

Analyzing Policy Risk and Accounting for Strategy: Auctions in the National Airspace System

James W. Bono
Department of Economics
American University

David H. Wolpert
NASA Ames Research Center

February 15, 2010

Abstract

We examine the potential for simple auction mechanisms to efficiently allocate arrival and departure slots during Ground Delay Programs (GDPs). The analysis is conducted using a new approach to predicting strategic behavior called Predictive Game Theory (PGT). The difference between PGT and the familiar Equilibrium Concept Approach (ECA) is that PGT models produce distribution-valued solution concepts rather than set-valued ones. The advantages of PGT over ECA in policy analysis and design are that PGT allows for decision-theoretic prediction and policy evaluation. Furthermore, PGT allows for a comprehensive account of risk, including two types of risk, systematic and modeling, that cannot be considered with the ECA. The results show that the second price auction dominates the first price auction in many decision-relevant categories, including higher expected efficiency, lower variance in efficiency, lower probability of significant efficiency loss and higher probability of significant efficiency gain. These findings are despite the fact that there is no *a priori* reason to expect the second price auction to be more efficient because none of the conventional reasons for preferring second price over first price auctions, i.e. dominant strategy implementability, apply to the GDP slot auction setting.

Keywords: auction, ground delay program, entropy, predictive game theory, strategic risk

1 Introduction

NASA and the FAA are exploring market mechanisms to increase the efficiency of the National Airspace System (NAS) so that it can accommodate the projected increase in demand over the next half century. Of particular interest is the efficient reassignment of takeoff and arrival slots for periods during which bad weather reduces NAS capacity. We introduce an auction mechanism (slot auction) that is designed to serve this purpose.

We would like to thank George Judge.

We analyze this auction using a distribution-valued solution concept, an approach broadly referred to as Predictive Game Theory (PGT) [see [Wolpert and Bono \(2009\)](#); [Bono and Wolpert \(2009\)](#)]. This is in contrast to the standard game theoretic approach, which uses set-valued equilibrium solutions to analyze strategic situations. We refer to this approach as the Equilibrium Concept Approach (ECA). Examples of the ECA are the well-known Nash equilibrium (NE), quantal response equilibrium (QRE) [see [McKelvey and Palfrey \(1995\)](#)] and epsilon equilibrium [see [Radner \(1980\)](#)]. The PGT approach is a non-equilibrium-based approach that formulates a probability distribution over all possible behaviors from available information. Because the output of a PGT model is a probability distribution, researchers, decision makers and stakeholders are capable of using the well-established theory of decision-making under uncertainty to make predictions about airline behavior and to choose among allocation mechanisms. With the PGT distribution over behavior, decisions are fully informed of all relevant risk information. In contrast, the ECA has none of these capabilities. This paper represents the first application of PGT to a real-world problem of policy analysis.

1.1 Roadmap of the Paper

The paper proceeds as follows. First, we introduce the problem of allocating arrival slots during GDPs. We discuss the current mechanism, Ration-by-Schedule (RBS), and introduce first and second price versions of a simple auction alternative to RBS.

Next, we provide a background on PGT. In particular, we point out the advantages of using a distribution-valued solution concept over the ECA for purposes of real world policy and risk analysis. Among the most important of these advantages are the ability to do mechanism design in a manner that is consistent with decision theory and the availability of the PGT mode as a universal refinement. Here we also include a brief summary of the Bayesian PGT model of this paper.

We next introduce our PGT model of GDP slot auctions. First we specify the profit functions of the airlines in terms of delay costs. Next, we introduce the QR-rationality likelihood and the entropic prior. Finally, we discuss our importance sampling procedure for estimating the posterior distribution.

In the next section we present our predictions for airline behavior, profits and efficiency gains. We demonstrate that the first and second price versions of the GDP slot auction obtain higher expected efficiency than the RBS allocation. Furthermore, we demonstrate that the second price auction dominates the first price auction in many decision-relevant categories, including higher expected efficiency, lower variance in efficiency, lower probability of significant efficiency loss and higher probability of significant efficiency gain. These findings are despite the fact that there is no *a priori* reason to expect the second price auction to be more efficient because none of the conventional reasons for preferring second price over first price auctions, i.e. dominant strategy implementability, apply to the GDP slot auction setting.

We conclude with a discussion of the results.

We present a summary equilibrium analysis of the first and second price slot auctions

in the appendix. We show that both auctions Nash implement the efficient allocation. Here we also demonstrate that standard properties of the airline revenue functions render the concept of dominant strategy implementability meaningless in the context of GDP slot auctions. Several computational issues are also discussed in the appendix.

1.2 GDP Slot Reassignment

The FAA expects U.S. air traffic to more than double by the year 2025. In an attempt to accommodate such growth, it has launched a campaign called NextGen to transform the National Airspace System (NAS) with 21st century technologies. The major goals of NextGen are to ensure that safety, capacity and environmental goals are met [see [FAA \(2009a,b\)](#)]. Much of NextGen is aimed at expanding NAS resources. However, it is recognized that these resources will always be finite, and so efficient allocation of resources is also a major part of the NextGen plan. Through the Collaborative Air Traffic Management (CATM) component of NextGen, the FAA is focusing on the role of decentralized market approaches to efficient allocation of NAS resources.

Of all NAS resources, departure and arrival slots have perhaps received the most attention for potential conversion to market allocation. The departure and arrival slots for a given airport are the scheduled times at which aircraft are permitted to takeoff and land. Currently, the International Air Transport Association (IATA) provides slot allocation procedures that rely on airlines voluntary cooperation through IATA coordination at biannual conferences. Hence, airlines do not currently pay for slots. Special consideration is given to High Density Rule (HDR) airports, which include New York – Kennedy and LaGuardia, Chicago – OHare, and Washington – Ronald Reagan. In HDR airports, airlines are permitted to sell and lease slots. However, there is a “use-it-or-lose-it” provision that the current holder of a slot must operate it at least 80% of the time or it is reclaimed to a pool of unused slots for reallocation [see [Le et al. \(2007\)](#); [Fan and Odoni \(2002\)](#); [Cholankeril et al. \(2003\)](#)].

There is particular interest in deploying market allocations of slots in the context of a Ground Delay Program (GDP). A GDP is a situation in which the scheduled demand for slots exceeds the safe supply of slots. Because the FAA does not schedule more arrivals or departures than the NAS can handle under normal conditions, this situation typically arises during bad weather, congestion or periods of heightened security. In such cases, the FAA institutes a GDP at the affected airports in order to limit the arrivals and departures to a safe level. Therefore, regardless of how the original schedule was determined, a GDP forces the FAA to quickly create a new schedule.

Consider the following example. JFK airport has 30 flights scheduled to depart over the next hour. Bad weather in the New York City area has made it so that JFK can safely handle only 15 departures. The problem is how to allocate the 15 GDP departure slots among the airlines that had one or more of the 30 scheduled pre-GDP departure slots (scheduled carriers).

Currently, the FAA employs Collaborative Decision Making (CDM) to allocate GDP slots among scheduled carriers. The CDM method involves Ration-by-Schedule (RBS).

This is exactly as it sounds. In the example of the previous paragraph, the scheduled carriers of the first 15 pre-GDP scheduled flights will get the 15 GDP slots. The carriers are free to use the allocated GDP slots for whichever flights they please. They are not required to operate the pre-GDP scheduled flight in the GDP slot. However, if an airline is not able to operate one of the GDP slots that it is assigned, the airline can relinquish the slot to the FAA. In return for the relinquished slot, the airline will receive priority for slots that become available due to other airlines not being able to operate their slots.

The efficiency of the RBS allocation, however, depends entirely on the pre-GDP schedule. RBS can not be expected to be efficient even in a simple world where every airline's profit is decreasing in the sum of the delays of each of its flights at exactly the same rate. The reason is that flights at the GDP-affected airport have different downstream effects. For example, the value to United Airlines of its flight #111 arriving on time in Newark may depend on whether passengers from other United Airlines flights will be switching to #111 in Newark before #111 continues on to its next destination. Hence, if airline profit only depended on GDP-affected flights, then RBS might be efficient. But since airline profit depends on the airline's entire network of flights, RBS is not generally efficient even in the simple world where every airline's utility is just an identical decreasing function of the sum of the delays of its flights. We can loosely refer to the type of demand interdependencies discussed here as "network effects."

In consideration of these network effects, many researchers have discussed auction allocation of arrival and departure slots. [Grether et al. \(1979, 1981\)](#) and [Rassenti et al. \(1982\)](#) are the seminal papers in that line of research. They propose annual auction allocations of all departure and arrival slots for all airports (NAS-wide) to replace the biannual IATA procedure. This is a much more extensive implementation than the auctions analyzed in this paper, which are implemented just at affected airports during GDPs. For NAS-wide slot auctions, [Grether et al.](#) suggest a two-stage approach. The primary stage consists of independent sealed-bid competitive auctions at each airport. The authors suggest a secondary market consisting of a computerized form of the oral double auction to account for the slot demand interdependencies that are not captured in the primary market.

To overcome certain disadvantages of the two-stage approach, [Rassenti et al.](#) suggest a sealed-bid combinatorial auction. This auction allows airlines to submit conditional bids and therefore accounts for demand interdependencies in a single stage. Winners are determined by an algorithm that uses the sealed bids to maximize system surplus. In this way, efficiency is directly built into the scheme. They suggest a secondary trade market to follow the combinatorial auction which would serve to correct inefficiencies that arise as a result of new information. [Rassenti et al.](#) experimentally compare the combinatorial auction to the market proposed by [Grether et al.](#) and find that the efficiency gains are increasing in the degree of complexity of the assignment problem.

Unfortunately, these NAS-wide auctions are not a realistic policy choice for GDP slot allocations. The primary reason is that their complexity and multiple stages require time and resources that are not available during GDPs. For instance, if there are F GDP slots,

and an airline has just $f \leq F$ flights to operate during the GDP, then there are

$$\sum_{k=1}^f \frac{F!}{k!(F-k)!}$$

conditional bids. This number gets large quickly. For $F = 15$ and $f = 5$, this airline has 4943 conditional bids. When there are multiple airlines with strategy vectors of such length, it is a non-trivial problem to determine the system surplus maximizing allocation as is required by a combinatorial auction. Furthermore, it is unlikely that airlines will be able to accurately determine the value of all allocations of all subsets of their flights in such short time. However, a GDP slot allocation scheme needs to be implementable in real-time when NAS resources are stretched thin. If the scheme is a market mechanism, then the facilitation costs should be minimal. The lack of preparation time also suggests that the strategy space should be simple even though the airline profit functions rarely are.

These considerations motivate the design of a simple GDP slot auction. In this auction, the number of bids for a given airline is equal to the number of GDP slots. Furthermore, the algorithm which assigns slots from the bid vectors is trivial. A more detailed account of the auction design is presented in section 2.

In the appendix, we show that the first price version of the GDP slot auction has the property that every NE is efficient, and that the second price version Nash implements the efficient outcome. However, we also show that the simplification of the strategy space means that neither version of this auction can be dominant strategy implementable. In particular, the lack of an airline profit function that is additively separable in GDP slots renders the concept of incentive compatibility meaningless. This means there is a theoretical trade-off posed by reducing the strategy space from the set of conditional bids, which is the case in the full combinatorial auction, to the set of unconditional bids, which is the case in the GDP slot auction presented here.

But these types of theoretical equilibrium properties are not the basis upon which rational policy makers choose among policy alternatives. Rather, according to the axioms of [Savage \(1954\)](#), a rational decision maker chooses among available alternatives by comparing their expected utilities. There is no reason, in principle, that this should be different for a policy maker. For the problem of GDP slot auctions, this means the following. If the FAA policy maker's utility for a mechanism depends on the strategies of the airlines, and the strategies of the airlines depend on the mechanism, then the FAA policy maker requires a probability distribution over the airlines' strategies for each mechanism. The ECA is unable to provide this because the output is a *set* of equilibrium strategies without associated probabilities.

