-marginal types.

Unticketed Passengers

Consider now the virtual surplus of an unticketed passenger.

$$\underline{\text{VS}}(v_k) = v_k - \frac{1 - G(v_k|\tilde{\theta}_k < \theta^*) + \frac{1-F(\theta^*)}{F(\theta^*)}\left(1 - G(v_k|\theta^*)\right)}{g(v_k|\tilde{\theta}_k < \theta^*)} \tag{25}$$

Just as with ticketed passengers, the incentive to seat an unticketed passenger mixes the direct effect on revenues in the s-mechanism with an indirect effect on revenues from ticket sales. Indeed, the virtual surplus can be seen as the sum of two terms, the first of which,

$$v_k - \frac{1 - G(v_k|\tilde{\theta}_k < \theta^*)}{g(v_k|\tilde{\theta}_k < \theta^*)} \tag{26}$$

is the familiar expression for the marginal revenue from selling to buyer k in a standard monopoly problem. It has the one noteworthy difference: the expressions are conditioned on passenger k having a date 0 signal below the threshold θ^*. The remaining term in $\underline{\mathrm{VS}}$ is a correction which accounts for the effect on first period ticket sales from a date 1 decision to seat an unticketed passenger. A mechanism which allocates space to unticketed passengers reduces the value of holding a ticket, and consequently reduces the revenue from ticket sales. This term is conditional on θ^* because it is the marginal ticket holder's willingness to pay that determines the ticket price.

Optimal ticket price

The above analysis takes the marginal signal θ^* as given, then finds the optimal allocation conditional on this cut-off. The final step in deriving the profit-maximizing pricing mechanism is to deduce the optimal choice of θ^*. We begin by providing a sufficient condition to guarantee the profitability of setting $\theta^* < 1$ and hence selling tickets with positive probability.

Corollary 1. *Suppose that there exists* $\theta^* \in [0, 1)$ *such that (i)* $\max\left\{\overline{\mathrm{VS}}(v_j), 0\right\}$ *and* $\max\left\{\underline{\mathrm{VS}}(v_k), 0\right\}$ *are non-decreasing functions, and (ii)* $\overline{\mathrm{VS}}(v_i) \geq \underline{\mathrm{VS}}(v_i)$ *for all* $v_i \in \mathrm{Supp}\left[G\left(\cdot|\tilde{\theta} \geq \theta^*\right)\right] \cap \mathrm{Supp}\left[G\left(\cdot|\tilde{\theta} < \theta^*\right)\right]$*, a nonempty set, with a strict inequality for an interval of positive length.*[25] *The airline strictly prefers a pricing mechanism with ticket sales at date 0 and a marginal ticket holder* θ^* *over a date-1 spot auction (where passengers only contract with the airline at date 1). If* $n > m$*, then given marginal ticket holder* θ^* *such that the above conditions are satisfied, the airline strictly profits by overbooking the flight with positive probability.*

The optimality of selling tickets at date 1 under these conditions follows the same logic as in the stylized example of Section 3. The conditions guarantee that there exists some threshold θ^* such that the optimal mechanism conditional on θ^* (as determined above) implements a different allocation rule for ticketed versus unticketed passengers. For such θ^*, the airline finds tickets a useful screening device to distinguish between passengers arriving at date 0 with strong beliefs that their values for flying will be high and other passengers (those who believe their values are more likely to be low, or who arrive at date 1). Since signals are drawn independently, there is a positive probability that all n passengers observe signals above θ^* at date 1. Hence, if the number of seats m is less than n, the flight is overbooked with positive probability. Note here that, under the conditions in Corollary 1, the airline does strictly better by overbooking the flight (selling tickets to

[25]The conditions ensure that there are some values for which, under the optimal allocation given θ^*, holding a ticket *matters* in deciding whether a passenger is seated. While we find these conditions straightforward to give in light of Proposition 2; we expect one can also find weaker conditions to guarantee the profitability of offering tickets in advance of the flight.

any passengers with signals above some θ^*) than randomly rationing tickets to ensure the flight is not overbooked (by allocating any passenger with signal above θ^* a ticket with an appropriate probability of less than one).

We can contrast Corollary 1 with a particularly simple case in which ticket sales and overbooking do not improve profits. This is where all passengers in the market at date 0 are perfectly informed of their eventual values for flying at date 1.

Example 1. *When signals are fully informative (i.e., passengers arriving at date 0 know their values with certainty at that date), a spot auction at date 1 is an optimal mechanism.*

The reason for this result is related to the discussion of our stylized example in Section 3. As we explained in this example, the reason an airline finds advanced ticket sales useful in our model is that it allows the airline to capture some of the rents that passengers expect to earn due to the possibility that their values for flying may turn out to be higher than expected (in the stylized example, when their values equal $\overline{v}$). When passengers know their values in advance, this possibility does not arise.

Finally, calculating the precise value of θ^* is complicated because changing θ^* affects not only the beliefs of the marginal ticket holder (whose signal is θ^*) about his eventual value for flying, but also the distribution of values among ticketed and unticketed passengers (conditional on each number of tickets sold). This not only affects profits in the date-1 mechanism, but also the likelihood that the marginal ticket holder is finally seated. Nonetheless, the determinants of the optimal choice of θ^* are easier to understand in the context of a simple example like the following.

Example 2. *Suppose that, for each passenger i, $\tilde{\theta}_i$ and $\tilde{\varepsilon}_i$ are independently and uniformly distributed on $[0, 1]$. Suppose that passengers' values are determined by $v_i = \theta_i + a\varepsilon_i$.*

1. *Fix any $m, n \geq 1$. There exists $\overline{a}$ such that, for all $a \geq \overline{a}$, $\theta^* = 0$; i.e., all passengers purchase tickets.*

2. *Fix $m \geq n$ (so there are no capacity constraints), and assume $a \in (0, 1/4)$. Then $\theta^* = 1/2 - a/4$.*

3. *Fix any $a, \gamma \in (0, 1/2)$ and let $m = \lfloor \gamma n \rfloor$. There exists $\overline{n}$ such that, for all $n \geq \overline{n}$, $\theta^* \in [1 - \gamma - a, 1 - \gamma + a]$.*

This example features a setting where all n passengers arrive at date 0 with probability 1. The date-0 signal θ_i of passenger i adjusts his final value for a seat one-for-one, while ε_i represents a "shock" to the final value whose size is scaled by the parameter a. Hence, when a is large, passengers are highly uncertain at date 0 as to their values for seats.

Case 1 applies when the parameter a is sufficiently large, i.e., when the passengers are relatively uncertain at date 0 of their future value for flying. In this case, an optimal policy involves selling tickets to all passengers and overbooking occurs with probability 1 whenever $n > m$. The reason this is optimal is simply the rent extraction motive for ticket sales described above. In particular, as a becomes large, the rents that a passenger i can earn on account of a high realization of the shock ε_i become large. The seller can capture much of the expected rent by contracting with the passenger at date 0, but not if contracting is delayed until date 1.

As noted above, a second consideration in determining the threshold for ticket purchases, θ^*, is the aim of discriminating between passengers with favorable and unfavorable beliefs about their eventual values for flying. This can be achieved by selling tickets to those passengers with favorable signals but not to those with unfavorable signals. For instance, when $m \geq n$, so that capacity is unconstrained, as in Case 2, we can calculate θ^* exactly. Passengers with signals above θ^* receive tickets and these passengers are favored in the optimal allocation: if an unticketed passenger with value v_i is seated, then so is a ticketed passenger with the same value.

When capacity constraints bind, calculating the optimal threshold θ^* is more complicated for the reasons explained above. However, the two principles described above continue to apply. Note that Case 1 applies whether or not capacity constraints bind. Case 3 addresses a scenario where the number of passengers is large and capacity constraints are highly likely to bind. In this case, we can observe the following regarding the optimal pricing mechanism conditional on θ^*. If θ^* falls outside the specified bounds, then whether a passenger holds a ticket is highly unlikely to affect the optimal allocation (for large enough n). For example, if $\theta^* < 1 - \gamma - a$, passengers with signals no greater than θ^* would be highly unlikely to be seated in the optimal pricing mechanism (conditional on θ^*), simply because their realized value will be relatively low. Hence, ticket sales do a poor job of screening customers based on their initial beliefs, suggesting $\theta^* < 1 - \gamma - a$ cannot be optimal. This appears to reflect a general principle: The airline would like to choose θ^* so that whether a passenger holds a ticket is likely to matter for the eventual allocation of seats.

