

UC Berkeley

Working Papers

Title

Reducing Auxiliary Power Unit Reliance through Decision-Sufficient Information: A Game-Theoretic Framework for Airport Gate Electrification

Permalink

<https://escholarship.org/uc/item/70271188>

Authors

Shahriar, Muhammad Shajid

Rakas, Jasenka

Publication Date

2026-10-01

Data Availability

The data associated with this publication are available upon request.



Berkeley
UNIVERSITY OF CALIFORNIA

Reducing Auxiliary Power Unit Reliance through Decision-Sufficient Information: A Game-Theoretic Framework for Airport Gate Electrification

Muhammad Shajid Shahriar
Jasenka Rakas, Ph.D.

University of California, Berkeley
National Center of Excellence for Aviation Operations Research Airport
Design Studio

October 2026

ABSTRACT

When parked at airport gates, commercial aircraft need electricity and conditioned air to power avionics, lighting, cabin services, and passenger comfort. Aircraft typically use an auxiliary power unit (APU), a small gas-turbine engine located at the rear of the aircraft that burns jet fuel as a power source. Ground power (GP) and preconditioned air (PCA) can reduce auxiliary power unit (APU) use during aircraft turnaround, but their availability alone does not ensure adoption. Reliability, coordination, timing and perceived delay risk can make continued APU reliance a rational airline response. This study develops a game-theoretic framework that links diagnosis of persistent APU reliance to the evidence needed for airport intervention. Each airline–gate pair is represented as a two-player airport–airline game. Four information packages are ordered by expected acquisition burden, with decision sufficiency defined as preserving both the equilibrium classification and selected intervention obtained under a full-information benchmark. Two synthetic demonstrations illustrate the framework. The first evaluates baseline equilibria and one-at-a-time intervention thresholds across 16 airline–gate pairs. Fourteen pairs initially remained in a weak-support/APU equilibrium. Coordination improvement shifted six pairs toward strong support and GP/PCA use, the broadest response among the five tested interventions. Equipment reliability had the smallest mean positive threshold among successfully shifted pairs. Intervention responses varied across pairs, with the broadest lever shifting 37.5% of all pairs within the tested range, supporting a pair-specific approach to intervention selection. The second demonstration uses 60 synthetic operation-style records across the same 16 pairs to illustrate bottleneck diagnosis and candidate-response mapping. Weak economic incentives relative to operational risk were the dominant modeled constraint for 10 pairs. Together, the demonstrations motivate a staged pathway that begins with existing operational records and adds technical or strategic evidence when needed to clarify intervention choice. The framework establishes a structured basis for targeted information acquisition and provides a pathway for empirical evaluation of decision-sufficient stopping packages, acquisition costs, and operational effectiveness.

Keywords: Airport gate electrification; Auxiliary power units; Ground power; Game theory; Decision-sufficient information; Targeted information acquisition.

1. INTRODUCTION

Commercial aircraft require electricity and conditioned air while parked at airport gates to support avionics, lighting, cabin services, and passenger comfort. An aircraft commonly operates its auxiliary power unit (APU), a small gas turbine engine that burns jet fuel and produces no useful thrust. APU operation remains necessary under certain technical and safety conditions, but extended routine use produces carbon dioxide, nitrogen oxides, carbon monoxide, hydrocarbons, and particulate matter near terminal buildings, ramp personnel, and surrounding communities. It also increases noise, fuel consumption and aircraft maintenance burden (Kinsey et al. 2012; Padhra 2018).

The location of these emissions is operationally important. Unlike cruise emissions, gate emissions occur repeatedly within a confined apron environment where aircraft, workers, passengers, ground-support equipment and terminal ventilation systems are in close proximity. A short period of APU operation may appear small when evaluated for one aircraft, but the cumulative effect becomes material when similar cycles occur across hundreds of daily turnarounds. The environmental burden is also unevenly distributed. Airlines generally purchase fuel and maintain the APU, while airport operators, ramp workers and nearby communities experience a substantial share of the local noise and air-quality impacts. This separation between who controls the operating decision and who experiences the external effects is one reason that technically efficient solutions may remain underused (Kinsey et al. 2012; Padhra 2018).