Consider the imaginary case where, for every single mechanism, the ECA produces a unique equilibrium. One *might* interpret the unique equilibrium of a mechanism as a Dirac delta function about the equilibrium profile. However, this implies that all non-equilibrium strategy profiles do not occur, which is in strict conflict with virtually every experimental data set ever compiled. That is, *all* strategy profiles occur with some probability, including non-equilibrium strategy profiles. So if the policy maker chooses

among multiple mechanisms (each with a unique equilibrium) by evaluating her expected utility using the Dirac delta function about the equilibrium profile of each, her decision carries the assumption that the only possible strategy profile for each mechanism is the equilibrium behavior. It is not hard to see that such a model can lead to very bad decision making.

In contrast to the imaginary situation where all mechanisms have a unique equilibrium, in practice, at least some mechanisms will have multiple equilibria. Such is the case with GDP slot auctions, where the first and second price versions of the GDP slot auctions have many equilibria. Here, the equilibrium concept does not allow the policy maker to calculate expected utility at all. The reason is that the equilibrium concept simply lists the strategy profiles that satisfy the equilibrium condition. It does not provide the relative likelihoods of the elements of the equilibrium set. When there are multiple equilibria, there is no principled way to calculate the policy maker’s expected utility from a given mechanism. Without calculating expected utility for each mechanism, the policy maker cannot make use of the axioms of [Savage \(1954\)](#) to make a rational choice from her policy alternatives. Note also that the presence of multiple equilibrium profiles does not fix the problem that non-equilibrium profiles are implicitly given probability zero of occurring.

Because the ECA can only in rare special cases produce degenerate distributions over airline strategies, the FAA cannot use it to conduct a risk analysis. In particular, the ECA cannot be used to answer the following basic questions posed by NAS stakeholders and policy makers:

- “Which mechanism produces the least variance in efficiency?”
- “What is the probability that mechanism A is more efficient than mechanism B?”
- “Which mechanism has a lower probability of being less efficient than current practices?”
- “What is the 95% confidence interval for the efficiency of each mechanism?”
- “Which mechanism has the highest expected airline profits?”
- “What is the probability that profits will be below current profits?”

Fortunately, the PGT approach used in this paper produces a probability distribution over airline strategies. This means it can always be used to compare mechanisms in terms of the policy maker’s expected utility. It can also always be used to conduct a thorough risk analysis and answer the risk-relevant questions of stakeholders listed above. Furthermore, the PGT approach explicitly accounts for two types of risk that the ECA simply assumes away.

systematic risk: when there is uncertainty about which mixed strategy profile the players will choose, this is the risk that they will choose a mixed strategy profile that produces undesirable outcomes on average.

modeling risk: this is the risk associated with the researcher’s uncertainty about the appropriateness of a particular model for describing the strategic setting.

The way in which PGT is able to account for systematic and modeling risk is discussed in section 1.3.2.

For these reasons, and others yet to be mentioned, we choose to analyze our slot auction with PGT, a model of strategic behavior that outputs a probability distribution over all strategy profiles instead of an equilibrium set. The following subsection formally introduces the PGT approach.

1.3 PGT

1.3.1 Background on PGT

Say one wishes to predict some characteristic of interest y concerning some physical system, based on some information \mathcal{I} concerning the system. Statistics provides many ways to convert such a \mathcal{I} into a probability distribution over y .¹ Such a distribution is far more informative than a single “best prediction”.

However if needed we can synthesize the distribution with a single prediction. One way to do that is to use the mode of the distribution as the prediction. When the distribution is a Bayesian posterior probability, $P(y | \mathcal{I})$, this mode is called the Maximum A Posterior (MAP) prediction. Alternatively, say there is a real-valued loss function, $L(y, y')$ that quantifies the penalty we will incur if we predict y' and the true value is y . Then Bayesian decision theory counsels us to predict the “Bayes optimal” value, which is the y' that minimizes the posterior expected loss, $\int dy L(y, y') P(y | \mathcal{I})$ [see Jaynes and Bretthorst (2003); Gull (1988); Loredó (1990); Bernardo and Smith (2000); Berger (1985); Zellner (2004); Paris (1994); Horn (2003)].²

A priori, there is no reason that this standard approach to predicting the behavior of physical systems is not appropriate when the physical system in question is some human beings playing a game. To do this we would identify y with the joint choice made by the players in the game. For example, if the players are engaged in a conventional strategic form game, the choice of each player i is i ’s mixed strategy, which we will sometimes just call i ’s “strategy” for short [see Fudenberg and Tirole (1991); Myerson (1991); Aumann and Hart (1992); Basar and Olsder (1999)]. In such strategic form noncooperative games the moves of the players are independent, so the joint choice of the players – y – is the product of their mixed strategies, which we write as $q \equiv \prod_i q_i$. In this example, \mathcal{I} is the details of the game (e.g., the profit functions of the airlines), perhaps in conjunction with other information, like quantifications of how rational each player is. So the Bayesian posterior is a distribution over strategy profiles, $P(q | \mathcal{I})$.

¹In this paper we will sometimes be loose in distinguishing between probability distributions, probability density functions, etc., and will generically write any of them as “ $P(\dots)$ ” with the context making the meaning clear.

²In this paper, we will write integrals with the measure implicit. So if the set being “integrated over” is countable, we implicitly mean a point measure, in which the integral is equivalent to a sum.

We use the term **Predictive Game Theory** (PGT) to refer to any application of statistical inference (Bayesian or otherwise) *to* games, in contrast to the use of statistical inference by some players *within* a game. The ultimate goal of PGT is to use the same kinds of statistical tools to exploit all information about a system being predicted, whether that information is the utility functions of some players in the system, or some more conventional kind of statistical data concerning inanimate subsets of the overall system. In this paper we focus on PGT for GDP slot auctions. These are noncooperative strategic form games (although PGT is also applicable to cooperative games, unstructured bargaining, etc.). PGT replaces the ECA issue of how to specify a set of equilibrium strategy profiles for a specified game \mathcal{J} , $\mathcal{E}(\mathcal{J})$, with the issue of how to specify a density function over all possible joint strategies of that game, e.g., a Bayesian posterior $P(q \mid \mathcal{J})$. Therefore, PGT should be viewed as an alternative to the ECA, rather than as an alternative to any particular instance of the ECA, e.g. Nash equilibrium.

In the usual way, a loss function can be used to distill PGT’s density function over strategy profiles into a single predicted strategy, via decision theory. This mapping of a game to a single Bayes optimal strategy profile can be viewed as an “solution concept”. This solution concept depends on the loss function of the statistician who is making the prediction. (Note that this loss function is *not* specified in the game — the statistician making the prediction is external to the game.) Accordingly, this solution concept will vary with the external statistician who is making the prediction. This contrasts with the ECA, which ignores the concerns of the external statistician when telling that statistician what prediction to make. Another contrast with the ECA is that this solution concept typically produces a single strategy profile, without any need for a refinement.

Furthermore, often under the Bayes optimal strategy profile no player’s strategy is a best response to the strategies of the other players. Assuming there is more than one NE of the game, this is true even if the players are all fully rational, i.e., if the support of the density over strategy profiles is restricted to the NE. In this sense, “predictive” bounded rationality is automatic under PGT, in contrast to the case when using the ECA.

Finally, there are substantial computational differences between PGT and the ECA. In the ECA, numerical techniques are often needed to solve sets of simultaneous nonlinear equations. In contrast, under PGT numerical techniques are often needed to solve constrained maximization problems (e.g., if one wishes to find $\operatorname{argmax}_q P(q \mid \mathcal{J})$) or to solve integrals (e.g., if the loss function is quadratic, so that the Bayes optimal prediction is the average $\int dq q P(q \mid \mathcal{J})$). Especially in large problems like GDP slot auctions, the computational burdens of the numerical techniques used in PGT might be smaller than those under the ECA.

1.3.2 Advantages of PGT

PGT has all the usual benefits inherent in using a statistical approach to predict the real world. In contrast, many of these benefits are absent in the ECA. As a result PGT has some inherent advantages over the ECA; in this subsection we highlight three of them in the context of GDP slot auctions, ranging from the abstract to the highly practical.

1. One benefit of PGT is that, in general, it assigns non-zero probability to all mixed strategy profiles. In contrast, as mentioned above, the ECA generally assigns probably zero to almost all profiles, in the sense that it treats all strategy profiles outside a measure-zero equilibrium set as physically impossible.³ This means that unlike the ECA, PGT respects the fact that in the real world, all mixed strategy profiles can occur with some non-zero probability.

Indeed, if a given strategy profile is not an equilibrium strategy profile for a given equilibrium concept, then an observation of that profile invalidates that equilibrium concept, strictly speaking. In this narrow sense, every equilibrium concept suggested to date has been experimentally invalidated [see [Camerer \(2003\)](#)]. In contrast, under the PGT approach, the effect of observing any particular profile is simply to modify the posterior distribution over future profile observations.

In fact, the situation is worse than this for the ECA. Even if one restricts attention to the profiles in some equilibrium set, whenever there is more than one such profile, the ECA provides no information about the relative probabilities of those profiles. In this, every equilibrium concept is an incomplete predictive theory. In contrast, by definition PGT provides the relative probabilities of all profiles.

It is in this manner that PGT accounts for the systematic risk discussed above.

2. Another advantage of PGT is that, being a fully statistical model, it can combine multiple types of information / data into an associated posterior. This ability is necessary to properly express the uncertainty the modeler still has about the strategy profile after all that information. As an example, say the modeler is uncertain about the airline profit functions, so that \mathcal{J} is a distribution over possible profit functions. (Note that the modeler may have such uncertainty about the airlines' profit functions even for a complete information game, where the airlines have no such uncertainty about one another's profit functions.) Then the proper way for the modeler to express her associated uncertainty over mixed strategy profiles is by averaging over that distribution.

As a simple illustration, suppose m is the probability that the profit functions are \mathcal{J}' , and $1 - m$ the probability that they are instead \mathcal{J}'' . Then PGT says we must average over those two sets of utility information to properly express modeler uncertainty. Formally, we write $\mathcal{J} = \{\mathcal{J}', \mathcal{J}''\}$ and break the posterior into two terms:

$$P(q \mid \mathcal{J}) = mP(q \mid \mathcal{J}') + (1 - m)P(q \mid \mathcal{J}'').$$

In contrast, in the ECA, trying to address uncertainty about the profit functions in a similar fashion would entail averaging over the associated equilibrium sets somehow. It is not at all clear that the axiomatic foundations of the ECA provide a principled way of doing such averaging.

Futhermore, often we will have types of information that are relevant to our prediction of the strategy profile but that do not directly concern the game specification.

³One notable exception is the epsilon equilibrium concept of [Radner \(1980\)](#)

Examples of such information are demographic data, observational data concerning a particular player’s idiosyncracies (e.g., in the form of a Bayes net stochastic model of that player’s behavior in the absence of utility functions), and empirical data about the relative probabilities of various focal points. Again, a statistical approach like PGT is necessary to use these types of information to refine our prediction in a principled manner. (For example, given a distribution over focal points, one should (!) use it to average the posteriors given each possible focal point, in exact analogy to the average over utility functions described above.) In contrast, there is nothing in the ECA that would allow us to incorporate this information in such a principled way.