Implementation via a double auction

Assuming Proposition 2 applies, it is straightforward to specify a "handicap" double auction which implements the optimal pricing mechanism. The transfers that are made in the auction we propose are identical to those given by Eq. (15).

Some new notation will be convenient for formalizing the rules of the auction. For any ticket-holder value v_j, define the matching value $v_k(v_j)$ for an unticketed passenger to be the highest value satisfying

$$\overline{\mathrm{VS}}(v_j) = \underline{\mathrm{VS}}(v_k(v_j)). \tag{27}$$

Note that $v_k(v_j)$ is increasing in v_j. Conversely, define $v_j(v_k)$ as the matching value of the ticket holder. The rules of the double auction are as follows. Each passenger submits a bid and the airline announces the reserve price R defined to be the highest value satisfying $\overline{\mathrm{VS}}(R) = 0$. Passengers are ranked in descending order of bids and the q highest bidding passengers are allocated a seat provided their bids exceed the reserve R.

Payments are determined as follows. Let b^q and b^{q+1} denote the q and $q+1$st highest bids. In case the number of bids exceeding R is smaller than either q or $q+1$, then b^q and/or b^{q+1} are set equal to R. Any ticketed passenger who is not seated receives compensation equal to b^q. Any unticketed passenger who wins a seat is charged $v_k(b^{q+1})$. The transfers are zero for all ticketed passengers who fly and unticketed passengers who do not.

Proposition 3. *Suppose that the conditions of Proposition 2 hold. Then there exists a dominant-strategy equilibrium of the double auction which implements the optimal pricing mechanism.*

Note that the spread between the price paid by unticketed passengers for seats and the compensation to the unseated ticket holders need *not* be positive. Although seats are transferred from ticketed passengers to unticketed passengers only if the latter have sufficiently higher values than the former, it does not mean that the net transfers are always positive. To see this, consider an example with one seat, one ticketed passenger and one unticketed passenger. Suppose the ticketed passenger has a value close to the lower bound $\underline{v}$ and the unticketed passenger a value close to the upper bound $\overline{v}$. Then the ticketed passenger may receive compensation close to $\overline{v}$, whereas the unticketed passenger may pay a significantly smaller amount for the seat.[26] However, the negative ex-post net revenue occurs only if the difference between b^q and b^{q+1} is sufficiently large and these bids are made by an unticketed and a ticketed passenger, respectively. The probability of this event is very low when the number of passengers is higher than two.

5 Multiple fare classes and unrestricted mechanisms

In this section, we briefly discuss the restriction to pricing mechanisms with a single price at date 1. As we noted, the restriction to a single price, with all passengers making their

[26]The argument is analogous to Myerson and Satterthwaite (1983).

ticket purchasing decisions simultaneously, corresponds to an assumption that passengers can communicate a single message (say "purchase") at date 0. Here, we are motivated by the observation that passengers, in visiting an airline's website, for example, often have the choice of a limited number of options, or "fare classes", and remain uninformed about other passengers' purchasing decisions. Our focus on a single fare class is aimed at simplifying the key trade-offs. We now discuss how our results should be expected to extend to cases where the airline is permitted additional fare classes, or faces no restriction on the number of classes. While the details of how we can adapt our formal analysis are presented in Appendix B, here we outline the main qualitative conclusions that we anticipate from this analysis.

To begin, consider the extreme case in which the airline faces no restriction so that each signal $\theta_i \geq 0$ can be assigned distinct ticket terms. Since the airline has an array of ticket classes at its disposal, it is (at least weakly) optimal to contract with all passengers in the market at date 0. As we noted above, this follows simply from the revelation principle: any outcome which has some passengers delay their participation until date 1 can be replicated by a mechanism which induces participation (and revelation of θ_i) at date 0. In contrast, when there is a single fare class (as we have analyzed in this paper), the choice of signals for which passengers will purchase tickets (as captured by θ^*) plays an important role in screening customers based on their initial information. Setting ticket prices high enough that some passengers do not purchase is a way for the airline to reduce the likelihood of being seated for those passengers who are initially most pessimistic about their values for flying. This reduces the information rents for those passengers who are initially more optimistic.

The optimal mechanism with no restriction on fare classes is easiest to characterize under certain restrictions on the joint distribution of θ_i and v_i. For instance, if all passengers arrive at date 0, the relevant restrictions are familiar from the literature following Baron and Besanko (1984) and Courty and Li (2000). Under such restrictions, the optimal allocation is monotone in the sense that a passenger is more likely to be seated for higher realizations of both the initial signal θ_i and final value v_i. Suppose we focus on an implementation in the spirit of that described in Proposition 1. In particular, passengers reporting high enough values v_i will obtain a seat with probability 1 and make no payments. We can then observe that passengers reporting higher values of θ_i obtain, on average, less compensation for not flying. Indeed, this is necessarily the case in an incentive-compatible mechanism to ensure passengers with the highest values of θ_i have incentives to keep their seats at date 1.[27] That passengers with higher values of θ_i wish

[27] In fact, the observation is true both conditionally on the realization of v_i and unconditionally (the former by incentive compatibility at date 1, and the latter because higher values of θ_i realize higher values

to purchase tickets with less compensation for not flying follows from the fact that such passengers are more certain of having a high value for flying (and hence expect to benefit less from the possibility of compensation).[28]

The case in which airlines offer only a few fare classes is perhaps most closely in line with pricing practice. This case can be understood as intermediate between the extremes of a single fare class and having no restriction on fare classes. The date-0 mechanism would involve passengers simultaneously deciding which among several fare classes (with different prices) to purchase.

Recall our restriction that a passenger is more likely to fly if he purchases a ticket than if he does not purchase one (for each realization of $(v_i, \theta_i)_{i=1}^n$). The natural analogue in a setting with multiple fare classes would feature the tickets purchased by higher signals θ_i increasing the likelihood of flying for a given passenger. We then anticipate that the optimal mechanism conditional on a given number of fare classes would behave as follows, subject to appropriate restrictions on the distribution of signals and values (and hence corresponding virtual surpluses). As with the unrestricted mechanism, passengers with high signals θ_i would purchase cheaper tickets offering less compensation for giving up their seat. These passengers would be willing to purchase tickets with low compensation simply because they are less likely to have low values. This would be consistent with the idea that so-called "flexible" tickets are particularly valuable to passengers who are unsure whether they will enjoy a high payoff from taking the flight.

Finally, note that, under the proposed restriction, the signals purchasing different fare classes would be determined simply according to an ordered partition. For instance, with two fare classes, we would now anticipate two thresholds, θ_1^* and θ_2^*, satisfying $0 \leq \theta_1^* < \theta_2^*$. Passenger i would not purchase a ticket if $\theta_i < \theta_1^*$, would purchase a relatively cheap fare with little compensation if $\theta_i \geq \theta_2^*$, and would purchase a more expensive fare with high compensation in case $\theta_i \in [\theta_1^*, \theta_2^*)$. Overbooking would then occur whenever the number of passengers purchasing airfares, i.e., the number with $\theta_i \geq \theta_1^*$, exceeds the number of available seats m.

Similar principles to those we have seen would then apply to setting the prices and compensation policies for the various fare classes. For instance, the airline could lower the probability of being seated for a passenger with the less expensive fare class (i.e., with $\theta_i \geq \theta_2^*$) and in doing so raise compensation to an unseated ticket holder. This would raise the price the airline could charge for this fare class and reduce the rents earned by infra-

v_i in the sense of first-order stochastic dominance).

[28]Naturally, this conclusion reflects that the distribution over values v_i increases with θ_i in the sense of first-order stochastic dominance. Our focus on this case simplifies the analysis, but a natural alternative may be where θ_i captures passengers' degree of certainty regarding future values (for example, a second-order stochastic dominance relation, as analyzed in Courty and Li (2000)).

marginal holders of tickets in this fare class (i.e., those with $\theta_i > \theta_2^*$). Again, the reason would be that infra-marginal passengers are less likely to benefit from the compensation. Similarly, the airline could lower the probability of being seated for a passenger with the more expensive fare (i.e., those with $\theta_i \in [\theta_1^*, \theta_2^*)$). This would increase the compensation available for this fare and permit an increase in prices of both fares (since the option of type θ_2^* to purchase the more expensive fare class becomes less appealing).