A technically mature alternative is already available at many airports. Fixed or mobile 400-Hz ground power (GP) supplies electricity, while preconditioned air (PCA) supplies heated or cooled cabin air. Together, these systems allow the APU to be shut down after arrival and restarted only when operationally required before departure. The Federal Aviation Administration recognizes gate electrification as an airport-emission-reduction strategy (Federal Aviation Administration 2021). Greer, Rakas, and Horvath (2021) estimated life-cycle greenhouse gas reductions of approximately 63% to 97% per gate operation and average payback periods of about one to two years, while Padhra (2018) reported APU-related emission reductions approaching 47.6% in the turnaround context analyzed.

Gate electrification is therefore a near-term measure that complements aviation's long-term net-zero objective without requiring aircraft-fleet replacement (International Civil Aviation Organization 2022). Its benefit still depends on the electricity source, equipment suitability, PCA capacity, and operating timing, and broader adoption can materially change airport electricity demand profiles (Sadati and Cetin 2020). The relevant comparison is therefore among specific combinations of equipment, energy source and turnaround practice, not simply APU versus ground power.

The operational problem is that installed infrastructure does not guarantee consistent use. GP/PCA may be unavailable, incompatible, damaged or slow to restore. Cables and hoses may be difficult to access, crew response may be late, and rapidly changing departure conditions may reduce the available connection window. Pilots and airlines may also retain or restart the APU early when they are uncertain that external systems will remain reliable or when a possible departure delay is perceived as more consequential than the fuel cost of conservative APU use. Analysis of more than 25,000 turnarounds at 125 airports found that departure-cycle APU emissions were approximately three times greater than arrival-cycle emissions, partly because uncertainty in departure timing encouraged early APU startup (Padhra 2018). ACRP guidance similarly identifies equipment design, maintenance, human factors, communication and utilization tracking as interdependent determinants of effective GP/PCA operation (National Academies of Sciences, Engineering, and Medicine 2019; Schmidt 2017).

These constraints are closely coupled. A cable that is functional but difficult to position may be recorded as available while remaining operationally unattractive. A maintenance program may achieve high

average availability but still fail during the short interval in which an aircraft needs service. A crew may connect GP correctly but disconnect it too early because the expected pushback time is unreliable. Conversely, a strict APU policy can increase operational tension if the airport cannot provide rapid support when external equipment fails. The same underuse indicator can therefore emerge from physical, procedural, behavioral or institutional causes. This is why a higher-resolution diagnosis is needed before an airport can rank potential actions.

Detailed monitoring has improved the visibility of this problem. Prior work at San Francisco International Airport showed that meter and operational records can identify GP use, APU duration, gate-level and airline-level variation, emissions and cost. These systems reveal where and when underuse occurs, but a similar observable pattern can arise from very different causes. Low GP utilization may reflect equipment unreliability, delayed connection, weak coordination, perceived delay risk, insufficient economic advantage or unclear accountability (Achatz Antonelli 2023; Rakas et al. 2025). Consequently, the airport may which intervention will change behavior.

Operational monitoring can support screening, while further validation is needed before its outputs guide intervention deployment (Pan et al., 2026). A performance dashboard may accurately flag a gate with long APU duration and low GP consumption. That evidence supports prioritization, yet it does not reveal whether the airport should replace equipment, improve preventive maintenance, revise crew responsibilities, provide better departure forecasts, introduce an incentive or strengthen accountability. Each action changes a different component of the airport-airline payoff structure. If the cause is incorrectly inferred, the airport can spend money while leaving the airline's preferred response unchanged. The analytical problem is therefore not only to predict energy use, but to identify the minimum strategic explanation required for intervention selection (Achatz Antonelli 2023; Howard 1966).

This creates an airport-airline dilemma. Airports benefit from lower emissions, improved infrastructure utilization and reduced community and worker exposure. Airlines also benefit from lower fuel use and APU maintenance, but they remain responsible for departure readiness and delay propagation. If GP/PCA is perceived as uncertain, APU reliance can remain the rational operational response even when it is environmentally and economically inferior at the system level. Game theory is well suited to this setting because the preferred action of one party depends on the operating conditions and expected response of the other. Transportation applications have used game-theoretic models to examine congestion and pricing, freeway interactions, environmental decisions in airline-airport systems and traffic-law reform (Nash 1950; Osborne and Rubinstein 1994; Levinson 2005; Cortes-Berruero, Gershenson, and Stephens 2016; Dixit et al. 2024; Shahriar 2026).