In fact, in PGT the principled integration of uncertainty extends even to uncertainty about what model of strategic behavior to use. Just as a modeler does not need to make a choice between utility information \mathcal{J}' and \mathcal{J}'' , she also does not need to make a choice among models that describe player behavior. As an example, she does not need to choose between a posterior $P(q \mid \mathcal{J}')$ motivated by the quantal response equilibrium (QRE) and a posterior $P(q \mid \mathcal{J}'')$ motivated by a level- k model [see [Costa-Gomes and Crawford \(2006\)](#)]. In fact, she *should not* make such a choice. Rather she should average over both posteriors, according to the relative probabilities that she assigns to the possibilities that each of those two models applies to her particular prediction problem. No such averaging is possible with the ECA.

This is precisely the manner in which PGT accounts for the modeling risk discussed above.

3. Perhaps the most important benefit of PGT’s statistical approach is that it not only allows us to address point prediction in a principled, decision theoretic manner (as described above), but also to address mechanism design problems this way. In fact, PGT allows us to extend the scope of “mechanism design” far beyond its usual domain, into a full-fledged theory of “game control”.

More precisely, say the FAA policy maker can set a parameter λ specifying some aspect of the GDP slot auction played by the GDP-affected airlines, whose mixed strategy profile is q , as usual. Let $G(q, \lambda)$ be the “efficiency” of the slot auction allocation (i.e. the utility function of the policy maker), and indicate the game specified by λ as Γ_λ . Let \mathcal{J} be some other information that the controller has concerning the game and/or player behavior, in addition to the value λ that she will choose. Then the standard approach of optimal control (i.e., Bayesian decision theory) says that the FAA policy maker should set λ to

$$\operatorname{argmax}_\lambda \left[\mathbb{E}(G \mid \mathcal{J}, \lambda) \right] = \operatorname{argmax}_\lambda \left[\int dq G(q, \lambda) P(q \mid \mathcal{J}, \lambda) \right] \quad (1)$$

So for example, if the policy maker’s utility function only depends on the pure strategy profile of the players, we can write $G(q, \lambda) = \int dx q(x) W(x)$ for some

function W . In this case the policy maker should set λ to

$$\operatorname{argmax}_{\lambda} \left[\int dq G(q, \lambda) P(q | \mathcal{J}, \lambda) \right] = \operatorname{argmax}_{\lambda} \left[\int dq dx W(x) q(x) P(q | \mathcal{J}, \lambda) \right] \quad (2)$$

There are many ways to extend the foregoing. For example, consider the case where the policy maker’s utility is not efficiency $G(q, \lambda)$, but rather $G(\theta, \lambda)$ where $\theta \in \Theta$ is set stochastically by $P(\theta | q, \lambda)$. In other words, the policy maker does not directly care about the mixed strategy profile, but rather about the ramifications of that profile on the state of some other system with state space Θ . As an example, say the auction allocation stochastically sets NAS-wide delays, θ , and the policy maker only cares about that value θ . Then the policy maker should set λ to maximize $\int dq dx P(\theta | x) q(x) P(q | \mathcal{J}, \lambda) W(\theta)$.

Furthermore, by using PGT we can compare choices of λ (i.e., choices of “mechanism”) based on other considerations beside the associated values of expected welfare. In particular, we can use PGT’s posterior to answer many of the questions that real-world NAS stakeholders will ask concerning the possible policy choices of a regulator, such as the questions suggested above:

- “Which mechanism produces the least variance in efficiency?”
- “What is the probability that mechanism A is more efficient than mechanism B?”
- “Which mechanism has a lower probability of being less efficient than current practices?”
- “What is the 95% confidence interval for the efficiency of each mechanism?”
- “Which mechanism has the highest expected airline profits?”
- “What is the probability that airline profits will be below current profits?”

In addition to all these advantages, the PGT approach to the control of games allows the policy maker to seamlessly incorporate constraints on her decision. For example, if she wants not to over-burden airlines, she can choose the slot allocation mechanism that maximizes efficiency subject to a constraint on airline profits. This constraint can be on the variance of total airline profits, the distribution of individual airline profits or the probability that each airline achieves some lower bound profit. In PGT, these types of calculations are straight-forward in principle, since the constraint uses only \mathcal{J} and the PGT posterior.

It is important to emphasize that PGT point predictions and mechanism choices are determined by the researcher’s and/or policy maker’s objectives, which are external to the game specification. (This is *exactly* like in conventional Bayesian decision theory and statistics.) This central importance of the modeler’s objectives is completely missed by the ECA.

In general, the axiomatic foundations of Bayesian decision theory tell us that modelers seeking to make a point prediction should minimize expected loss over the posterior, and policy makers seeking to choose a mechanism should maximize their own expected utility over the posterior. To do so they need to have a distribution over strategic behaviors. The focus of PGT is on providing such a distribution. Given such a distribution, the modeler / policy-maker can meet their goals, regardless of issues like whether there are multiple equilibria.

None of this is possible with the ECA.

1.3.3 The PGT Model of this Paper

As mentioned, the PGT approach encompasses Bayesian and non-Bayesian methods of formulating a distribution over strategy profiles, $P(q|\mathcal{I})$, from information \mathcal{I} . In this paper, we consider a specific example of a Bayesian PGT model. In this model, we specify an entropic prior $P(q)$. This prior is based on the maximum entropy principle of information theory and embodies the principle of insufficient reason [see [Shannon \(1948\)](#); [Jaynes \(1957, 2003\)](#)]. In short, if two strategy profiles are equally probable according to the profit function information, the entropic prior gives more weight to the profile that contains less information.

The likelihood, $\mathcal{L}(\mathcal{I}|q)$, is based on the QRE model of [McKelvey and Palfrey \(1995\)](#) and is called the QR-rationality likelihood (short for quantal response rationality). The QR-rationality likelihood assigns weight to a given q according to the degree to which each airline's strategy q_i is a profit-maximizing response to the profile of other airlines' strategies q_{-i} . This allows us to incorporate bounded rationality in a way that makes sense given the complex and rushed environment of a GDP. That is, we make the assumption that airlines seek to maximize profits but are not always capable of doing so. However, they are not completely incapable. In particular, given q_{-i} , the likelihood assigns greater weight to q_i than q'_i if and only if q_i achieves higher expected profit against q_{-i} than does q'_i .

2 Description of Slot Auctions and Notation

In this section we describe the simple GDP slot auction in both its first and second price versions.

Consider a one-shot game in which each of I airlines simultaneously submits a vector of bids, x_i , with the property that each element, $x_i(j)$ is nonnegative. The length of x_i is \hat{F} , the number of GDP slots. In other words, each airline submits a bid for each GDP slot.

Starting from the earliest slot, the FAA proceeds to allocate each slot to the highest bidder that has flights to operate in the current slot as well as all previously allocated GDP slots. Hence, the highest bidder for the j th slot will not be allocated slot j if that airline cannot operate slot j in addition to all the slots prior to j that it has already been allocated. In that case, the FAA will allocate slot j to the next highest bidder that

can. This information is known by the FAA from the pre-GDP schedule which lists the earliest slot each scheduled flight can fill. Each of the slots are allocated in this fashion, and no airline will have inoperable slots.

There are two versions of this auction, first price and second price. In the first price version, an airline that is awarded a slot pays exactly what it bid for that slot. In the second price version, an airline that is awarded a slot pays the next highest bid for that slot. When the winner of a slot is the airline that submitted the lowest bid, the cost is zero.

As mentioned above, the value to an airline of operating a given GDP slot depends on the other GDP slots it is operating. In other words, airline utility is generally not additively separable in GDP slots operated. This is due to the downstream effects of delays and cancellations.

Therefore, to explore the properties of the GDP slot auction via the PGT framework, it is informative to think about a very general utility form. We adopt the following,

$$\pi_i = R_i - C_i$$

where R_i is airline i 's revenue, which includes all operational costs and benefits, and C_i is the amount it pays to the FAA for its allocated flights as per the rules of the auction.

Let \mathcal{F} represent the set of pre-GDP slots. \mathcal{F}_i is the set of i 's pre-GDP slots, for which i presumably associates an operable flight to each element. Let $\hat{\mathcal{F}}$ be the set of GDP slots and $\Omega(\hat{\mathcal{F}})$ is the set of all subsets of GDP slots. Therefore, airline i 's revenue is a function, $R_i(\cdot|\mathcal{F}) : \Omega(\hat{\mathcal{F}}) \rightarrow \mathbb{R}$, that assigns a monetary value to each subset of GDP slots given the pre-GDP schedule. This is the value to i of optimally operating its assigned GDP slots. $\hat{\mathcal{F}}_i(x)$ is the set of GDP slots assigned to i under the profile of bids x .

3 PGT Model

We are interested in formulating a distribution over the space of mixed strategy profiles. The set of bid vectors (pure strategies) for airline i is X_i . The set of mixed strategies for airline i is $\Delta(X_i)$. A generic element of $\Delta(X_i)$ is q_i , a mixed strategy. The set of mixed strategy profiles is $\Delta_{\mathcal{X}} = \times_i \Delta(X_i)$. A generic element of $\Delta_{\mathcal{X}}$ is $q = \prod_i q_i$, a mixed strategy profile.

The central focus of the PGT approach, from which all predictive information is derived, is the posterior distribution, $P(q|\mathcal{J})$, over mixed strategy profiles $q \in \Delta_{\mathcal{X}}$:

$$P(q|\mathcal{J}) \propto P(q)\mathcal{L}(\mathcal{J}|q), \quad (3)$$

where $P(q)$ is the prior distribution over mixed strategy profiles, \mathcal{J} is information about the game specification, including profit functions, and $\mathcal{L}(\mathcal{J}|q)$ is the likelihood of \mathcal{J} given q .

3.1 Airline Profit Functions

To specify the airline profit functions, we closely follow the detailed account of airline cost functions presented in [Sherali et al. \(2006\)](#). In particular, we model airline i 's

revenue $R_i(\hat{\mathcal{F}}_i|\mathcal{F}_i)$ from operating a subset of its pre-GDP slots, $\hat{\mathcal{F}}_i$, as the cost of delays incurred by the airline for being forced to deviate from its pre-GDP schedule. Formulating revenues as delay costs carries the assumption that the (positive) portion of revenue that airlines collect in ticket sales is unaffected by GDP operations. This is reasonable given the standard legal clause on all ticket sales that claims the airline is not responsible for delays. Hence the maximum revenue for airline i is 0, which i receives from operating all of its pre-GDP flights without delay. The minimum revenue for airline i is $R_i(\emptyset|\mathcal{F}_i) < 0$, the delay cost of canceling all of i 's pre-GDP flights.

Let t_i be the minimum total delay in minutes for the flights that i operates during the GDP, and let t_0 be the expected delay per passenger for a canceled flight. Intuitively, passengers on a canceled flight will be delayed longer, on average, than passengers on a canceled flight. Therefore t_0 is greater than the longest possible delay for a pre-GDP flight operated during the GDP. Denote the number of canceled flights for airline i by \hat{F}_{i0} . Let \hat{l}_i be the average number of passengers per flight for airline i . Then $\hat{l}(t_i + t_0\hat{F}_{i0})$ is the total number of passenger minutes of delay on i 's pre-GDP flight schedule \mathcal{F}_i .

However, the airline cares not only about delays on \mathcal{F}_i . If delays on \mathcal{F}_i cause passengers to miss connections, then the connections depart with empty seats, and delayed passengers must be rescheduled on later flights. In this way, delayed flights impact other flights as the delayed passengers continue on their journey. Hence, the delays from \mathcal{F}_i can be multiplied through the system. Let d_i be the factor by which a single passenger delay minute is multiplied to give the total downstream effect, called the connection delay cost factor. For small regional and private airports, d_i is around 1.0. For medium hub airports, d_i is around 1.5. And for large hub or international airports, d_i is as much as 2.0.