6 Conclusions

This paper proposed an incentive-based rationale for overbooking and shed light on a novel interplay between screening ex-ante via ticket prices and screening ex-post through auctions. Compensation offered to ticketed passengers in return for their seats helps improve demand for tickets, allowing the airline to raise ticket prices. At the same time, passengers who strongly believe they will have a high value for flying do not expect to benefit from such compensation. These passengers, who are infra-marginal when it comes to buying tickets, pay the higher ticket prices but expect to benefit little from compensation, so they expect less surplus. Repurchased tickets may or may not be sold to unticketed passengers. Selling seats to unticketed passengers can improve efficiency, but it also improves the outside option of passengers who may consider purchasing tickets in advance of the flight, lowering the demand for tickets.

To highlight our main ideas, we have abstracted from several relevant features of airline markets. As noted, we focused for the most part on a single fare class, although most airlines offer several. We also considered a monopolist airline, although overbooking is likely to have implications for competition. A passenger who has purchased a ticket from one airline will find it costly to then purchase a seat from a competitor unless she expects a full refund for the original ticket price. In turn, competition would be likely to influence airlines' choices of compensation policies. For instance, an airline may wish to dissuade a ticketed passenger from giving up his seat at the last minute and flying with a cheaper competitor.

We assumed that passengers are forward-looking and strategic and abstracted from possible behavioral biases. For instance, the extent to which passengers internalize possible compensation when considering ticket purchases clearly has important implications for our theory. Passengers' perceptions about likely compensation may not depend simply on contractually agreed terms, but also on airlines' marketing policies.

A related concern from which we abstracted in the analysis is transaction costs. In practice, overbooking is usually resolved at the airport, with passengers volunteering their seats right before the flight (in response to appropriate compensation). However, turning up to a flight on time and then not taking it (typically delaying travel until a later time) is costly for the passenger. This concern may well limit airlines' use of overbooking at present, but this could conceivably change with the evolution of information technology. For instance, airlines could adopt technology to solicit volunteers to satisfy capacity constraints before passengers incur the cost of visiting the airport on time for the flight. Our analysis suggests how this could be profitable for the airline, and what trade-offs it may face in implementing such a scheme. While increases in overbooking might seem a concern for regulators, we have illustrated how overbooking may enhance efficiency. The reason is that the airline can extract passenger rents through early ticket sales rather than charging high prices to unticketed passengers, generating inefficient exclusion.

Finally, it is worth pointing out that overbooking arises in settings other than airline markets. Hotels, restaurants, and car rental agencies are also known to overbook, although mechanisms for eliciting customer preferences and compensating them in the event of overbooking seem much less developed. One key difference with airlines may be that there is a predetermined flight date, so all passengers who are ticketed for the flight gather in the same place. Soliciting volunteers according to willingness to pay then becomes a natural possibility, and one which we observe in practice.

Overbooking can also be understood as occurring in some labor market settings, although it does not go by that name.[29] For instance, tenure-track professors may compete for a limited number of tenured positions, while junior employees in professional services firms compete for a limited number of senior posts (under what are often termed "up or out" contracts). In fact, our framework might prove readily adaptable to these labor market settings, given that employees may learn over time about their ability to deliver high quality outputs to the employer. For instance, by hiring more junior staff members than there are available senior staff positions, and by foregoing hires of more seasoned workers from outside the company, a firm might reduce the information rents associated with workers' private understandings of their own abilities.

References

Baron, D. P., and D. Besanko (1984): "Regulation and information in a continuing relationship," *Information Economics and Policy*, 1(3), 267–302.

Battaglini, M. (2005): "Long-Term Contracting with Markovian Consumers," *American Economic Review*, 3, 637–658.

[29] We are grateful to Makoto Watanabe for sharing this idea and suggesting how our framework fits this application.

Beckmann, M. J. (1958): "Decision and Team Problems in Airline Reservations," *Econometrica*, 26(1), 134–145.

Bergemann, D., and M. Said (2011): "Dynamic Auctions," *Wiley Encyclopedia of Operations Research and Management Science.*

Biehl, A. R. (2001): "Durable-Goods Monopoly with Stochastic Values," *The RAND Journal of Economics*, 32(3), 565–577.

Board, S., and A. Skrzypacz (2015): "Revenue Management with Forward-Looking Buyers," *Journal of Political Economy, forthcoming.*

Courty, P. (2003): "Ticket Pricing under Demand Uncertainty," *The Journal of Law and Economics*, 46(2), 627–652.

Courty, P., and H. Li (2000): "Sequential Screening," *The Review of Economic Studies*, 67(4), 697–717.

Dana, J. D. (1998): "Advance-Purchase Discounts and Price Discrimination in Competitive Markets," *Journal of Political Economy*, 106(2), 395–422.

Deb, R. (2014): "Intertemporal Price Discrimination with Stochastic Values," .

Deb, R., and M. Said (2015): "Dynamic screening with limited commitment," *Journal of Economic Theory*, 159, 891–928.

DeGraba, P. (1995): "Buying Frenzies and Seller-Induced Excess Demand," *The RAND Journal of Economics*, 26(2), 331–342.

Eso, P., and B. Szentes (2007): "Optimal Information Disclosure in Auctions and the Handicap Auction," *The Review of Economic Studies*, 74(3), 705–731.

Fu, X., P. Gautier, and M. Watanabe (2012): "Overbooking in Markets with Rationing," .

Gallego, G., S. G. Kou, and R. Phillips (2008): "Revenue Management of Callable Products," *Management Science*, 54(3), 550–564.

Gallego, G., and O. Sahin (2010): "Revenue Management with Partially Refundable Fares," *Operations Research*, 58(4-Part-1), 817–833.

Garrett, D. (2014): "Dynamic Mechanism Design: Dynamic Arrivals and Changing Values," .

——— (2015): "Intertemporal price discrimination: dynamic arrivals and changing values," .

Gershkov, A., J. K. Goeree, A. Kushnir, B. Moldovanu, and X. Shi (2012): "On the equivalence of Bayesian and dominant strategy implementation," *Econometrica.*

Gershkov, A., and B. Moldovanu (2009): "Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach," *American Economic Journal: Microeconomics*, 1(2), 168–198.

——— (2012): "Dynamic allocation and pricing: A mechanism design approach," *International Journal of Industrial Organization*, 30(3), 283–286.

Hinnosaar, T. (2015): "Calendar Mechanisms," .

Lazarev, J. (2012): "The Welfare Effects of Intertemporal Price Discrimination: An Empirical Analysis of Airline Pricing in U.S. Monopoly Markets," .

McAfee, R. P., and V. te Velde (2006): "Dynamic pricing in the airline industry," in *Handbook on Economics and Information Systems*, ed. by T. J. Hendershott. Elsevier, Amsterdam.

Mierendorff, K. (2016): "Optimal dynamic mechanism design with deadlines," *Journal of Economic Theory*, 161, 190–222.

Möller, M., and M. Watanabe (2010): "Advance Purchase Discounts Versus Clearance Sales," *The Economic Journal*, 120(547), 1125–1148.

Myerson, R. B., and M. A. Satterthwaite (1983): "Efficient mechanisms for bilateral trading," *Journal of Economic Theory*, 29(2), 265–281.

Nocke, V., M. Peitz, and F. Rosar (2011): "Advance-purchase discounts as a price discrimination device," *Journal of Economic Theory*, 146(1), 141–162.

Pai, M. M., and R. Vohra (2013): "Optimal Dynamic Auctions and Simple Index Rules," *Mathematics of Operations Research*, 38(4), 682–697.

Pavan, A., I. Segal, and J. Toikka (2013): "Dynamic Mechanism Design: A Myersonian Approach," .

——— (2014): "Dynamic Mechanism Design: A Myersonian Approach," *Econometrica*, 82(2), 601–653.

Rothstein, M. (1971): "An Airline Overbooking Model," *Transportation Science*, 5(2), 180–192.