In the present setting, the airport acts as the designer of the operating environment rather than as a direct controller of the aircraft energy source. It can improve reliability, staffing, communication, monitoring or incentives, but the airline, pilot and ground operation ultimately respond to those conditions. A stable Weak/APU outcome can therefore persist without either party explicitly preferring higher emissions. The airport may be unwilling to incur support costs without confidence that airlines will use the system, while the airline may be unwilling to depend on GP/PCA without confidence that the support will be available. A game-theoretic formulation makes this reciprocal hesitation visible and allows the model to test which airport-controlled change is sufficient to alter the airline's best response.

A second problem follows directly from this strategic uncertainty: how much information should the airport acquire before acting? Many variables may be useful, including equipment incidents, connection attempts, maintenance response, crew timing, departure forecasts, pilot confidence, delay exposure, pricing and enforcement history. Collecting all of them may require sensors, database integration, access

agreements, manual observation, surveys, governance and continuing staff effort. Acting only on conventional dashboard measures, however, may lead the airport to treat a visible symptom rather than the binding cause. Replacing equipment will not solve a coordination failure, while stronger policy pressure may not work if the ground system is unreliable.

Value-of-information theory provides an economic basis for resolving this tension. Information is valuable because it improves a decision, not because it increases the volume of available data. Its importance depends jointly on uncertainty and the consequences of acting incorrectly (Howard 1966). More recent work also shows that information rankings can depend on risk attitude and the decision context in which the evidence is used (Abbas and Hazen 2025). Applied to gate electrification, a data source becomes strategically valuable when it changes the bottleneck diagnosis, changes the preferred intervention or reduces the expected loss from intervention mismatch.

This definition is narrower than the conventional definition of data completeness. An equipment-failure log may be highly valuable for a pair that appears close to a reliability-driven equilibrium shift, but add little value for a pair already performing well. A pilot-confidence survey may be necessary when technical records cannot distinguish delay aversion from weak economic incentive, yet unnecessary when a repeated hardware outage is already evident. Decision sufficiency is consequently pair-specific and action-specific. The airport should not ask whether a variable is generally useful; it should ask whether resolving that uncertainty is likely to change what it does next.

The bottleneck-to-evidence mapping converts that question into an operational data request. A reliability diagnosis calls for incident records, outage durations, failed connection attempts, and repair response times. A timing diagnosis calls for connection and disconnection timestamps, turnaround milestones, departure forecasts and schedule slack. A coordination diagnosis calls for crew response, communication and responsibility assignment. Economic or delay-risk diagnoses require evidence of cost, maintenance, operational confidence, and accountability. This mapping narrows the scope of the collection task and makes the rationale for each requested dataset auditable.

The reviewed literature provides a strong foundation in GP/PCA engineering, gate-energy monitoring, strategic interaction, and information value, but these streams remain largely separate. Airports, therefore, lack a practical method that jointly determines what they should do and what they need to know before doing it. This study addresses that gap by combining pair-specific game-theoretic diagnosis with a burden-ordered information-acquisition rule. A full-information game provides the benchmark diagnosis and intervention, while lower-burden information packages are evaluated according to whether they reproduce that decision. Synthetic data are used to test the decision logic, not to claim airport-specific field estimates. The study therefore links descriptive monitoring, strategic diagnosis, and burden-ordered evidence acquisition into a single decision sequence. Accordingly, this study pursues two objectives:

- i. To develop a resource-efficient game-theoretic framework for diagnosing the strategic cause of persistent APU reliance and selecting an appropriate airport intervention under incomplete information.
- ii. To develop a procedure for identifying the lowest burden, decision-sufficient information set needed for reliable intervention selection while minimizing unnecessary data acquisition and the risk of intervention mismatch.

2. METHODS

The study applied a decision-first design, in which additional data were treated as an intermediate action rather than the final research objective. The unit of analysis was the airline-gate pair because GP/PCA use can vary substantially across airlines, gates, aircraft, and individual turnarounds. A separate strategic game was therefore constructed for each pair rather than estimating a single airport-wide equilibrium.

The framework was demonstrated through two connected analytical components. Demonstration A included 16 synthetic airline-gate pairs formed from four airlines and four gates. It tested whether Weak/APU lock-in could emerge under controlled assumptions and whether selected airport-controlled levers could shift the interaction toward Strong/GP. Demonstration B included 60 operation-style records distributed across the same 16 pairs and illustrated how operational, reliability, coordination, economic and airline-risk inputs could be translated into bottleneck diagnoses. These deliberately constructed scenarios provide conceptual and computational validation, not field validation for a specific airport (Sargent 2013).