This multiplicative factor, d_i , is calculated for a single flight. However, when there are multiple flights delayed and canceled, this factor grows. This is due to the fact that passengers arrive on one flight and connect to another. As the total number of impacted flights grows, there is greater downstream effect of a passenger delay minute because there are fewer rescheduling options available to the airline. For example, suppose an airline has one delayed arrival that results in missed connections. However, the airline has a later departure that can accommodate some of the passengers from the delayed arrival with the missed connections. However, if that departure is delayed because *that* aircraft's arrival is delayed, then the connection delay cost factor increases. Hence, the connection delay cost factor is given by $D_i = d_i + a(\hat{F}_{i0} + \hat{F}_{id})$, where \hat{F}_{id} is the number of delayed flights for airline i and $a > 0$ is the marginal impact of an additional delayed or canceled flight on the connection delay cost factor.

Therefore, the total number of downstream passenger delay minutes incurred by operating the GDP slots $\hat{\mathcal{F}}_i$ is given by $\hat{l}[t_i + \hat{F}_{i0}t_0]D_i$. Finally, [Chang et al. \(2001\)](#) estimate that the cost of a passenger delay minute is approximately \$0.20. Hence, airline i 's revenue (total delay cost) from operating the GDP slots $\hat{\mathcal{F}}_i$ is given by:

$$R_i(\hat{\mathcal{F}}_i|\mathcal{F}) = -0.2\hat{l}[t_i + \hat{F}_{i0}t_0]D_i.$$

Unless otherwise indicated, we use the following default values:

- $\hat{l} = 75$, from 809 million passengers and 10.7 million flights per year as reported by [US Department of Transportation Bureau of Transportation Statistics \(2009\)](#).
- $d_i = 1$, from [Sherali et al. \(2006\)](#)
- t_0 is twice the maximum delay for any flight for any airline.
- $a = .5$, so that the connection delay cost factor D_i ranges from 1 when zero flights are delayed or canceled to 2.5 when three flights (the maximum number of flights in our model) are delayed or canceled [see [Sherali et al. \(2006\)](#)].

This means airline i 's profit function is

$$\pi_i(\hat{\mathcal{F}}_i | \mathcal{F}_i) = -20[t_i + \hat{F}_{i0}t_0](1 + .5(\hat{F}_{i0} + \hat{F}_{id})) - C_i(x)$$

where $C_i(x)$ gives the auction costs airline i incurs from bids x .

As a final note, we assume delay cost functions are the same for all airlines in order to isolate the source of efficiency gains. In particular, if there are efficiency gains for slot auctions over RBS, it is due to asymmetries in the pre-GDP schedule of flights rather than asymmetries in delay cost functions. Allowing for asymmetric delay cost functions would only reduce the likelihood that RBS is efficient, thereby raising the relative efficiency of GDP slot auctions. Therefore, the assumption that all airlines have the same delay cost function is a conservative one. If anything, it will reduce efficiency gains of slot auctions.

3.2 QR-rationality Likelihood

In most strategic settings it is reasonable to think that players *seek* to maximize expected utility. However, it is also reasonable to think that players do not, in practice, always succeed in maximizing expected utility. In the context of GDP slot auctions, such reasons include the complexity of the network effects and the hurried environment in which airlines make strategic decisions. So, we would like to incorporate some notion of bounded rationality in our PGT model of GDP slot auctions. Our likelihood model, called QR-rationality (short for quantal response rationality), incorporates bounded rationality by borrowing from the concept of a logit quantal response. Under the logit quantal response, a player's rationality is given by the degree to which that player responds optimally to the other players' strategies. This degree of rationality is the criterion upon which our likelihood differentiates between q 's.

Before we formally introduce our measure of QR-rationality, we first need more notation. Let U_{q-i}^i be the vector of expected profits that airline i gets from each of its pure strategy bids against the mixture q_{-i} . We call this airline i 's *environment*. The logit mixed strategy distribution for airline i facing environment U_{q-i}^i is

$$\mathbb{L}_{U_{q-i}^i, \beta_i}(x_i) \propto e^{\beta_i \mathbf{E}_q(\pi_i | x_{i,j})}$$

where $\mathbf{E}_q(\pi_i | x_{i,j})$ is airline i 's expected profit from choosing its j 'th pure strategy against the mixture q_{-i} . The constant β_i is a measure of airline i 's rationality because as β_i

increases, the mixed strategy \mathbb{L} assigns greater probability to those pure strategies of i with greatest expected profit. As shown in McKelvey and Palfrey (1995), as $\beta_i \rightarrow \infty$, the logit mixed strategy is a best response to q_{-i} , i.e. it selects the strategy x_i that corresponds to the maximal element of $U_{q_{-i}}^i$.

So given any q (with finite support), the question is how to calculate β_i for each i . One method of doing so is to find the β_i that minimizes the Kullback-Leibler (KL) divergence from q_i to the logit distribution parameterized by β_i . The KL divergence is a concept from information theory that is used to measure the difference between two distributions [see Kullback and Leibler (1951); Kullback (1951, 1987)]. The KL divergence is:

$$\begin{aligned} KL(q_i(x_i), \mathbb{L}_{U_{q_{-i}}^i, \beta_i}(x_i)) &= \sum_{x_{i,j} \in X_i} q(x_{i,j}) \ln \left(\frac{q(x_{i,j})}{\mathbb{L}_{U_{q_{-i}}^i, \beta_i}(x_i)} \right) \\ &= \sum_{x_{i,j} \in X_i} q(x_{i,j}) \ln \left(\frac{q(x_{i,j}) \sum_{x_{i,l} \in X_i} e^{\beta_i \mathbf{E}_q(\pi_i | x_{i,l})}}{e^{\beta_i \mathbf{E}_q(\pi_i | x_{i,j})}} \right). \end{aligned} \quad (4)$$

By minimizing the KL divergence from q_i to the logit distribution parameterized by β_i , we are finding the logit distribution that most accurately models q_i in an information theoretic sense. Then we borrow the common interpretation of β_i as i 's rationality when playing $\mathbb{L}_{U_{q_{-i}}^i, \beta_i}(x_i)$ in response to q_{-i} .

In the special case where q_{-i} is such that all entries of $U_{q_{-i}}^i$ are identical, the QR-rationality parameter β_i can be any real number. This is the case in a mixed strategy NE with full support. For completeness we define $\beta_i = \infty$ when q_{-i} is such that all entries of $U_{q_{-i}}^i$ are identical. In other words, when i is indifferent among his pure strategies, he is perfectly rational by default.

This gives us the following characterization of rationality.

Definition 3.1. The *QR-rationality* of q_i against q_{-i} is the value of β_i that minimizes the KL distance from q_i to $\mathbb{L}_{U_{q_{-i}}^i, \beta_i}(x_i)$, equation 4.

One potentially worrisome property of the QR-rationality parameter, that is also shared by the logit-QRE, is that it is not invariant to positive rescalings of profits. In other words, airline QR-rationality parameters depend on currency units.

The next question is, given the choice of QR-rationality to measure how smart airlines are, what should the functional form of the likelihood, $\mathcal{L}(\mathcal{J}|q)$, be? In other words, in the absence of data about the particular airlines involved in the auction, how strongly do we believe they are likely to be smart, as measured by QR-rationality? One simple parameterized form is the following:

$$\mathcal{L}(\mathcal{J}|q) \propto \prod_i [\tanh(\alpha_i \beta_i(q)) + 1]^\gamma \quad (5)$$

where each α_i measures how much more likely i is to be smart rather than dumb. Physically, $\mathcal{L}(\mathcal{J}|q)$ quantifies how likely it is that out of all auctions a set of real-world airlines

could have just played, that they played the auction with profit information \mathcal{J} , given that they chose joint mixed strategy q when they participated in that auction.

When looking at equation 5, it may be tempting to think of it as a type of averaging of QRE's. This is not the case. Rather, not every $q \in \Delta_{\mathcal{X}}$ is a QRE for properly chosen β . That's because equation 5 is defined for every q , while only an infinitesimal subset of product distributions are logit distributions.

It should be emphasized that equation 5 is not the only reasonable choice for a QR-rationality likelihood. In general, when predicting player behavior in a game, ultimately the modeler must choose how to quantify their insight into how the system's state is related to what information they have concerning it, in terms of a likelihood. Like all models, this likelihood must be vetted with real world data.

Ultimately, the QR-rationality likelihood describes the underlying distribution of QR-rationalities in the set of airlines. The true distribution cannot be known with certainty, so any functional form will be wrong. The important point is that a non-degenerate distribution over rationalities is, in many settings, an improvement over an assumption of perfect rationality (as in NE) or a point mass assigned to a specific imperfect rationality (as in QRE). This is particularly true for settings in which learning has not yet converged to equilibrium, multiple equilibria exist, or computational complexities are involved (which covers most real-world settings). All of the above apply to the problem of GDP slot auctions.⁴

3.3 Entropic Prior

The role of the prior distribution, $P(q)$ is to quantify our subjective beliefs about the relative probabilities of mixed strategy profiles without regard to the profit information used by the likelihood function, $\mathcal{L}(\mathcal{J}|q)$. At first glance, the task of formulating any beliefs about a distribution of mixed strategy profiles without the benefit of utility information may seem difficult and/or unproductive.

However based on any of several separate sets of simple desiderata, there is a unique real-valued quantification of the amount of syntactic information in a distribution $q(x)$ [see Shannon (1948), Mackay (2003), Cover and Thomas (1991)]. That quantification, the Shannon entropy of a density q , is written as $S(q) = -\sum_x q(x) \ln(q(x))$. The entropic prior density is written as $P(q) \propto \exp(\delta S(q))$ for real-valued parameter δ .

For $\delta > 0$, the entropic prior assigns greater probability to mixed strategy profiles that are more diffuse. This is attractive from a modeling perspective because it represents an agnostic way of differentiating between q 's that have the same likelihood. More precisely, say we have two mixed strategy profiles, q and q' , that have the same QR-rationality, and

⁴Since we are using QR-rationality, the choice of the form of the likelihood function has implications for the convergence rate of Monte Carlo estimates of the posterior. This is because the QR-rationality parameter, β_i , can diverge to infinity for q_i that are best responses to q_{-i} . Infinite values of β are unlikely a problem in practice because best response correspondences are often of measure zero in the space of $\Delta(X_i)$. However, large β are quite possible. Therefore, if $\mathcal{L}(\mathcal{J}|q)$ is unbounded as $\beta(q)$ grows, Monte Carlo estimates of the posterior may never converge.

therefore the same likelihood. With $\delta > 0$, the posterior then favors the mixed strategy profile that has a smaller influence on the distribution over the support of the q 's,

$$P(x|\mathcal{J}) = \int_q q(x)P(q|\mathcal{J})dq. \quad (6)$$

The entropic prior is not the only candidate for prior distribution. Indeed, it just one member of the Cressie-Read family of distributions [see [Cressie and Read \(1984\)](#); [Read and Cressie \(1988\)](#)]. However, we adhere to the entropic prior ($\delta > 0$) as it is consistent with the principle of maximum entropy [see [Jaynes \(1957\)](#)], which can itself be derived from the principle of insufficient reason [see [Jaynes \(2003\)](#)]. The principle of insufficient reason tends that when faced with a set of possibilities that are indistinguishable based on the data at hand, each possibility should be equally likely [see [Poincare \(1912\)](#)]. The entropic prior upholds that principle because it says that for a given $\hat{\beta}$ the mixed strategy profile q that comes closest to putting equal weight on each pure strategy (i.e. maximizes entropy), subject to $\beta(q) = \hat{\beta}$, is the most likely.

3.4 Computational Methods

3.4.1 Sampling the Posterior

In order to make predictions about the outcome of the slot auction, we need to know $P(x|\mathcal{J})$ from equation 6, the posterior probabilities of each of the pure strategy profiles. More generally, researchers may want to know the expected value of any function $f(q)$ of the players' strategies

$$\mathbf{E}[f(q)] = \int_{\Delta_{\mathcal{X}}} f(q)P(q|\mathcal{J})dq. \quad (7)$$

This includes expected profits, expected welfare, expected covariance, etc.