Said, M. (2012): "Auctions with dynamic populations: Efficiency and revenue maximization," *Journal of Economic Theory*, 147(6), 2419–2438.

Shugan, S. M., and J. Xie (2000): "Advance Pricing of Services and Other Implications of Separating Purchase and Consumption," *Journal of Service Research*, 2(3), 227–239.

Simon, J. L. (1994): "Origins of the Airline Oversales Auction System," *Regulation*, 2, 48–52.

Tirole, J. (1988): *The Theory of Industrial Organization*. The MIT Press, 1st edn.

Vickrey, W. (1972): "Airline Overbooking: Some Further Solutions," *Journal of Transport Economics and Policy*, 6, 257–270.

Xie, J., and S. M. Shugan (2001): "Electronic Tickets, Smart Cards, and Online Prepayments: When and How to Advance Sell," *Marketing Science*, 20(3), 219–243.

A Proofs of results

This appendix collects proofs not given in the main text.

Proof of Lemma 1

Proof. Let Θ_T be the set of signals for which passengers purchase tickets, which may be any measurable subset of $[0,1]$. We need to show that there exists $\theta^* \in [0,1]$ such that $\Theta_T = [\theta^*, 1]$. First note that the stochastic process described by F and G necessarily admits an "independent-shock" representation as follows (see Eso and Szentes (2007) and Pavan, Segal, and Toikka (2014)). Define, for each $(\theta, \varepsilon) \in [0,1]^2$, $z(\theta,\varepsilon) = G^{-1}(\varepsilon|\theta)$. That is, for each $(\theta,\varepsilon) \in [0,1]^2$, $z(\theta,\varepsilon)$ is the unique value in the support of $G(\cdot|\theta)$ satisfying $G(z(\theta,\varepsilon)|\theta) = \varepsilon$. Suppose that $\tilde{\varepsilon}$ is uniformly distributed on $[0,1]$. Then, for each $\theta \in [0,1]$, $z(\theta,\tilde{\varepsilon})$ is distributed according to $G(\cdot|\theta)$. Without loss of generality, we focus on direct mechanisms at date 1, which induce truthtelling irrespective of the decision to buy a ticket at date 1. Let $\overline{q}^s(v_j)$ be the probability a ticketed passenger j is seated, conditional on s tickets being sold and his reported value being v_j; and let $\underline{q}^s(v_k)$ be the probability of being seated for an unticketed passenger k. Let $\overline{V}^s(\theta_j)$ be the expected payoff (gross of the ticket price) to a ticketed passenger conditional on s tickets being purchased and let $\underline{V}^s(\theta_j)$ be the expected payoff to an unticketed passenger. Let $\tilde{r}$ be distributed according to a binomial distribution with parameters $(n-1, 1-F(\Theta_T))$, where $F(\Theta_T)$ is the probability that any other passenger purchases a ticket. Applying Theorem 1 of Pavan, Segal, and Toikka (2014), we have that the expected payoff for a passenger who purchases a ticket satisfies, for any signals $\theta', \theta'' \in [0,1]$ with $\theta' < \theta''$,

$$\mathbb{E}_{\tilde{r}}\left[\overline{V}^{\tilde{r}+1}(\theta'')\right] = \mathbb{E}_{\tilde{r}}\left[\overline{V}^{\tilde{r}+1}(\theta')\right] + \int_{\theta'}^{\theta''} \mathbb{E}_{(\tilde{r},\tilde{\varepsilon})}\left[\frac{\partial z(y,\tilde{\varepsilon})}{\partial \theta}\overline{q}^{\tilde{r}+1}(z(y,\tilde{\varepsilon}))\right] dy. \tag{28}$$

The expected payoff for a passenger who does not purchase a ticket satisfies

$$\mathbb{E}_{\tilde{r}}\left[\underline{V}^{\tilde{r}}(\theta'')\right] = \mathbb{E}_{\tilde{r}}\left[\underline{V}^{\tilde{r}}(\theta')\right] + \int_{\theta'}^{\theta''} \mathbb{E}_{(\tilde{r},\tilde{\varepsilon})}\left[\frac{\partial z(y,\tilde{\varepsilon})}{\partial \theta}\underline{q}^{\tilde{r}}(z(y,\tilde{\varepsilon}))\right] dy. \tag{29}$$

That higher signals imply higher values in the sense of first-order stochastic dominance means that, for all (θ,ε), $\frac{\partial z(\theta,\varepsilon)}{\partial \theta} \geq 0$. Our assumption that ticket holders are favored implies that, for all $s \in \{0,\ldots,n-1\}$ and all $v_i \in [\underline{v},\overline{v}]$, $\underline{q}^s(v_i) \leq \overline{q}^{s+1}(v_i)$. Hence, for θ', θ'' with $\theta' < \theta''$,

$$\int_{\theta'}^{\theta''} \mathbb{E}_{(\tilde{r},\tilde{\varepsilon})}\left[\frac{\partial z(y,\tilde{\varepsilon})}{\partial \theta}\overline{q}^{\tilde{r}+1}(z(y,\tilde{\varepsilon}))\right] dy \geq \int_{\theta'}^{\theta''} \mathbb{E}_{(\tilde{r},\tilde{\varepsilon})}\left[\frac{\partial z(y,\tilde{\varepsilon})}{\partial \theta}\underline{q}^{\tilde{r}}(z(y,\tilde{\varepsilon}))\right] dy. \tag{30}$$

Now, suppose p is the price of the ticket and that

$$\mathbb{E}_{\tilde{r}}\left[\overline{V}^{\tilde{r}+1}(\theta')\right] - p \geq \mathbb{E}_{\tilde{r}}\left[\underline{V}^{\tilde{r}}(\theta')\right] \tag{31}$$

so that a passenger with signal θ' does better by purchasing the ticket at price p than by not purchasing. Then Equations (28), (29), and (30) imply

$$\mathbb{E}_{\tilde{r}}\left[\overline{V}^{\tilde{r}+1}(\theta'')\right] - p \geq \mathbb{E}_{\tilde{r}}\left[\underline{V}^{\tilde{r}}(\theta'')\right] \tag{32}$$

so that θ'' also prefer does better to purchase the ticket at price p. Hence, given our assumption that the passenger purchases a ticket whenever willing, if θ' purchases, then so does θ''. This establishes that ticket sales have a threshold property: the fact that this set is closed—i.e., it includes the threshold value θ^*—follows from the continuity of $\mathbb{E}_{\tilde{r}}\left[\overline{V}^{\tilde{r}+1}(\cdot)\right]$ and $\mathbb{E}_{\tilde{r}}\left[\underline{V}^{\tilde{r}}(\cdot)\right]$, which is immediate from Equations (28) and (29). □

Proof of Lemma 2

Proof. The proof is a minor adaptation of Gershkov, Goeree, Kushnir, Moldovanu, and Shi (2012), and here we provide only a sketch. Let $\chi_i = 1$ if i holds a ticket (i.e., $i \leq s$)

and $\chi_i = 0$ otherwise (i.e., if $i > s$).[30] We let $(\hat{\mathbf{q}}, \hat{\mathbf{t}})$ be a Bayesian incentive-compatible mechanism and consider the set C of alternative allocation rules $\mathbf{q}$ satisfying

$$q_i^s(\mathbf{v}) \geq 0 \quad \forall \mathbf{v}, s, i$$

$$\sum_i q_i^s(\mathbf{v}) \leq m \quad \forall \mathbf{v}, s$$

$$\mathbb{E}\left[q_i^{\tilde{s}}(v_i, \tilde{\mathbf{v}}_{-i})\chi_i\right] = \mathbb{E}\left[\hat{q}_i^{\tilde{s}}(v_i, \tilde{\mathbf{v}}_{-i})\chi_i\right] \quad \forall v_i, \chi_i \tag{33}$$

$$\mathbb{E}\left[\mathbf{q}^{\tilde{s}}(\tilde{\mathbf{v}})\right] = \mathbb{E}\left[\hat{\mathbf{q}}^{\tilde{s}}(\tilde{\mathbf{v}})\right]. \tag{34}$$