Demonstration A preserved heterogeneity across GP utilization, equipment reliability, coordination quality, timing performance, monitoring strength, policy pressure, and airline delay sensitivity. Demonstration B added operation-style variation within the same airline-gate structure. Each record contained a pair identifier and inputs representing GP utilization, APU duration, equipment condition, coordination, stakeholder weights and airport-controlled actions. Records were aggregated to the pair level before bottleneck classification. Arithmetic means summarized GP utilization and APU duration, while strategic inputs were standardized before payoff construction. This design separates equilibrium and perturbation analysis from structured bottleneck diagnosis.

2.1 Strategic Variables and Information Packages

A focused review of ACRP Research Report 207 and prior gate-power research was used to organize relevant variables into three groups: (1) airport-controlled conditions, including equipment reliability, maintenance responsiveness, coordination, monitoring, pricing and policy pressure; (2) airline-side conditions, including fuel and maintenance savings, delay risk, operational confidence and uncertainty cost; and (3) shared operational measures, including GP/PCA use, APU duration, turnaround time, connection timing, aircraft characteristics and departure uncertainty (National Academies of Sciences, Engineering, and Medicine 2019; Padhra 2018; Achatz Antonelli 2023; Rakas et al. 2025).

Existing monitoring may omit information needed to explain persistent APU reliance. We refer to these unmonitored but strategically relevant variables as missing strategic evidence. Existing dashboards may have broad descriptive coverage while still omitting the evidence needed to distinguish among competing causes. Candidate evidence was therefore organized into the progressive packages shown in **Table 1**. Package burden reflects the expected effort for access, integration, and collection rather than the number of variables. The packages form an ordinal burden sequence, $D_0 < D_1 < D_2 < D_3$, and should not be interpreted as estimates of monetary acquisition cost.

This distinction is important because a database with many routinely captured fields may be less burdensome than a database with a new behavioral variable. D_0 primarily uses records that many airports or airlines already maintain. D_1 may require integration of maintenance systems, unsuccessful connection attempts, and turnaround milestones that exist but are not linked to gate-energy performance. D_2 is potentially more burdensome because confidence, perceived delay exposure, and stakeholder weights may require interviews, surveys, or targeted observation. D_3 is not proposed as the default operating requirement; it is the benchmark against which lower-burden decisions can be evaluated.

TABLE 1 Information Packages Used in the Decision Framework

Information Package	Principal Inputs	Primary Decision Supported	Relative Acquisition Burden
D_0 : Existing operational screen	Gate, airline, aircraft, block-in/out, GP/PCA use, APU time, delay and cost	Identify underperforming and already-aligned pairs	Negligible when records exist
D_1 : Technical/process diagnosis	Incidents, failed connections, outages, maintenance response, crew timing and milestones	Distinguish reliability, timing and coordination	Low to moderate
D_2 : Strategic confirmation	Delay risk, confidence, stakeholder weights, pricing, monitoring and accountability	Distinguish economic, behavioral and policy constraints	Moderate
D_3 : Full-information benchmark	All operational, technical, behavioral and policy inputs	Establish benchmark for unresolved cases and future validation	Highest; validation only

2.2 Decision-Sufficient Information Rule

For airline-gate pair i and information package D , define the decision output as

$$z_i(D) = [NE_i(D), k_i(D)] \quad (1)$$

where i indexes the airline-gate pair, $D \in \{D_0, D_1, D_2, D_3\}$ denotes the information package being evaluated and $z_i(D)$ is the resulting decision output for pair i under package D . The term $NE_i(D)$ is the Nash-equilibrium classification and $k_i(D)$ is the selected airport lever. Let Q denote the full set of strategically relevant information, operationally represented by the full-information benchmark package D_3 . Package D is decision-sufficient relative to Q when

$$z_i(D) = z_i(Q) \quad (2)$$

Here, $z_i(Q)$ denotes the benchmark decision output obtained under full information. Equality in **Equation 2** requires both components of the decision output, $NE_i(D)$ and $k_i(D)$, to match their corresponding full-information values, $NE_i(Q)$ and $k_i(Q)$.