Unfortunately, we cannot evaluate the posterior in closed form for any of the likelihoods discussed in this paper. Therefore, we must numerically estimate it. We use the Monte Carlo method of importance sampling to do so. Importance sampling relies on taking draws from a known distribution $H(q)$ in order to estimate an unknown distribution $P(q|\mathcal{J})$. This means we need a population of mixed strategy profiles, q 's, from the space of mixed strategy profiles, $\Delta_{\mathcal{X}}$.

In GDP slot auctions, airline i 's bids, $x_i(j)$, can take on any value in $[0, \bar{x}_i(j)]$. So a single q_i is a vector of infinite length. For obvious computational reasons, we cannot work directly with such vectors. Therefore, we need to discretize the space of mixed strategy profiles. That is, the actual value of $\mathbf{E}[f(q)]$ is an integral over an infinite-dimensional space, $\Delta_{\mathcal{X}}$, but we want to estimate this integral over a finite-dimensional space. However, we must be careful to do so in a way such that our estimate of the posterior approximates the actual posterior.

Our solution is to form a population of q 's by randomly drawing mixtures of Gaussian distributions. The q 's are drawn from the sampling distribution $H(q) = H(\rho, \mu, \Sigma)$, where ρ gives the convex weights for each component of each player's mixture, μ gives the mean for each component of each player's mixture, and Σ gives the covariance matrix for each

component of each player's mixture. Without much information about the space of joint distributions q , it is safest to explore the space of triples (ρ, μ, Σ) uniformly. Hence, each ρ_i is sampled uniformly from the \mathcal{M}_i -dimensional simplex, where \mathcal{M}_i is the number of mixture components in q_i . The means, μ_i , are sampled uniformly from the hypercube given by lower and upper bounds μ_{il} and μ_{ih} . Finally, Σ_i^j is the covariance matrix of the j 'th component of i 's mixture distribution. It is determined by random Jacobi rotations of a diagonal matrix with eigenvalues λ . These eigenvalues are drawn from a uniform distribution with lower bound λ_l and upper bound λ_h . In order to guarantee positive definiteness of Σ_i^j , λ_l is non-negative.

Specifically, to obtain each q , we draw a mixture of truncated multivariate normal distributions for each player,

$$q_i(x_i) = \begin{cases} \sum_{j=1}^{\mathcal{M}_i} \frac{\rho_i^j \phi_i^j(x_i)}{Z_i} & \text{if } L_i \leq x_i \leq B_i \\ 0 & \text{otherwise} \end{cases}$$

where

$$\phi_i^j(x_i) = \frac{1}{2\pi^{\mathcal{D}_i/2} |\Sigma_i^j|^{.5}} \exp \left[-.5(x_i - \mu_i^j)' (\Sigma_i^j)^{-1} (x_i - \mu_i^j) \right].$$

and

$$Z_i = \int_{\mathcal{B}_i} \sum_{j=1}^{\mathcal{M}_i} \rho_i^j \phi_i^j(x_i) dx_i.$$

The constant Z_i normalizes the mixture to the space of i 's allowable bids, \mathcal{B}_i . This region is bounded below by the origin and above by airline i 's budget. In particular, the sum of the components of i 's bid vector will not exceed the difference between i 's minimum and maximum possible revenue. \mathcal{D}_i is the dimensionality of i 's mixed strategy vector, i.e. the length of x_i as determined by the number of flights in the GDP schedule.

3.4.2 Estimating Statistics

Now that we have a method for sampling the posterior, it is possible to form Monte Carlo estimates of statistics that come from the posterior.

Let $q^{\rho, \mu, \sigma}$ be the parameterized mixed strategy profile distribution and $f(q^{\rho, \mu, \sigma})$ be any function of $q^{\rho, \mu, \sigma}$. The posterior expectation of $f(\cdot)$ is then:

$$\begin{aligned} \mathbf{E}_{\rho, \mu, \sigma}[f(q)] &= \int_{\rho, \mu, \sigma} f(q^{\rho, \mu, \sigma}) P(q^{\rho, \mu, \sigma} | \mathcal{J}) d\rho d\mu d\sigma \\ &= \int_{\rho, \mu, \sigma} f(q^{\rho, \mu, \sigma}) \frac{V(q^{\rho, \mu, \sigma})}{Z} d\rho d\mu d\sigma \end{aligned} \tag{8}$$

where

$$V(q^{\rho, \mu, \sigma}) = e^{\alpha S(q^{\rho, \mu, \sigma})} \mathcal{L}(\mathcal{J} | q^{\rho, \mu, \sigma})$$

and

$$Z = \int_{\rho, \mu, \sigma} V(q^{\rho, \mu, \sigma}) d\rho d\mu d\sigma$$

is the normalizing constant.

As an example, choose $f(q) = q$. Then $\mathbf{E}_{\rho,\mu,\sigma}(f(q) \mid \mathcal{J}) = \mathbf{E}_{\rho,\mu,\sigma}(q \mid \mathcal{J})$ is the expected mixed strategy profile. Now each mixed strategy profile q is a distribution $P(x \mid q)$. Accordingly, for this choice of f , $\mathbf{E}_{\rho,\mu,\sigma}(f(q) \mid \mathcal{J})$ is just the posterior expected pure strategy profile, $P(x \mid \mathcal{J})$.

We estimate the numerator integral in equation 8 with T i.i.d. samples

$$\{\rho(t), \mu(t), \Sigma(t)\}_{t=0}^T$$

from H . In the usual way with importance sampling [see Robert and Casella (2004)], we write

$$\int_{\rho,\mu,\sigma} f(q^{\rho,\mu,\sigma}) V(q^{\rho,\mu,\sigma}) d\rho d\mu d\sigma \simeq \frac{1}{T} \sum_{t=1}^T \frac{f(q^{\rho(t),\mu(t),\sigma(t)}) V(q^{\rho(t),\mu(t),\sigma(t)})}{H(q^{\rho(t),\mu(t),\sigma(t)})}$$

Similarly, we estimate the denominator integral by

$$\int_{\rho,\mu,\sigma} V(q^{\rho,\mu,\sigma}) d\rho d\mu d\sigma \simeq \frac{1}{T} \sum_{t=1}^T \frac{V(q^{\rho(t),\mu(t),\sigma(t)})}{H(q^{\rho(t),\mu(t),\sigma(t)})}.$$

4 Predictions and Decision-Ready Results

We confine our investigation to a situation in which there are two airlines (A and B) and five equally-spaced pre-GDP arrival slots. The slots are at minutes 10, 20, 30, 40 and 50. A one-hour GDP will begin ten minutes before the first slot, at minute 0. This means it will expire ten minutes after the fifth slot, at minute 60. Hence, only these five slots are impacted by the GDP. During the GDP, arrival capacity is reduced so that the airport can handle at most three equally-spaced arrivals, occurring at minutes 15, 30 and 45. The goal of a GDP slot reallocation scheme is to efficiently allocate the three slots to the two airlines. This situation is depicted in figure 1.

The pre-GDP arrival schedule lists the airline that operates each slot, the flight that is scheduled for each slot and the earliest runway time of arrival (ERTA) for each scheduled flight. Table 1 provides an example. We will indicate pre-GDP arrival schedules as a string of capital letters representing the airlines in the order of the pre-GDP schedule. For example, we would write *ABBAA* to indicate the schedule in table 1. Here *A* has the 10, 40 and 50 minute slots.

There are ten pre-GDP schedules in which airline A has three pre-GDP slots and airline B has two pre-GDP slots (3-2 schedules). There are five pre-GDP schedules in which airline A has four pre-GDP slots and airline B has one pre-GDP slot (4-1 schedules). We do not bother with the trivial case in which airline A has all five pre-GDP slots.

4.1 Summary of Strategy and Efficiency Predictions

We first present our predictions for the bid behavior of airlines A and B under the first and second price versions of auction 1 for all 3-2 and 4-1 schedules. We generate these

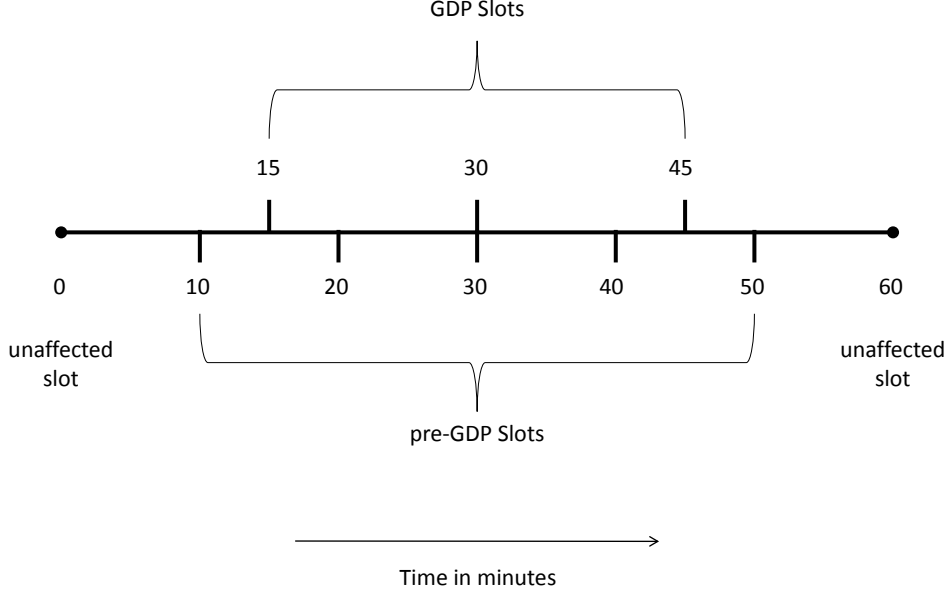


Figure 1: An illustration of the one-hour GDP affecting five arrival slots

predictions by minimizing expected quadratic loss over the posterior. This means that our prediction for airline i 's bid vector for a given pre-GDP schedule and price scheme is the expected value of i 's bid vector. This is given by

$$\mathbf{E}[x_i|\mathcal{J}] = \int_q \int_x dx dq x_i q(x) P(q|\mathcal{J}),$$

where \mathcal{J} contains the price scheme specification and the pre-GDP schedule in addition to airline delay cost information. For example, for the second price auction with pre-GDP schedule $AAAB B$ and profit functions π , we have that $\mathcal{J} = \{\text{Second Price}, AAAB B, \pi\}$.

Our predictions for the efficiency gain from the GDP slot auction are also generated by minimizing expected quadratic loss over the posterior. For a given GDP schedule and price scheme, the efficiency gain from the GDP slot auction is the expected value of total slot auction revenue minus RBS revenue. This is given by

$$\mathbf{E}[W(q|\mathcal{F}) - W_{RBS}|\mathcal{J}] = \int_q dq [W(q|\mathcal{F}) - W_{RBS}] P(q|\mathcal{J})$$

where $W(q|\mathcal{F})$ is the sum of expected revenue for airlines A and B when the mixed strategy profile is q , i.e.

$$W(q|\mathcal{F}) = \int_x dx [R_A(x|\mathcal{F}_A) + R_B(x|\mathcal{F}_B)] q(x).$$

Slot	Airline	Flight	ERTA
10	A	111	0
20	B	222	10
30	B	333	20
40	A	444	30
50	A	555	40

Table 1: Example pre-GDP arrival schedule showing the airline, flight and ERTA for each pre-GDP slot. This schedule can be summarized as ABBAA.

$W_{RBS}(\mathcal{F})$ is the sum of revenues to airlines A and B that arises from the RBS allocation when the pre-GDP schedule is \mathcal{F} .