Let $\mathbf{q}^*$ solve

$$\min_{\mathbf{q} \in C} \mathbb{E}\left[\|\mathbf{q}^{\tilde{s}}(\tilde{\mathbf{v}})\|^2\right]. \tag{35}$$

By following exactly the same steps as in Gershkov, Goeree, Kushnir, Moldovanu, and Shi (2012), one can show the existence of $\mathbf{q}^*$ which satisfies, for each i,

$$q_i^{s*}(\cdot, \mathbf{v}_{-i}) \quad \text{is non-decreasing for all } s \text{ and } \mathbf{v}_{-i}.$$

That is, the probability with which i is awarded a seat is ex-post monotone in v_i. By standard arguments, this implies that $\mathbf{q}^*$ is implementable in dominant strategies. Eq. (33) implies $\mathbf{q}^*$ yields the same interim expected utilities as $\mathbf{q}'$, and Eq. (34) is used to show that expected revenues are also the same (again following the steps in Gershkov, Goeree, Kushnir, Moldovanu, and Shi (2012)). □

Proof of Lemma 3

Proof. The passenger with the marginal signal is indifferent between purchasing and not. Therefore, the expected payoff $\mathbb{E}[\overline{V}^{\tilde{r}+1}(\theta^*)] - p$ must equal $\mathbb{E}[\underline{V}^{\tilde{r}}(\theta^*)]$, where the number of other ticket holders $\tilde{r}$ has binomial distribution with parameters $(n-1, 1-F(\theta^*))$. This gives

$$p = \sum_{r=0}^{n-1} \binom{n-1}{r} (1 - F(\theta^*))^r F(\theta^*)^{n-1-r} [\overline{V}^{r+1}(\theta^*) - \underline{V}^r(\theta^*)]. \tag{36}$$

[30] Clearly, each passenger knows whether he holds a ticket, i.e., whether he is ordered below or above s, but faces a distribution over s conditional on whether he holds a ticket.

Expected total profit is $\pi = n(1-F(\theta^*))p + \mathbb{E}[\tilde{s}\overline{\Pi}^{\tilde{s}} + (n-\tilde{s})\underline{\Pi}^{\tilde{s}}]$, where $\tilde{s}$ has binomial distribution with parameters $(n, 1-F(\theta^*))$. Therefore,

$$\pi = n(1-F(\theta^*))p + \sum_{s=0}^{n} \binom{n}{s} (1-F(\theta^*))^s F(\theta^*)^{n-s} [s\overline{\Pi}^s + (n-s)\underline{\Pi}^s]. \tag{37}$$

From Eq. (5), $n(1-F(\theta^*))p = n(1-F(\theta^*))\mathbb{E}[\overline{V}^{\tilde{r}+1}(\theta^*)] - n(1-F(\theta^*))\mathbb{E}[\underline{V}^{\tilde{r}}(\theta^*)]$. Rewriting the two parts of this expression,

$$n(1-F(\theta^*))\mathbb{E}[\overline{V}^{\tilde{r}+1}(\theta^*)] = \sum_{r=0}^{n} \frac{n!}{r!(n-r)!}(1-F(\theta^*))^r F(\theta^*)^{n-r} r \overline{V}^r(\theta^*), \text{ and}$$

$$n(1-F(\theta^*))\mathbb{E}[\underline{V}^{\tilde{r}}(\theta^*)] = \frac{1-F(\theta^*)}{F(\theta^*)} \sum_{r=0}^{n} \frac{n!}{r!(n-r)!}(n-r)(1-F(\theta^*))^r F(\theta^*)^{n-r} \underline{V}^r(\theta^*).$$

Inserting $n(1-F(\theta^*))\mathbb{E}[\overline{V}^{\tilde{r}+1}(\theta^*)]$ and $n(1-F(\theta^*))\mathbb{E}[\underline{V}^{\tilde{r}}(\theta^*)]$ into the profit equation gives Eq. (8). □

Proof of Lemma 4

Proof. We consider a direct mechanism without any restrictions. Here, with an abuse of notation, $q_i(\theta_i, v_i; \theta_{-i}, v_{-i})$ specifies, for each passenger i, the probability i is seated given the sequence of own reports (θ_i, v_i) and the reports of the other passenger (θ_{-i}, v_{-i}). Let $x_i(\theta_i, v_i) = \mathbb{E}_{(\tilde{\theta}_{-i}, \tilde{v}_{-i})}\left[q_i\left(\theta_i, v_i; \tilde{\theta}_{-i}, \tilde{v}_{-i}\right)\right]$. Let $W(\theta)$ be the expected payoff to a passenger with signal θ who reports optimally in the direct mechanism (and whose only information is his own type θ and his initial prior belief about the other passenger's type), and $W(\theta, v)$ the expected payoff when he receives signal θ and his value is v (again, given that he does not update his beliefs about the other passenger's information). First, note that, because the passenger with signal $\overline{\theta}$ can follow a strategy of reporting $\underline{\theta}$ and then $\underline{v}$, we must have

$$W(\overline{\theta}) \geq W(\underline{\theta}) + x_i(\underline{\theta}, \underline{v}) \left(1 - \frac{1}{2}\right)(\overline{v} - \underline{v}). \tag{38}$$

Because the passenger with signal $\underline{\theta}$ can follow a strategy of waiting and reporting only at date 1, and at that point reporting value $\underline{v}$, we have

$$W(\underline{\theta}) \geq W(\varnothing, \underline{v}) + x_i(\varnothing, \underline{v}) \frac{1}{2}(\overline{v} - \underline{v}). \tag{39}$$

Because a passenger who arrives at date 1 and has value $\overline{v}$ can report $\underline{v}$, we must have[31]

$$W(\varnothing,\overline{v}) \geq W(\varnothing,\underline{v}) + x_i(\varnothing,\underline{v})(\overline{v}-\underline{v}). \tag{40}$$

As a result, using that $W(\varnothing,\underline{v}) \geq 0$, the expected contribution of passenger i to the airline's profit is at most $\mathbb{E}\left[\text{VS}\left(\tilde{\theta},\tilde{v}\right) x_i\left(\tilde{\theta},\tilde{v}\right)\right]$, where $\text{VS}(\underline{\theta},\underline{v}) = \underline{v} - (\overline{v}-\underline{v})$, $\text{VS}(\varnothing,\underline{v}) = \underline{v} - 3(\overline{v}-\underline{v})$, and $\text{VS}(\theta,\overline{v}) = \overline{v}$ for all θ. The allocation proposed in the text maximizes the sum of expected virtual surpluses VS across both passengers, and thereby attains the maximum possible expected profit for the airline in an incentive-compatible mechanism. It remains to consider the ticket price and date-1 payments which implement this allocation. These are chosen so that, on path, an unticketed low-value passenger i pays $q_i(\varnothing,\underline{v};\theta_{-i},v_{-i})\,\underline{v}$ and a high-value passenger $q_i(\varnothing,\overline{v};\theta_{-i},v_{-i})\,\overline{v} - q_i(\varnothing,\underline{v};\theta_{-i},v_{-i})(\overline{v}-\underline{v})$ for the possibility of being seated. Also, a ticketed passenger i receives payments at date 1 equal to $(1-q_i(\theta_i,v_i;\theta_{-i},v_{-i}))\,\overline{v}$. In particular, a ticketed passenger receives compensation equal to $\overline{v}$ times the probability he loses his seat, conditional on reported values. Via Eq. (5), this pins down the ticket price:

$$\overline{v} - x_i(\underline{\theta},\underline{v})\frac{1}{2}(\overline{v}-\underline{v}) - x_i(\varnothing,\underline{v})\frac{1}{2}(\overline{v}-\underline{v}). \tag{41}$$

Given this ticket price and transfers, the only incentive constraint that requires some consideration is that $\overline{\theta}$ wishes to purchase a ticket at date 0. Note that, by construction of ticket prices and compensation, the expected payoff for $\overline{\theta}$ by obediently purchasing a ticket and then reporting truthfully is

$$x_i(\varnothing,\underline{v})\frac{1}{2}(\overline{v}-\underline{v}) + x_i(\underline{\theta},\underline{v})\frac{1}{2}(\overline{v}-\underline{v}). \tag{42}$$

The deviation is unprofitable if and only if this is at least what can be obtained by not purchasing a ticket, i.e., $x_i(\varnothing,\underline{v})(\overline{v}-\underline{v})$. Therefore, purchasing the ticket is incentive compatible if and only if $x_i(\underline{\theta},\underline{v}) \geq x_i(\varnothing,\underline{v})$. This is guaranteed by our assumptions that $f(\underline{\theta}) = g(\varnothing) = 1/3$ and $g(\overline{v}|\underline{\theta}) = g(\overline{v}|\varnothing) = 1/2$. □

Proof of Proposition 2

Proof. The proof proceeds by specifying an allocation $\mathbf{q}$ which maximizes the virtual surplus Eq. (23) and which is implementable by a pricing mechanism which satisfies the restrictions set out in Section 3. For reported values consistent with equilibrium

[31] We focus initially on interim incentive constraints, and then verify that the proposed optimal mechanism respects ex-post incentive constraints.

ticket purchasing, the allocation is determined simply through a comparison of the virtual surpluses. What must still be specified is the allocation when one of the passengers deviates by reporting a value inconsistent with his decision to purchase or not to purchase a ticket (the allocation in case two or more passengers make such reports is irrelevant because we need not consider joint deviations from equilibrium play).