Table 1 orders the information packages from lower to higher expected acquisition burden: D_0, D_1, D_2 and D_3 . The aim is to identify the first package in this sequence that produces the same equilibrium classification and selected airport intervention as the full-information benchmark D_3 in **Equation 3**:

$$D_i^{LB} = \text{first package } D \text{ in } (D_0, D_1, D_2, D_3) \text{ for which } z_i(D) = z_i(D_3) \quad (3)$$

Here, D_i^{LB} denotes the lowest-burden decision-sufficient package for airline-gate pair i . “First” refers to the stated burden order, not to a comparison of monetary costs. The benchmark D_3 represents the full information set Q used in **Equation 2**.

Operationally, the procedure begins with D_0 . If the pair is already Strong/GP, acquisition stops. If it remains Weak/APU, D_1 is evaluated and the procedure can stop when a clear low-threshold reliability, timing or

coordination lever emerges. D_2 is requested only when competing strategic explanations remain plausible, while D_3 is reserved as the full-information benchmark.

For a pair requiring further diagnosis, a candidate information set combines existing operational records with additional evidence relevant to the suspected bottleneck:

$$D_i^{candidate} = D_0 \cup E_i \quad (4)$$

Here, D_0 contains the existing operational records and E_i contains the additional evidence collected to distinguish plausible bottlenecks for airline–gate pair i . The union symbol $\cup$ means that these information sets are combined. **Equation 4** describes how a candidate information set is assembled; it does not establish that the set is sufficient or has the lowest acquisition burden. Those properties require evaluation against the benchmark and package order specified in **Equations 2–3**.

The equality criterion is intentionally decision-based. A partial package does not need to reproduce every payoff component or descriptive statistic. It is sufficient when it preserves the equilibrium classification and selected intervention. Future empirical applications can extend this exact-agreement rule by allowing a tolerance for decision regret or confidence. **Figure 1** summarizes the analytical sequence and its stopping rule.




Figure 1 Decision-first framework and proposed burden-ordered information pathway

2.3 Game-Theoretic Model

The airport-airline interaction was represented as a two-player non-cooperative game $G = (N, S, U)$, where $N = \{P, L\}$ includes the airport P and airline L . The airline selected between retaining APU reliance (A) and using GP/PCA during the applicable portion of turnaround (G). The airport did not directly select the aircraft energy source; it shaped the operating environment through reliability, equipment availability, coordination, monitoring, pricing and accountability. For the baseline 2 x 2 game, that environment was simplified to weak support (W) or strong support (S), producing four profiles: (W,A) , (W,G) , (S,A) and (S,G) . Airport utility U_P , where P denotes the airport, balanced emission-reduction, infrastructure-utilization and transparency benefits against operating-support, implementation and reliable-power-provision costs:

$$U_P = \alpha_1 B_{em} + \alpha_2 B_{ut} + \alpha_3 B_{tr} - \alpha_4 C_{op} - \alpha_5 C_{im} - \alpha_6 C_{ps} \quad (5)$$

In **Equation 5**, B_{em} is the emission-reduction benefit, B_{ut} is the infrastructure-utilization benefit and B_{tr} is the transparency benefit. The terms C_{op} , C_{im} and C_{ps} denote the operating-support cost, implementation cost and reliable-power-provision cost, respectively. The coefficients $\alpha_1, \dots, \alpha_6$ are airport-side weights representing the relative importance assigned to the corresponding benefit and cost components.

Airline utility U_L , where L denotes the airline, balanced fuel-savings, maintenance-savings and sustainability or reputational benefits against perceived delay-risk, coordination and uncertainty-related costs:

$$U_L = \beta_1 B_{fs} + \beta_2 B_{ms} + \beta_3 B_{sr} - \beta_4 C_{dr} - \beta_5 C_{cc} - \beta_6 C_{ur} \quad (6)$$

In **Equation 6**, B_{fs} is the fuel-savings benefit, B_{ms} is the maintenance-savings benefit and B_{sr} is the sustainability or reputational benefit. The terms C_{dr} , C_{cc} and C_{ur} denote the perceived delay-risk cost, coordination cost and uncertainty-related cost, respectively. The coefficients $\beta_1, \dots, \beta_6$ are airline-side weights representing the relative importance assigned to the corresponding benefit and cost components.