Our estimates of $\mathbf{E}[x_i|\mathcal{J}]$ and $\mathbf{E}[W(q|\mathcal{F}) - W_{RBS}|\mathcal{J}]$ are generated using the importance sampling procedure detailed in section 3.4.2 above. These estimates serve as our predictions for slot bids and efficiency gain, and they are given for each pre-GDP schedule in tables 2 and 3. The likelihood and entropic prior parameter values are as follows, $\alpha = .5$, $\gamma = 4$, $\delta = 1$.⁵

First note that for some pre-GDP schedules, one of the airlines cannot use the first GDP slot (slot 15) because its earliest ERTA is after the first slot. This means the first GDP slot is allocated automatically to the airline that *can* use it. In this case, there are no bids reported for the first slot (indicated by “-” in the table). Instead, for these schedules, the bid vector for each airline is two dimensional rather than three dimensional. In addition, for one of the pre-GDP schedules (AAAAB), airline B cannot use the first *or* the second slot. In this special case, both slots are automatically allocated to airline A , and the only bids reported are for the third slot. As expected, corresponding bids are greater for the second price auction than they are for the first price auction.

We calculate average efficiency gain for each price scheme, first and second price, by invoking the principle of insufficient reason to assume that each pre-GDP schedule is equally likely. From that point it is trivial to average efficiency gains across schedules. This is one example of the way in which PGT explicitly accounts for modeling risk. That is, using the ECA, there is no way to average across efficiency gains for each pre-GDP schedule because there is no way to average across the equilibrium sets of each schedule.

Under the assumption that each pre-GDP schedule is equally likely, we find that both price schemes will provide efficiency gains, \$105.65 for first price auctions and \$149.65 for second price auctions. When tabulating efficiency gains for 3-2 and 4-1 schedules independently, we also find significant efficiency gains for both schemes. For the first price auction, 3-2 schedules average \$97.22 and 4-1 schedules average \$122.5. For the second price auction, 3-2 schedules average \$137.39 and 4-1 schedules average \$174.18. The efficiency gains are greater in every category for the second price auction. This means that if the policy maker’s objective function equals expected total airline revenue, then

⁵Naturally, one could average over alternative parameterizations to properly account for the modeling risk associated with specifying model parameters. The same can be said for the profit functions of the airlines.

pre-GDP Schedule	airline	slot 15 bid	slot 30 bid	slot 45 bid	efficiency gain
AAABB	A	-	1,031.41	823.79	143.71
	B	-	914.44	744.33	
AABAB*	A	-	1,120.20	940.06	-404.18
	B	-	1,002.95	826.54	
AABBA	A	-	992.00	1,289.91	282.97
	B	-	893.78	1,079.38	
ABABA	A	983.83	1,013.42	1,046.64	-107.98
	B	807.81	844.61	883.34	
ABBAA	A	1,026.75	1,137.16	1,101.41	420.80
	B	839.65	953.31	930.72	
BABAA	A	1,076.09	1,129.54	1,092.11	346.19
	B	904.73	934.62	933.99	
BBAAA	A	-	780.65	756.18	424.19
	B	-	517.26	556.42	
BAAAB	A	897.94	886.23	969.85	129.63
	B	782.79	856.54	790.21	
BAABA	A	1,066.33	1,040.94	1,066.52	-100.03
	B	887.31	848.59	898.67	
ABAAB	A	851.35	904.41	986.61	-163.08
	B	703.51	835.76	777.99	
AAAAB*	A	-	-	1,049.63	-219.58
	B	-	-	645.74	
AAABA*	A	-	888.19	932.87	-705.42
	B	-	468.47	518.28	
AABAA	A	-	994.55	957.19	676.61
	B	-	545.53	536.95	
ABAAA	A	607.42	691.51	665.92	522.23
	B	338.71	398.52	403.00	
BAAAA	A	663.88	690.82	663.35	338.69
	B	364.62	396.78	396.46	

Average First Price Auction Efficiency Gain: 105.65

Table 2: First Price Auction Estimates. An asterisk (*) next to the schedule indicates that the RBS allocation for this schedule is efficient. Average efficiency gain for 3-2 schedules is 97.22. Average efficiency gain for 4-1 schedules is 122.51.

she ranks the policy alternatives as $SP \succ FP \succ RBS$.

It is important to note that the efficiency gains presented here arise only from the asymmetry of pre-GDP schedules. That is, we assume the same delay cost function for both airlines, so the pre-GDP schedule is the only reason that some GDP schedules are more efficient than the RBS allocation. Allowing for asymmetry in delay costs would

pre-GDP Schedule	airline	slot 15 bid	slot 30 bid	slot 45 bid	efficiency gain
AAABB	A	-	1,330.37	983.51	239.71
	B	-	1,068.52	944.60	
AABAB*	A	-	1,414.03	1,144.99	-238.47
	B	-	1,208.81	1,039.96	
AABBA	A	-	1,153.92	1,738.75	323.94
	B	-	1,203.26	1,259.25	
ABABA	A	1,121.62	1,167.28	1,224.85	-103.40
	B	895.29	960.56	1,006.06	
ABBAA	A	1,164.50	1,341.29	1,296.91	421.16
	B	968.16	1,105.17	1,056.99	
BABAA	A	1,290.31	1,314.19	1,236.12	356.45
	B	1,044.19	1,047.41	1,046.64	
BBAAA	A	-	921.53	908.74	487.60
	B	-	560.11	578.68	
BAAAB	A	931.37	919.33	1,260.91	139.07
	B	858.58	983.20	877.54	
BAABA	A	1,231.00	1,176.87	1,274.05	-97.65
	B	1,005.98	982.90	1,012.29	
ABAAB	A	847.25	984.79	1,261.38	-154.47
	B	792.27	947.56	854.72	
AAAAB*	A	-	-	1,284.14	-155.53
	B	-	-	704.40	
AAABA*	A	-	995.44	1,065.66	-646.90
	B	-	517.02	531.20	
AABAA	A	-	1,130.65	1,083.88	755.15
	B	-	578.07	552.45	
ABAAA	A	659.61	787.12	769.70	548.05
	B	373.10	418.73	407.62	
BAAAA	A	732.41	766.77	741.82	370.12
	B	389.32	404.57	407.36	

Average Second Price Auction Efficiency Gain: 149.65

Table 3: Second Price Auction Estimates. An asterisk (*) next to the schedule indicates that the RBS allocation for this schedule is efficient. Average efficiency gain for 3-2 schedules is 137.39. Average efficiency gain for 4-1 schedules is 174.18.

decrease the likelihood that the RBS allocation is efficient and increase the efficiency gain of the slot auction. Similarly, expanding the analysis to include more airlines and more pre-GDP slots increases the potential for asymmetry in pre-GDP schedules. Hence, we should also expect efficiency gains to increase as the pre-GDP schedule expands in number of airlines and slots. Therefore, our setup with only two airlines and five pre-GDP

slots is conservative in terms of expected efficiency gains.

4.2 Posterior Distributions and Risk Information

The PGT posterior provides much more information than the point predictions reported in tables 2 and 3. For example, we can provide an estimate of the posterior expected mixed strategy profile given a price scheme and pre-GDP schedule,

$$P(x|\mathcal{J}) = \mathbf{E}[q(x)|\mathcal{J}] = \int_q dq q(x)P(q|\mathcal{J}).$$

This distribution contains a wealth of information, including risk information such as covariance. Figure 2 below shows the marginalization of $P(x|\mathcal{J})$ into separate distributions $P(x_A|\mathcal{J}) = \mathbf{E}[q_A(x_A)|\mathcal{J}]$ and $P(x_B|\mathcal{J}) = \mathbf{E}[q_B(x_B)|\mathcal{J}]$ for the second price auction with pre-GDP schedule AAABB. We display these marginalizations rather than the full distribution, $P(x|\mathcal{J})$, because x is four-dimensional (two airlines, two GDP slots) and cannot be displayed.




Figure 2: Contour representation of posterior expected mixed strategies, $P(x_i|\mathcal{J})$ for $i = A, B$, given second price auction and pre-GDP schedule AAABB. Darker colors represent greater probability density. Axes are in thousands.

The covariance matrix of $P(x|\mathcal{J})$ is given by

$$\text{Cov}(x) = \int_q \int_x dq dx (x - \bar{x})(x - \bar{x})' q(x) P(q|\mathcal{J}),$$

where \bar{x} is $\mathbf{E}[x|\mathcal{J}]$. For the second price auction with pre-GDP schedule AAABB, $\text{Cov}(x)$ is given by equation 9. The nonzero entries tell us that, in expectation, bids *appear* to be correlated across airlines. This is despite the fact that airline strategies are chosen independently. The correlation is a result of averaging over q 's weighted by the posterior. That is, the posterior couples the strategies of the airlines. Hence, observing airline A 's bid tells us something, *in expectation over all strategies*, about airline B 's bid. This is in contrast to what one finds under an equilibrium concept approach such as NE or QRE. Under these approaches, there is zero correlation between the bids of airlines A and B because there is zero uncertainty about the equilibrium mixed strategy profile, q^* . Note that the covariance between bids from different airlines (northeast and southwest quadrants) is smaller in magnitude than the covariance between bids from the same airline.

$$\begin{aligned} \text{Cov}(x) &= \begin{bmatrix} \text{Var}(x_{A,30}) & \text{Cov}(x_{A,30}x_{A,45}) & \text{Cov}(x_{A,30}x_{B,30}) & \text{Cov}(x_{A,30}x_{B,45}) \\ \text{Cov}(x_{A,30}x_{A,45}) & \text{Var}(x_{A,45}) & \text{Cov}(x_{A,45}x_{B,30}) & \text{Cov}(x_{A,45}x_{B,45}) \\ \text{Cov}(x_{A,30}x_{B,30}) & \text{Cov}(x_{A,45}x_{B,30}) & \text{Var}(x_{B,30}) & \text{Cov}(x_{B,30}x_{B,45}) \\ \text{Cov}(x_{A,30}x_{B,45}) & \text{Cov}(x_{A,45}x_{B,45}) & \text{Cov}(x_{B,30}x_{B,45}) & \text{Var}(x_{B,45}) \end{bmatrix} \\ &= \begin{bmatrix} 462.66 & -243.49 & -96.25 & 85.78 \\ -243.49 & 421.05 & 92.85 & -88.85 \\ -96.25 & 92.85 & 392.57 & -202.41 \\ 85.78 & -88.85 & -202.41 & 349.86 \end{bmatrix}. \end{aligned} \quad (9)$$

Similarly we can report the posterior distribution over firm profits, $P(\pi|\mathcal{J})$. The posterior distributions over airline profits for first and second price auctions are given in figure 3 for the pre-GDP schedule AAABB. From the posterior over airline profits, it is straightforward to calculate the variance of profits, the probability that each airline achieves some threshold value of profits and many other policy-relevant quantities. For example, under the second price auction, the probability that airline A earns profits greater than $-\$4,000$ and airline B earns profits greater than $-\$3,500$ is approximately 0.2451. We get this by integrating $P(\pi|\mathcal{J}_{SP})$ over the region $\{(\pi_A, \pi_B) \in \mathbb{R}^2 : \pi_A \geq -4,000, \pi_B \geq -3,500\}$. In point of comparison, the same probability for the first price auction with pre-GDP schedule AABBA is 0.0016.

Figure 3 is a prime example of the way in which the PGT approach accounts for systematic risk. In particular, figure 3 shows the posterior distribution over the *expected* profits for q , i.e. $\pi(q) = \mathbf{E}_q[\pi] = \int_x dx q(x)\pi(x)$, where q is chosen randomly according to $P(q|\mathcal{J})$. In contrast, the analogous distribution under the ECA is not a distribution at all because the ECA conveys no uncertainty about the choice of q . When there is a unique equilibrium, q^* , the ECA reports a unique, deterministic value of $\pi(q^*)$. This obviously hides the uncertainty that really exists about q and π . When there are multiple equilibria, as is the case in GDP slot auctions, there is an equilibrium set of profits $\{\pi(q^{*,1}), \dots, \pi(q^{*,E})\}$, but no associated probability distribution.