1. First, for any s and for any ticket holder j reporting v_j above the support of $G(\cdot|\tilde{\theta} \geq \theta^*)$, and for any $\mathbf{v}_{-j}$, $q_j^s(v_j, \mathbf{v}_{-j}) = 1$ while $q_i^s(v_j, \mathbf{v}_{-j}) = 0$ for all $i \neq j$.

2. Second, for any s and for any ticket holder j reporting v_j below the support of $G(\cdot|\tilde{\theta} \geq \theta^*)$, and for any $\mathbf{v}_{-j}$, the allocation is the same as for passenger j reporting the minimal value of the support of $G(\cdot|\tilde{\theta} \geq \theta^*)$.

3. Third, for any s and for any non-ticket holder k reporting v_k above the support of $G(\cdot|\tilde{\theta} < \theta^*)$, and for any $\mathbf{v}_{-k}$, the allocation is the same as in case k reports the maximum of the support of $G(\cdot|\tilde{\theta} < \theta^*)$.

4. Finally, for any s and for any non-ticket holder k reporting v_k below the support of $G(\cdot|\tilde{\theta} < \theta^*)$, and for any $\mathbf{v}_{-k}$, $q_j^s(v_k, \mathbf{v}_{-k}) = 0$ and $q_i^s(v_k, \mathbf{v}_{-k}) = 0$ for all $i \neq k$.

Using Condition (i) of the proposition, the allocations defined above satisfy, for any s, any passenger i (whether or not he has purchased a ticket), and any $\mathbf{v}_{-i}$ consistent with equilibrium ticket purchasing, $q_i^s(\cdot, \mathbf{v}_{-i})$ is non-decreasing over $[\underline{v}, \overline{v}]$. This guarantees that the period-1 mechanism, defined by the above allocations and the transfers Eq. (15), is ex-post incentive compatible irrespective of whether the passenger follows the equilibrium strategy in purchasing (or not purchasing) a ticket. Second, Condition (ii), together with the allocations defined above for out-of-equilibrium play, ensures that the allocation rule gives precedence to ticket holders. Since the price satisfying Eq. (5) keeps the passenger with signal θ^* indifferent between purchasing and not purchasing a ticket, the same argument as in the proof of Lemma 1 implies that all passengers with signals above θ^* are willing to purchase tickets, while passengers with signals less than θ^* are not. □

Proof of Corollary 1

Proof. Let $\mathbf{q}^*$ be a symmetric allocation specifying, for each passenger i, each $\mathbf{v}$, a probability of flying $q_i^*(\mathbf{v})$ and maximizing the airline's profit in a mechanism *without* ticket sales (i.e., where passengers contract only at date 1). Pick $\theta^* \in [0, 1)$ satisfying the conditions in the Corollary. Consider now the pricing mechanism which is optimal conditional on (i) selling tickets to passengers if and only if their signals exceed θ^*, and (ii) implementing the symmetric allocation rule $\mathbf{q}^*$ (which depends only on values reported at date

1 and not on which passengers hold tickets). From Eq. (28) and Eq. (29), we see that, given the passenger with the marginal signal θ^* is indifferent between purchasing and not purchasing a ticket, the passengers with all other signals in $[0, 1]$ must be indifferent as well. Hence, the optimal pricing mechanism with threshold θ^* and allocation $\mathbf{q}^*$ and the mechanism without tickets both provide the same expected payoff to passengers. Since total welfare is also the same for each mechanism, total profit must be identical as well. Note, however, that by Condition (ii) of the Corollary, the allocation $\mathbf{q}^*$ fails to maximize the expression in Eq. (23). Hence, both the pricing mechanism that implements this allocation and the optimal mechanism without tickets deliver a strictly lower profit than the one achievable by selling tickets to passengers with signals above θ^* and by implementing an allocation which optimally distinguishes ticketed from unticketed passengers (as described in the proof of Proposition 2). □

Proof of Example 1

Proof. This follows from noticing that, for any $\theta^* \in [0, 1]$, the expected virtual surplus Eq. (23) can be maximized by an allocation rule which is symmetric and independent of which passengers are ticketed. The optimal mechanism without ticket sales, where passengers contract only date 1, gives passengers the same surplus as in the optimal pricing mechanism with ticket sales to passengers with signals above θ^*. Hence it delivers the airline the same profit. □

Proof of Example 2

Proof.

1. Fix $\theta^* > 0$ and call the corresponding pricing mechanism "Mechanism A". There exists a threshold $v^\#$, such that, in the optimal pricing mechanism conditional on θ^*, if all n passengers purchase tickets, a passenger j flies if and only if he is among the m highest values, and his value v_j exceeds $v^\#$ (note that $\overline{VS}\left(v^\#\right) = 0$). Now define "Mechanism B" to be the alternative mechanism which sells tickets to all passengers, and then allocates a seat to a passenger if and only if he is among the m highest values, and his value v_j exceeds $v^\#$. Passengers who do not purchase a ticket do not fly, and the price of the ticket is set to ensure that the passenger with the lowest date-0 signal expects a payoff zero from purchasing a ticket (and hence is indifferent between purchasing and not).

   Now, one can verify that the additional expected rent in Mechanism B is larger than in Mechanism A by no more than $n\theta^*$. On the other hand, whenever a is sufficiently large, expected welfare losses in Mechanism A exceed those in Mechanism B by an amount exceeding $n\theta^*$. To see this, note that, under Mechanism B, a passenger is seated unless possibly (i) $v_j - 1$ is less than the m^{th} highest realized value of the other passengers, or (ii) $v_j \leq 1$. However, under Mechanism A, there exists $a^{\#}$ such that the following holds. For all $a \geq a^{\#}$, if a passenger k is unticketed, he is not seated with probability 1 whenever $v_k < \frac{a}{2}$. The corresponding efficiency loss then becomes unboundedly large with a.

2. Given the absence of capacity constraints, we can focus on the case with $n = 1$. For each $\theta^* \in [0, 1)$, let $\overline{q}(v)$ and $\underline{q}(v)$ indicate whether the passenger is seated under the optimal allocation when his value is v, and when he is ticketed and unticketed, respectively. Following the same steps as in the derivation of the unrestricted mechanism in Appendix B, the airline's expected profit, conditional on θ^*, can be shown to equal

$$\mathbb{E}_{(\tilde{\theta},\tilde{\varepsilon})}\left[\begin{array}{c} \mathbf{1}_{\tilde{\theta}\geq\theta^*}\overline{q}(\tilde{\theta}+a\tilde{\varepsilon})\left(2\tilde{\theta}+a\tilde{\varepsilon}-1\right) \\ +\left(1-\mathbf{1}_{\tilde{\theta}\geq\theta^*}\right)\underline{q}(\tilde{\theta}+a\tilde{\varepsilon})\left(2\tilde{\theta}+a\tilde{\varepsilon}-1\right) \end{array}\right]. \tag{43}$$

The first-order necessary condition for an optimum is

$$\mathbb{E}_{\tilde{\varepsilon}}\left[-\overline{\text{VS}}\left(\theta^*+a\tilde{\varepsilon}\right)\overline{q}\left(\theta^*+a\tilde{\varepsilon}\right)+\underline{\text{VS}}\left(\theta^*+a\tilde{\varepsilon}\right)\underline{q}\left(\theta^*+a\tilde{\varepsilon}\right)\right]=0. \tag{44}$$

This yields $\theta^* = \frac{1}{2} - \frac{a}{4}$. It can further be verified that the expected profit in Eq. (43) is quasi-concave in θ^*, so that $\frac{1}{2} - \frac{a}{4}$ achieves the optimum.