The categories and signs were grounded in prior GP/PCA and turnaround research, while numerical magnitudes in this demonstration were synthetic. Components were linearly standardized to $[-10, +10]$ before aggregation so that benefits and costs could be represented on a common utility scale. For raw component x_{ij} , representing component j for pair i , the transformation was

$$\tilde{x}_{ij} = -10 + 20 \left[\frac{x_{ij} - x_{j,min}}{x_{j,max} - x_{j,min}} \right] \quad (7)$$

Here, x_{ij} is the raw value of component j for pair i , $\tilde{x}_{ij}$ is its standardized value and $x_{j,min}$ and $x_{j,max}$ are the minimum and maximum assigned values of component j , respectively. The constants -10 and 20 transform the original component range to the common scale $[-10, +10]$.

The transformation preserved the ordering of assigned values while preventing variables expressed on different conceptual scales from dominating solely because of their units. Standardized components were then combined into a 2×2 payoff matrix containing the airport and airline utility for each strategy profile. The numerical values should be interpreted as comparative utilities within the synthetic environment, not as monetary estimates of airport or airline welfare.

A pure-strategy Nash equilibrium was identified by direct best-response comparison across the four profiles. A pair was classified as Weak/APU when (W,A) was stable and Strong/GP when (S,G) was stable. The model retained the ability to identify alternative profiles if they emerged. The airport intervention

process was then treated sequentially: the airport changed the operating environment first and the airline selected its best response to the revised conditions (Fudenberg and Tirole 1991).

This airport-first sequence is consistent with the institutional structure of the problem. The airport can invest in reliability, change a process or introduce a policy before observing whether the airline alters APU behavior. For each feasible airport intervention, the airline compares the utilities of continued APU reliance and GP/PCA use and selects the option with the higher utility. The airport anticipates this response when choosing an intervention.

The formulation does not imply that the airport has complete control or that all airlines respond identically. It formalizes the narrower claim that airport-controlled conditions can change the payoff comparison under which the airline chooses between APU and GP/PCA.

2.4 Intervention Perturbation, Thresholds and Universal No-Shift Rate

Five airport-controlled levers were evaluated separately: equipment reliability, timely performance, coordination quality, monitoring transparency and policy pressure. One-at-a-time perturbation was used as a local, interpretable screening method while holding other conditions constant. Because this method does not capture interactions among simultaneous or bundled interventions, the resulting thresholds are local leverage points rather than a global sensitivity ranking (Hamby 1994; Saltelli et al. 2008). For lever k and pair i , the minimum tested positive change that altered the equilibrium was

$$\Delta_{ik}^* = \min\{\Delta_k > 0: NE[M_i(x_k + \Delta_k)] \neq NE[M_i(x_k)]\}. \quad (8)$$

In **Equation 8**, i indexes the airline-gate pair and k indexes the airport-controlled intervention lever. The term x_k is the baseline value of lever k , Δ_k is a tested positive change applied to that lever and $M_i(\cdot)$ is the payoff matrix for pair i under the specified lever value. The operator $NE[\cdot]$ returns the Nash-equilibrium classification of the corresponding payoff matrix, while Δ_{ik}^* is the minimum tested positive change in lever k that changes the equilibrium for pair i . The superscript $*$ identifies the minimum successful threshold rather than multiplication.

A zero-value indicated that the pair was already Strong/GP and was excluded from intervention-success counts and mean positive thresholds. For Weak/APU pairs, the lowest positive successful threshold was selected as the pair-specific lever.

For universal application, let k index an airport-controlled lever and let n_k denote the number of the 16 airline-gate pairs shifted by lever k . The success rate was:

$$S_k = \frac{n_k}{16} \quad (9)$$

where S_k is the success rate of lever k , n_k is the number of pairs shifted by that lever and 16 is the total number of airline-gate pairs evaluated. The corresponding universal no-shift rate is $R_k = 1 - S_k$, where R_k represents the share of all pairs that did not shift under lever k within the tested perturbation range.

This measure does not distinguish a mismatched lever from a correct lever applied at insufficient magnitude and it is not an empirical misclassification probability. Intervention breadth and mean positive threshold were therefore reported separately. A smaller normalized threshold does not automatically imply a lower monetary implementation cost in field application.