We also verify that expected total airline profits under second price auctions are greater than under first price auctions for every pre-GDP schedule. Averaging across all pre-GDP schedules, the expected difference between second price and first price profits is




Figure 3: Contour plots of posterior distributions over airline profits $\pi = (\pi_A, \pi_B)$ for first and second price slot auctions with pre-GDP schedule AAABB.

approximately \$1,416.20. For 3-2 pre-GDP schedules, the expected difference is approximately \$1,562.40. For 4-1 pre-GDP schedules, the expected difference is approximately \$1,123.90.

Obviously though, it is not the case that second price slot auction profits are expected to be greater than RBS profits, even though second price auctions are expected to be more efficient than RBS. The reason is that, while second price auctions minimize delay costs, the fact that it is an auction means that airlines incur the additional costs of their bids. The auction costs do not represent a loss in efficiency like delay costs. Auction costs are simply a transfer from the airlines to the FAA. However, for airlines, the reduction in delay costs that result from the slot auction do not, on average, outweigh the auction costs they incur. Therefore, second price profits are generally lower than RBS profits.

From a modeling perspective, it is informative to look at the posterior distribution over our rationality measure β , $P(\beta|\mathcal{J})$. This is shown in figure 4 for the second price auction with pre-GDP schedule AAABB. $P(\beta|\mathcal{J})$ is the distribution that we imply when we specify the QR-rationality likelihood. One informative quantity that we can calculate from $P(\beta|\mathcal{J})$ is the probability that both airlines have rationality greater than zero. To find this probability, we integrate $P(\beta|\mathcal{J})$ over the region $\{(\beta_A, \beta_B) \in \mathbb{R}^2 : \beta_A \geq$

$0, \beta_B \geq 0\}$. Doing so for the distribution in 4, we get 0.9425. The interpretation is that there is approximately a 94% chance that neither of the airlines' strategies are anti-rational ($\beta \leq 0$) given the strategy of the other airline. In point of comparison, the same probability for the first price auction with pre-GDP schedule AAABB is 0.9214.




Figure 4: Contour plots of posterior distributions over airline rationality $\beta = (\beta_A, \beta_B)$ for the second price slot auction with pre-GDP schedule AAABB.

In the same way that the expected bids reported in tables 2 and 3 are statistics of the posterior expected mixed strategy profile $P(x|\mathcal{J})$, the expected efficiency gains reported in those tables are statistics of the posterior distribution over efficiency gains, ω ,

$$P(\omega|\mathcal{J}) = \int_{\{q: W(q|\mathcal{F}) - W_{RBS} = \omega\}} dq P(q|\mathcal{J}).$$

Figure 5 plots the distributions over expected efficiency gains for the first and second price auctions. These distributions are averaged over *all* 3-2 and 4-1 schedules. That is, to construct the distribution for the first price auction, we find $P(\omega|\mathcal{J})$ for each pre-GDP schedule. Since there are fifteen pre-GDP schedules, each considered to be equally likely, we summed the fifteen distributions and divided by fifteen. The distribution over efficiency gains in the second price auction is constructed in the same way. The first thing to note is that both distributions are bimodal around zero. This arises because the RBS




Figure 5: Plot of the posterior distributions over efficiency gains for first and second price auctions.

allocation is efficient for some pre-GDP schedules. For such schedules, $P(\omega|\mathcal{J}) = 0$ for all $\omega > 0$. Similarly, the RBS allocation is the least efficient allocation for some pre-GDP schedules. For such schedules, $P(\omega|\mathcal{J}) = 0$ for all $\omega < 0$. Averaging over all schedules results in lower density around zero.

It should be noted that one cannot produce these plots using equilibrium concepts like NE and QRE. To be clear, if there existed a unique equilibrium, then there would be no uncertainty about $W(q|\mathcal{J})$. Hence, the distribution over efficiency gains would be a Dirac delta function about the equilibrium value $W(q^*)$, where q^* is the unique equilibrium. However, there are generally multiple equilibria for each pre-GDP schedule in first and second price slot auctions. There is no principled way to average over equilibrium sets to get a distribution over efficiency gains.

From the distributions in figure 5 we can calculate other decision-relevant statistics. For example, as reported in table 3, the expected efficiency gains for the second price auction are \$323.94, while they are only \$282.97 for the first price auction. The variance of efficiency gains for the first price auction is approximately 163.66. The variance of efficiency gains for the second price auction is lower at approximately 155.33. Another important quantity is the probability of negative efficiency gains. We get this by integrating the distributions in figure 5 from $-\infty$ to 0. The probabilities are approximately

0.3764 and 0.3760 for the first and second price auctions respectively. This means that the probability of an efficiency loss is roughly the same for both auction schemes. However, the probability of a loss greater than \$250 is significantly higher for the first price auction (0.1806) than for the second price auction (0.1238). On the other side, the probability of a gain greater than \$250 is significantly higher for the second price auction (0.4385) than for the first price auction (0.3828).

5 Discussion and Conclusions

This study represents the first application of PGT to predicting strategic behavior and evaluating risk in a real-world policy domain. PGT overcomes the shortcomings of conventional equilibrium approaches and produces decision-relevant information regarding policy risk and efficiency far beyond what is possible under conventional approaches. This includes full distributions over airline profits, slot bids and system efficiency as well as the statistics associated with those distributions. We account for two types of risk, systematic risk and modeling risk. Neither can be accounted for with the ECA.

Using the PGT approach, we find that the second price GDP slot auction dominates the first price auction in *every* decision-relevant category. It yields greater expected efficiency gains with lower downside risk and greater upside potential. It also yields lower variance in efficiency gains. Expected profits are greater under second price GDP slot auctions than under first price auctions for every pre-GDP schedule. So it seems that FAA policy-makers should have a strong preference for second price GDP slot auctions over first price auctions. This is despite the fact that none of the conventional arguments for second price auctions, such as dominant strategy implementability, even apply to GDP slot auctions.

We also find that both the first and second price auctions have higher expected efficiency than RBS. This efficiency gain arises even though we take a conservative approach in modeling airline GDP revenue as delay costs. In particular, we assume that delay costs are the only source of inefficiency, and we let the delay cost functions be symmetric across airlines. Therefore, efficiency gains over RBS are solely a result of the asymmetric pre-GDP schedules of the airlines. Introducing asymmetric delay cost functions, more airlines and more pre-GDP slots will only increase overall asymmetry thereby decreasing the likelihood that RBS is efficient.

Future work in this area should expand on the degree to which we account for modeling risk. In particular, this means expanding the point estimates of airline profit functions used in our analysis to a full-fledged distribution over profit functions. This also means combining the particular QR-rationality likelihood model used here with models based on concepts like level-k thinking and alternative notions of rationality, some of which are discussed in [Bono and Wolpert \(2009\)](#).

References

- R. Aumann and S. Hart. *Handbook of Game Theory with Economic Applications*. North-Holland Press, 1992.
- T. Basar and G. Olsder. *Dynamic Noncooperative Game Theory*. Siam, Philadelphia, PA, 1999. Second Edition.
- J. M. Berger. *Statistical Decision theory and Bayesian Analysis*. Springer-Verlag, 1985.
- J. Bernardo and A. Smith. *Bayesian Theory*. Wiley and Sons, 2000.
- J. W. Bono and D. H. Wolpert. How to use decision theory to choose among mechanisms. submitted to American Economic Review, August 2009.
- C. F. Camerer. *Behavioral Game theory: experiments in strategic interaction*. Princeton University Press, 2003.
- K. Chang, K. Howard, R. Oiesen, L. Shisler, M. Tanio, and M. Wambsganss. Enhancements to the faa ground-delay program under collaborative decision making. *Interfaces*, 31:5776, 2001.
- E. Cholaneril, W. Hall, and J.-P. Clarke. Evaluation of an auction mechanism for allocating airport arrival slots. Presentation document <http://aviation-appl.section.informs.org/Auctions.ppt>, June 2003.
- M. Costa-Gomes and V. Crawford. Cognition and behavior in two-person guessing games: An experimental study. *American Economic Review*, 96(5):1737–1768, December 2006.
- T. Cover and J. Thomas. *Elements of Information Theory*. Wiley-Interscience, New York, 1991.
- N. Cressie and T. R. C. Read. Multinomial goodness of fit tests. *Journal of the Royal Statistical Society, Series B*, 46:440464, 1984.
- FAA. Nextgen commitments and fy09 work plan. Technical report, Federal Aviation Administration, 2009a. <http://www.faa.gov/about/initiatives/nextgen/resources/commitments.pdf>.
- FAA. Faa’s nextgen implementation plan 2009. Technical report, Federal Aviation Administration, 2009b. http://www.faa.gov/about/initiatives/nextgen/media/NGIP_0130.pdf.
- T. P. Fan and A. R. Odoni. A practical perspective on airport demand management. *Air Traffic Control Quarterly*, 10(3):285–306, 2002.
- D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, Cambridge, MA, 1991.

- D. Grether, M. Isaac, and C. Plott. Alternative methods of allocating airport slots: Performance and evaluation. Technical report, CAB Report. Pasadena, Calif.: Polynomics Research Laboratories, Inc., 1979.
- D. Grether, M. Isaac, and C. Plott. The allocation of landing rights by unanimity among competitors. *American Economic Review*, 71:166–171, May 1981.
- S. F. Gull. Bayesian inductive inference and maximum entropy. In *Maximum Entropy and Bayesian Methods*, pages 53–74. Kluwer Academic Publishers, 1988.
- K. S. V. Horn. Constructing a logic of plausible inference: a guide to cox’s theorem. *International Journal of Approximate Reasoning*, 34(1):3–24, 2003.
- E. Jaynes. Information theory and statistical mechanics. *Phys. Rev.*, 106:620, 1957.
- E. T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, 2003.
- E. T. Jaynes and G. L. Bretthorst. *Probability Theory : The Logic of Science*. Cambridge University Press, 2003.
- S. Kullback. *Information theory and statistics*. John Wiley and Sons, NY, 1951.
- S. Kullback. The kullback-leibler distance. *The American Statistician*, 41:340341, 1987.
- S. Kullback and R. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):7986, 1951.
- L. Le, G. Donohue, and C.-H. Chen. Auction-based slot allocation for traffic demand management at hartsfield atlanta international airport: A case study. *Transportation Research Record: Journal of the Transportation Research Board*, 1888(2004):50–58, 2007.
- T. Loredo. From laplace to sn 1987a: Bayesian inference in astrophysics. In *Maximum Entropy and Bayesian Methods*, pages 81–142. Kluwer Academic Publishers, 1990.
- D. Mackay. *Information theory, inference, and learning algorithms*. Cambridge University Press, 2003.
- R. D. McKelvey and T. R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38, July 1995.
- R. B. Myerson. *Game theory: Analysis of Conflict*. Harvard University Press, 1991.
- J. B. Paris. *The Uncertain Reasoner’s Companion: A Mathematical Perspective*. Cambridge University Press, 1994.
- H. Poincare. *Calcul des probabilités*, chapter 1. Paris, 1912. Reprinted Jacque Gabay, 1987.