3. First, suppose that $0 \leq \theta^* < 1 - \gamma - a$ and consider the optimal pricing mechanism conditional on θ^*. As $n \to \infty$, the probability that the flight is filled by ticket holders with values exceeding $1 - \gamma$ approaches 1 and hence the expected payoff of the passenger with date-0 signal equal to zero approaches zero. The same conclusion holds if instead the optimal date-1 spot mechanism is used (i.e., no tickets are sold), in which case the allocation is identical under the two mechanisms (i.e., the spot mechanism and the pricing mechanism with threshold for ticket purchases θ^*) with a probability that approaches 1 as $n \to \infty$.[32] This implies that the airline's expected profits per passenger converge under the two mechanisms.

Conversely, suppose $\theta^* > 1 - \gamma + a$, which we can take to be less than 1 (hence, $a < 1/2$). Under the optimal pricing mechanism, a passenger whose signal is at least

[32]Under the pricing mechanism, the m ticketed passengers with the highest values are all seated with probability approaching 1. With probability approaching 1, this set of passengers is also the set with the highest overall values (including ticketed and unticketed passengers). Such passengers receive seats in the optimal spot mechanism which ranks passengers in terms of the virtual surplus in the usual Myersonian static auction.

$(1-\gamma+a+\theta^*)/2$ is seated with a probability that approaches 1 as $n \to \infty$. The same is again true under the optimal spot mechanism, and the allocation rules under the two mechanisms coincide with a probability that approaches 1 as $n \to \infty$. Again it holds that expected profits per passenger converge under the two mechanisms as $n \to \infty$.

Now suppose instead that $\theta^* = 1 - \gamma$, for example. Then, ticketing affects the allocation of seats with a probability that remains bounded above zero as $n \to \infty$, and the additional profits per passenger relative to the optimal spot mechanism remains bounded above zero as well. This can be shown using the same logic as in Corollary 1.

□

Proof of Proposition 3

Proof. Consider the following strategy profile. Ticketed passengers with value v_j bid v_j. Unticketed passengers with value v_k bid $v_j(v_k)$. To show that this is a dominant-strategy equilibrium, first consider a ticketed passenger with value v_j when the profile of other bids is b_{-j}. Among those bids that exceed R, let b^q_{-j} denote the qth highest bid, set equal to R in case the number of such bids is smaller than n. If j's bid b_j exceeds b^q_{-j}, then j will be seated and receive no transfer. If he bids less than b^q_{-j}, he will volunteer his seat and receive compensation in the amount of $b^q = b^q_{-j}$. Passenger j wishes to volunteer his seat at this price if and only if $v_j \leq b^q_{-j}$. Thus by bidding according to the specified strategy profile, i.e., $b_j = v_j$, he volunteers his seat exactly when it is optimal to do so. Next consider an unticketed passenger with value v_k. Let b^q_{-k} denote the qth highest bid among the other passengers whose bid exceeds R, with b^q_{-k} equaling R if there are fewer than q. If k's bid b_k exceeds b^q_{-k}, then k will win a seat and pay $v_k(b^q_{-k})$. If he bids less than b^q_{-k}, he will lose and pay nothing. It is in the interest of passenger k to fly and pay $v_k(b^q_{-k})$ if and only if $v_k \geq v_k(b^q_{-k})$. By bidding $b_k = v_j(v_k)$ as dictated by the strategy profile, he flies if and only if $v_j(v_k) \geq b^q_{-k}$, i.e., if and only if $v_k \geq v_k(b^q_{-k})$, which is exactly when it is in his interest to do so. Finally, to show that this dominant-strategy equilibrium implements the optimal allocation of seats, it is enough to notice that the ranking of bids is the same as the ranking of virtual surpluses (with bids below the reserve R being only for values such that virtual surplus is negative). This follows directly from the construction, using Condition (i) of Proposition 2, which ensures virtual surpluses are monotone in values over the relevant range. □

B Unrestricted mechanism and multiple fare classes

In this Appendix we derive the unrestricted optimal mechanism and show how our analysis can be extended to multiple fare classes.

Unrestricted mechanism

Suppose for simplicity that $G(\cdot|\theta)$ has full support on $[\underline{v}, \overline{v}]$, for all signals θ (the arguments below, extend more generally, however). With no restriction on the space of possible mechanisms, the airline finds it optimal to contract with all passengers arriving in the market at date 0. Without loss of optimality, we consider direct mechanisms where each passenger i reports his initial signal θ_i in case he arrives at date 0, and then reports his value v_i in the second period. We can think of the report of the signal as the passenger choosing among the multiple fare classes available at date 0. We denote the null report, where the passenger fails to report at date 0, by $\varnothing$.

Let Ω_M denote a direct mechanism, comprising allocations $\mathbf{q} = (q_i)_{i=1}^n$ and transfers $\mathbf{t} = (t_i)_{i=1}^n$ such that, for all $\theta = (\theta_i)_{i=1}^n$ and $\mathbf{v}$, the probability each passenger is seated is given by $q_i(\theta, \mathbf{v})$ and the total payment by each payment over two periods is given by $t_i(\theta, \mathbf{v})$ (given the absence of discounting, the timing of payments does not affect payoffs; we could equivalently consider date-0 payments, i.e., prices for each fare class in a continuum, and compensation satisfying Proposition 1). The allocation of seats must satisfy the same feasibility constraint introduced above: no more passengers may be seated than the number of available seats.

Denote by $W^{\Omega_M}(\theta_i)$ the expected payoff of a passenger with signal θ_i given the opportunity to participate in the mechanism Ω_M at date 0 (but with no other information about his value or about the signals or values of other passengers). Recall the "independent-shock" representation in the proof of Lemma 1. By the envelope theorem (see Pavan, Segal, and Toikka (2013)), a necessary condition for the incentive compatibility of truthful reporting of signals is that for all $\theta_i \geq 0$,

$$W^{\Omega_M}(\theta_i) = W^{\Omega_M}(0) + \int_0^{\theta_i} \mathbb{E}\left[\frac{\partial z(y, \tilde{\varepsilon})}{\partial \theta_i} q_i\left(y, \tilde{\theta}_{-i}, z(y, \tilde{\varepsilon}_i), \tilde{\mathbf{v}}_{-i}\right)\right] dy. \tag{45}$$

We now conjecture that the value of $W^{\Omega_M}(0)$ is determined by the value a passenger with signal zero can obtain by deviating and reporting only at date 1. A passenger i arriving at date 1 and truthfully reporting his value v_i, given that other passengers report

θ_{-i} and $\mathbf{v}_{-i}$, earns a payoff

$$W^{\Omega_M}\left(\varnothing,\theta_{-i},v_i,\mathbf{v}_{-i}\right)=W^{\Omega_M}\left(\varnothing,\theta_{-i},\underline{v},\mathbf{v}_{-i}\right)+\int_{\underline{v}}^{v_i} q_i\left(\varnothing,\theta_{-i},y,\mathbf{v}_{-i}\right)dy. \tag{46}$$

We may optimally choose $W^{\Omega_M}\left(\varnothing,\theta_{-i},\underline{v},\mathbf{v}_{-i}\right)=0$ while ensuring the individual rationality of a passenger participating for the first time at date 1. A necessary condition for date-0 participation by a passenger with the 0 signal is therefore

$$\begin{aligned} W^{\Omega_M}(0) &\geq \mathbb{E}\left[\int_{\underline{v}}^{\tilde{v}_i} q_i\left(\varnothing,\tilde{\theta}_{-i},y,\tilde{\mathbf{v}}_{-i}\right)dy|\tilde{\theta}_i=0\right] \\ &= \mathbb{E}\left[\frac{1-G\left(\tilde{v}_i|0\right)}{g\left(\tilde{v}_i|0\right)}q_i\left(\varnothing,\tilde{\theta}_{-i},\tilde{v}_i,\tilde{\mathbf{v}}_{-i}\right)dy|\tilde{\theta}_i=0\right] \\ &= \mathbb{E}\left[\frac{1-G\left(\tilde{v}_i|0\right)}{g\left(\tilde{v}_i|\varnothing\right)}q_i\left(\varnothing,\tilde{\theta}_{-i},\tilde{v}_i,\tilde{\mathbf{v}}_{-i}\right)dy|\tilde{\theta}_i=\varnothing\right]. \end{aligned}$$