- R. Radner. Collusive behavior in noncooperative epsilon-equilibria of oligopolies with long but finite lives. *Journal of Economic Theory*, 22(2):136–154, April 1980.
- S. Rassenti, V. Smith, and R. Bulfin. A combinatorial auction mechanism for airport time slot allocation. *Bell Journal of Economics*, 13(2):402–417, Autumn 1982.
- T. R. C. Read and N. A. C. Cressie. *Goodness-of-Fit Statistics for Discrete Multivariate Data*. Springer-Verlag, New York, 1988.
- C. P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer, 2nd edition, 2004.
- L. J. Savage. *Foundations of Statistics*. John Wiley, New York, 1954.
- C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, 623–656, July, October 1948.
- H. D. Sherali, R. W. Staats, and A. A. Trani. An airspace-planning and collaborative decision-making model: Part ii cost model, data considerations, and computations. *TRANSPORTATION SCIENCE*, 40(2):147164, May 2006.
- US Department of Transportation Bureau of Transportation Statistics. Summary 2008 traffic data for u.s and foreign airlines: Total passengers down 3.5 percent from 2007. Research and Innovative Technology Administration Press-Release, April 2009. URL http://www.bts.gov/press_releases/2009/bts019_09/pdf/bts019_09.pdf.
- D. H. Wolpert and J. W. Bono. Distribution-valued solution concepts. working paper, October 2009.
- A. Zellner. Some aspects of the history of bayesian information processing. *Journal of Econometrics*, 2004.

A Auction Equilibrium Analysis

We establish several fundamental equilibrium properties of the first and second price auctions. The analysis in this section is superficial compared to the equilibrium analysis found in most conventional auction papers. This is because we are primarily concerned with the PGT analysis found in section 4; for all of the reasons mentioned above.

In what follows, we show that there are efficient NE of both first and second price GDP slot auctions. We also discuss the prospect of dominant-strategy implementation and incentive compatibility. Note that for individual rationality note that the FAA's procedures are required for all operating airlines. The alternative is to stop operating. We do not address individual rationality to the extent that airlines may choose to go out of business as a result of the implementation of GDP slot auctions. In the PGT analysis, we can easily incorporate the "going out of business" level of airline profits as a constraint faced by the policy maker.

We find the following restriction reasonable and useful in our analysis.

Definition A.1. A revenue function has the property of *nonnegative returns* if the marginal revenue to winning additional slots is always nonnegative. Alternatively,

$$R_i(\hat{\mathcal{F}}_i \cup \hat{\mathcal{F}}'_i | \mathcal{F}) - R_i(\hat{\mathcal{F}}_i | \mathcal{F}) \geq 0$$

for all subsets $\hat{\mathcal{F}}'_i$ of $\hat{\mathcal{F}}$.

The property of nonnegative returns says that the revenue of operating a subset of GDP slots does not decrease when the airline has rights to operate additional slots. Our reasoning is that an airline can simply refuse to operate or sell the additional slots without penalty if it finds that operating the slots is costly.

A.1 First Price Auction Equilibria

In order to discuss equilibria of GDP slot auctions in more general terms, assume the following notation:

- let $Y \subseteq \hat{\mathcal{F}}$ with elements y_j for $j = \{1, 2, \dots, |Y|\}$.
- let $Y(i, x)$ be the set of slots in Y that are won by airline i under the profile x . That is, $Y(i, x) = \hat{\mathcal{F}}_i \cap Y$. $Y(i, x)^c$ is the complement of $Y(i, x)$ in Y .
- let the marginal revenue of $Y \subseteq \hat{\mathcal{F}}$ to airline i under profile x be given by:

$$\Delta R_i(Y, x | \mathcal{F}) = R_i(\hat{\mathcal{F}}_i(x) \cup Y(i, x)^c | \mathcal{F}) - R_i(\hat{\mathcal{F}}_i(x) \setminus Y(i, x) | \mathcal{F})$$

- let $N(Y, x)$ be the set of players that win slots in Y under x .

To explore the efficiency properties of the first price auction, we assume that the FAA seeks to maximize the sum of airline revenues. In other words, we treat the costs (auction payments) of the airlines as taxes. Since they are paid to the FAA, an agency of the federal government, they enter the general pool of government spending. We also choose not to model passenger/customer welfare. There are many notions of how the FAA's objectives might differ from the notion of efficiency we adopt here. However, to facilitate the clear interpretation of results, we adhere strictly to this notion.

Theorem 1. *All Nash equilibria of the first price GDP slot auction where airline revenues exhibit nonnegative returns implement the revenue maximizing allocation.*

Proof. The sum of airline revenues is:

$$\sum_i R_i(\hat{\mathcal{F}}_i(x)|\mathcal{F}). \quad (10)$$

Suppose x^* is a NE that does not maximize 10. Then for some $Y \subseteq \hat{\mathcal{F}}$, there exists an airline i such that

$$\Delta R_i(Y, x|\mathcal{F}) > \sum_{l \in N(Y, x)} \Delta R_l(Y(l, x), x|\mathcal{F}).$$

This means either (1) that there exists an x'_i with $x'_i(Y) > x_i^*(Y)$ such that $\pi_i(x'_i, x_{-i}^*) > \pi_i(x_i^*, x_{-i}^*)$ or (2) that there exists a $j \in N(Y, x)$ and x'_j with $x'_j(Y(j, x)) < x_j^*$ such that $\pi_j(x'_j, x_{-j}^*) > \pi_j(x_j^*, x_{-j}^*)$. Both possibilities violate the equilibrium concept. \square

The above demonstrates that the first price GDP slot auction implements the revenue maximizing outcome in all Nash equilibria. However, as with other first price auctions, there is no dominant strategy in the first price GDP slot auction. Therefore, the revenue maximizing outcome is not dominant-strategy implementable.

Incentive compatibility, in the context of GDP slot auctions, cannot be handled in the usual way. Airlines bid on all slots simultaneously. However, as previously mentioned, airline utilities are not additively separable in GDP slots. Therefore, in general, airlines would not be able to bid their valuation for each subset of GDP slots simultaneously.

To be more concrete, for an airline's bid vector to represent its valuation of all subsets of GDP slots simultaneously we must have that

$$\sum_{j \in Y} x_i(j) = R_i(Y|\mathcal{F})$$

for all $Y \subseteq \hat{\mathcal{F}}$. Let $\hat{F} = |\hat{\mathcal{F}}|$ be the number of GDP slots. Then incentive compatibility means that each airline has a system of $2^{\hat{F}-1}$ equations in \hat{F} unknowns. Given that nonnegative returns is the only restriction on revenue, a solution does not generally exist. Therefore, when demand interdependencies are present, the first price GDP slot auction does not permit a Nash equilibrium in which airlines bid their valuations.

A.2 Second Price

In addition to the notation introduced for first price auctions above, let $\bar{i}(Y, x)$ be the airline with the highest marginal revenue to acquiring Y under the profile of bids x that does not win Y .

Theorem 2. *There exist Nash equilibria of the second price GDP slot auction where airline revenues exhibit nonnegative returns that implement the revenue maximizing allocation.*

Proof. Suppose \hat{F} is the number of GDP slots. Then the set of NE of the second price GDP slot auction includes a nonempty $N^{\hat{F}}$ -dimensional polytope of profiles x^e that result in the efficient allocation.

Fix an allocation, $\hat{\mathcal{F}} = \{\hat{\mathcal{F}}_1, \dots, \hat{\mathcal{F}}_N\}$, that maximizes the sum of airline revenues. $\hat{\mathcal{F}}_i$ is the set of slots allocated to airline i for all $i = 1, \dots, N$. Profiles, x^e , that satisfy the following system of inequalities give rise to $\hat{\mathcal{F}}$.

$$\begin{aligned} \sum_{j \in \hat{\mathcal{F}}_i} x_i^e(j) &\geq \Delta R_{\bar{i}(\hat{\mathcal{F}}_i, x^e)}(\hat{\mathcal{F}}_i, x^e) \\ \sum_{j \in \hat{\mathcal{F}}_k} x_i^e(j) &\leq \Delta R_k(\hat{\mathcal{F}}_i, x^e) \end{aligned}$$

for all i and all $k \neq i$. It follows directly that profiles that satisfy the above give rise to $\hat{\mathcal{F}}$.

Next, we show that such profiles are also NE. The first inequality requires that every airline i bids so that the sum of its bids for each subset $\hat{\mathcal{F}}_i$ exceeds the second highest valuation for that subset. Then, the second inequality requires that for every other subset $\hat{\mathcal{F}}_k$ for $i \neq k$ airline i bids so that the sum of its bids does not exceed airline k 's valuation of that subset.

Together these conditions establish for all i ; (i) that the marginal revenue of $\hat{\mathcal{F}}_i$ is greater than or equal to the marginal cost of $\hat{\mathcal{F}}_i$, and (ii) that the marginal revenue of slots Y is less than or equal to the marginal costs of Y for all $Y \subseteq \hat{\mathcal{F}}$ not equal to $\hat{\mathcal{F}}_i$.

The fact that no airline i finds it worthwhile to trade any subset of $\hat{\mathcal{F}}_i$ for any subset of slots Y follows from the fact that $\hat{\mathcal{F}}$ maximizes revenue.

If $n(\hat{\mathcal{F}})$ represents the number of airlines awarded slots under $\hat{\mathcal{F}}$, then this $N^{\hat{F}}$ -dimensional polytope is described by $n(\hat{\mathcal{F}})N$ equations in $\hat{F}N$ unknowns. Since, $n(\hat{\mathcal{F}}) \leq \hat{F}$, this polytope exists generically. \square

Unlike other second price auctions, the second price GDP slot auction is not dominant-strategy implementable. The reason follows from the discussion of incentive compatibility above. Suppose a dominant strategy x_i^d existed for player i . Then

$$R_i(\hat{\mathcal{F}}_i(x_i^d, x_{-i}) | \mathcal{F}) \geq R_i(\hat{\mathcal{F}}_i(x_i, x_{-i}) | \mathcal{F})$$

for all $x_{-i} \in X_{-i}$ and all $x_i \in X_i$ with strict equality for some pairs (x_i, x_{-i}) . This means that x_i^d must be optimal for airline i when x_{-i} is such that airline i wins Y and at the

same time when x_{-i} is such that airline i does not win Y for all $Y \subseteq \hat{\mathcal{F}}$. Therefore, x_i^d must be the bid vector for which airline i truthfully reports its valuation of every subset of flights. However, as already shown in the discussion of incentive compatibility, this is not generally possible in the presence of demand interdependencies.

B Computational Issues

Here we briefly describe two computational issues that arise in PGT modeling. The first concerns a “density of states phenomenon” that arises as the complexity of the game grows. The second concerns the choice of QR-rationality likelihood so that Monte Carlo estimates converge.

In many cases it will be very difficult to have any idea what the space of mixed strategy profiles, $\Delta_{\mathcal{X}}$, looks like. In particular, it will be difficult to know how to efficiently sample this space so that we draw with high probability the types of q ’s that get high probability under $P(q|\mathcal{J})$. Therefore, we resort to a proposal distribution $H(\rho, \mu, \Sigma)$ that is roughly uniform over the set of mixtures of Gaussians. This can be very inefficient. In addition, as the complexity of the game in question grows, the inefficiency of a uniform proposal distribution grows.

A final issue concerns the convergence of the Monte Carlo estimator given the unbounded nature of the QR-rationality likelihood. That is, the parameter β can vary from $-\infty$ to ∞ . Therefore, if our Monte Carlo estimates of the posterior and its moments are to converge, then we must worry about the specific form of the likelihood function. As we established in the discussion of density of states above, the probability of drawing a q with high QR-rationality under the proposal distribution $H(\cdot)$ can be vanishingly small. So if the likelihood is not bounded above for large β , then the ratio

$$\frac{\mathcal{L}(\mathcal{J}|q^{\rho(t), \mu(t), \sigma(t)})}{H(q^{\rho(t), \mu(t), \sigma(t)})}$$

will diverge for q that give rise to large β . This would lead our Monte Carlo estimator to have infinite variance [see [Robert and Casella \(2004\)](#)]. Note that the likelihood in equation 5 is bounded above and below.