We conjecture that this is the only relevant participation constraint at date 0 and therefore that the inequality must bind in an optimal mechanism. Taking expectations and integrating by parts, the airline's expected profit is then equal to

$$\sum_{i=1}^{n}\left(\begin{array}{c}(1-f(\varnothing))\,\mathbb{E}\left[\begin{array}{c}\left(z\left(\tilde{\theta}_i,\tilde{\varepsilon}_i\right)-\frac{1-F\left(\tilde{\theta}_i\right)}{f\left(\tilde{\theta}_i\right)}\frac{\partial z\left(\tilde{\theta}_i,\tilde{\varepsilon}_i\right)}{\partial\theta_i}\right) \\ \times q_i\left(\tilde{\theta},z\left(\tilde{\theta}_i,\tilde{\varepsilon}_i\right),\tilde{\mathbf{v}}_{-i}\right)\end{array}\middle|\tilde{\theta}_i\geq 0\right] \\ +f(\varnothing)\,\mathbb{E}\left[\left(\tilde{v}_i-\frac{1-G(\tilde{v}_i|\varnothing)}{g(\tilde{v}_i|\varnothing)}-\frac{1-f(\varnothing)}{f(\varnothing)}\frac{1-G(\tilde{v}_i|0)}{g(\tilde{v}_i|\varnothing)}\right)q_i\left(\tilde{\theta},\tilde{v}_i,\tilde{\mathbf{v}}_{-i}\right)|\tilde{\theta}_i=\varnothing\right]\end{array}\right). \tag{47}$$

We can now conjecture an optimal allocation rule along the same lines as for the single fare class mechanisms. Let

$$\mathrm{VS}\left(\theta_i,\varepsilon_i\right)=\begin{cases} z\left(\theta_i,\varepsilon_i\right)-\frac{1-F(\theta_i)}{f(\theta_i)}\frac{\partial z(\theta_i,\varepsilon_i)}{\partial\theta_i} & \text{if } \theta_i\geq 0, \\ z\left(\theta_i,\varepsilon_i\right)-\frac{1-G(z(\theta_i,\varepsilon_i)|\varnothing)}{g(z(\theta_i,\varepsilon_i)|\varnothing)}-\frac{1-f(\varnothing)}{f(\varnothing)}\frac{1-G(z(\theta_i,\varepsilon_i)|0)}{g(z(\theta_i,\varepsilon_i)|\varnothing)} & \text{if } \theta_i=\varnothing \end{cases} \tag{48}$$

Let $\mathbf{q}^*$ be the allocation which maximizes the virtual surplus VS. The following result gives conditions under which this allocation is implementable by an appropriate system of transfers.

Proposition 4. *Suppose that* $\mathrm{VS}(\cdot,\cdot)$ *as defined by Eq. (48) is non-decreasing in both arguments. Then there exists a system of transfers* $\mathbf{t}^*$ *that implements the allocation* $\mathbf{q}^*$*. The mechanism* $(\mathbf{q}^*,\mathbf{t}^*)$ *is profit maximizing in the class of all possible mechanisms.*

Proof. The payoffs $W^{\Omega_M}(\theta,\mathbf{v})$ to be earned in equilibrium are pinned down by the allocation rule, the fact that the date-0 participation constraint binds at $\theta_i=0$, and that a passenger arriving at date 1 with value $\underline{v}$ expects zero surplus. We can define transfers simply through the identity $t_i(\theta, \mathbf{v}) = q_i(\theta, \mathbf{v})\,\theta_i - W^{\Omega_M}(\theta, \mathbf{v})$, so that passengers receive the intended payoffs provided they report truthfully. Using the fact each passenger's probability of receiving the seat is non-decreasing in his value for it, the mechanism described above can be shown to be ex-post incentive compatible at date 1. Incentive compatibility of truthful reporting of the signal at date 0, given that the passenger participates at that date, also follows from the monotonicity of the allocation in both signals and values (see Pavan, Segal, and Toikka (2013), for instance) and from the first-order stochastic dominance property of the stochastic process. What is left to check is that all passengers with signals in $[0, 1]$ prefer to participate at date 0 rather than delaying participation to date 1. This follows from essentially the same argument as in the proof of Lemma 1, after noting that, for each reported value of a passenger i (holding fixed the reports of the other passengers), the probability that the passenger receives the good, for each possible value v_i, is lower if he delays participation until date 1. □

The virtual surplus for the unrestricted mechanism is closely related to those for the single fare class mechanism studied above. The airline would like to distort downwards the probability that a passenger with a low signal is seated so as to reduce the rents that must be left to passengers with higher signals, and thereby dissuade them from mimicking the distribution of reports by passengers with the lower signals. This idea recalls other work on dynamic mechanism design. With the single fare class mechanism studied above, the airline lacks the flexibility to distinguish the allocations for passengers with different signals, but it faces an incentive to achieve a similar end by selling tickets only to passengers with high signals (conferring a higher probability of being seated on the holder).

The airline would also like to reduce the probability of flying for passengers who only arrive to the market at date 1 in order to reduce the rents available to them. The advantage in doing so is that it permits the airline to reduce the rent that must be left to passengers who arrive at date 0 in order to persuade them to participate at that date. This again parallels the airline's incentive to limit the rents of non-ticket holders in the single fare class mechanism.

The relationship between the virtual surplus $\mathrm{VS}(\cdot,\cdot)$ for the unrestricted mechanism and the virtual surpluses of the single fare class mechanism $\overline{\mathrm{VS}}(\cdot)$ and $\underline{\mathrm{VS}}(\cdot)$ for ticketed and unticketed passengers turns out to be simple. Given the threshold for ticket purchases θ^*, for each v_i, we have

$$\overline{\mathrm{VS}}(v_i) = \mathbb{E}\left[\mathrm{VS}\left(\tilde{\theta}_i, \tilde{v}_i\right) | \tilde{\theta}_i \geq \theta^*, \tilde{v}_i = v_i\right]$$

while

$$\underline{\mathrm{VS}}\left(v_i\right) = \mathbb{E}\left[\mathrm{VS}\left(\tilde{\theta}_i, \tilde{v}_i\right) | \tilde{\theta}_i < \theta^*, \tilde{v}_i = v_i\right].$$

Multiple fare classes

The above suggests how our analysis can be extended to allow the airline to use two or more fare classes. If L fare classes are permitted, then we consider pricing mechanisms with thresholds $\theta_1^*, \ldots, \theta_L^* \in [0,1]$, ordered from highest to lowest, such that all passengers with signals in $I_1 = [\theta_1^*, 1]$ purchase class-1 tickets, those with signals in $I_2 = [\theta_2^*, \theta_1^*)$ purchase class-2 tickets, and so forth, with passengers having signals in $I_{L+1} = (-\infty, \theta_L^*)$ not purchasing any ticket. Assuming that a passenger in a higher class is favored in the seating allocation in the sense introduced in Section 3, with all ticketed passengers favored over unticketed passengers, any pricing mechanism with L classes has this threshold property.

We can then conjecture an allocation of seats to passengers with the highest non-negative virtual surpluses, where these virtual surpluses are given, for each fare class l and value v_i by $\mathbb{E}\left[\mathrm{VS}\left(\tilde{\theta}_i, \tilde{v}_i\right) | \tilde{\theta}_i \in I_l, \tilde{v}_i = v_i\right]$. We can then look for conditions guaranteeing that these allocations are monotone in v_i for each fare class (and for unticketed passengers), and which favor passengers in a higher fare class conditional on having the same value v_i. Such conditions ensure the existence of ticket prices and a compensation rule which implements the proposed allocations.