2.5 Structured Bottleneck Diagnosis and Reporting Boundaries

Let B denote the set of candidate bottlenecks:

$B = \{\text{weak economic incentive relative to operational risk, timing uncertainty, equipment reliability weakness, airline delay risk, coordination gap}\}.$

Let g_{ib} denote the standardized strategic deficit associated with bottleneck b for airline-gate pair i . The dominant modeled bottleneck was

$$b_i^* = \arg \max_{b \in B} g_{ib}. \quad (10)$$

In **Equation 10**, b indexes a candidate bottleneck in set B , g_{ib} is the standardized strategic deficit associated with bottleneck b for pair i and b_i^* is the dominant modeled bottleneck selected for that pair. The operator $\arg \max$ over $b \in B$ identifies the bottleneck with the largest value of g_{ib} among all candidates in B . This rule identifies the largest modeled constraint on the relative payoff of GP/PCA; it does not imply that all other constraints are absent.

The bottleneck set and the Demonstration A lever set serve different roles and are not intended to be identical. Reliability, timing and coordination map directly across the two demonstrations. Weak economic incentives and airline delay risk are airline-side constraints for which candidate airport responses include pricing, incentive support, and operational assurance. Monitoring transparency and policy pressure are airport-side enabling levers rather than separate bottleneck labels. The economic and delay-risk responses were not threshold-perturbed in Demonstration A, so Demonstration B reports them as candidate responses rather than quantitatively validated intervention recommendations.

Computation included payoff construction, direct equilibrium enumeration, perturbation, aggregation and figure generation. Because the scenarios were deliberately constructed, no null-hypothesis tests, confidence intervals or population-level probability claims were applied. Internal checks focused on best-response reproduction, one-lever-at-a-time changes and agreement between pair-level and summary outputs. Airport-specific probabilities, causal effects, monetary information values and operational savings require empirical calibration.

3. RESULTS AND DISCUSSION

The baseline analysis showed a strongly unbalanced strategic condition. Fourteen of the 16 airline-gate pairs, or 87.5%, remained in Weak/APU, while two pairs, or 12.5%, were positioned in Strong/GP. **Figure 2a** presents this distribution. The result demonstrates that APU reliance can remain stable even when cleaner infrastructure is assumed to be available. Under the assigned payoff conditions, reliability uncertainty, coordination burden, timing risk or insufficient support made APU operation the airline's rational fallback. The two Strong/GP pairs are equally important because they show that the model does not mechanically classify every operation as problematic. The first decision is therefore whether intervention is needed, not which intervention should be imposed. No pair was classified in either mixed profile, (W,G) or (S,A) , under the tested synthetic assignments.

This is a nontrivial resource-saving result. A conventional improvement program may begin by identifying a low-utilization gate and immediately assigning corrective action. The equilibrium screen first separates operations where the airport-airline incentives are already aligned from those where APU reliance remains strategically stable. Preserving that distinction prevents additional sensors, staff attention or

enforcement from being directed toward pairs that do not require them. It also provides a baseline against which later deterioration can be detected.

The heatmap in **Figure 2b** shows substantial variation in GP utilization across airline-gate combinations. AA at Gate A1 exhibited comparatively high utilization, while DL at Gate B5 was among the weakest combinations. This heterogeneity confirms that an airport-wide average can locate an overall performance problem but cannot diagnose its cause. Similar APU use may arise from failed equipment in one pair, late crew response in another, and a strong perception of departure risk in a third. Pair-specific diagnosis is therefore necessary before resources are assigned. This disaggregation principle is consistent with transportation-finance research showing that aggregate system measures can conceal substantial spatial and population-level differences (Shahriar et al. 2026).

The heterogeneity also suggests that the airport should preserve the joint airline-gate identity in analysis. Gate-only reporting can hide airline practices, while airline-only reporting can hide local infrastructure and operational conditions. The pair is the smallest practical unit that captures both sides of the strategic interaction without requiring a separate game for every individual flight. Operation-level records remain useful for constructing the pair scores and detecting temporal patterns, but intervention selection is made at the level where airport-controlled conditions and airline behavior repeatedly interact.




Figure 2 (a) Baseline equilibrium distribution and (b) synthetic GP-utilization heterogeneity across airline-gate pairs

3.1 Universal Intervention Mismatch, Breadth and Threshold

The perturbation analysis showed that no airport-controlled lever was sufficient for the full set of pairs. Coordination quality had the greatest breadth, producing positive shifts in six Weak/APU pairs, followed by policy pressure with five. Timely performance and monitoring transparency each shifted four pairs, while equipment reliability produced positive shifts in two Weak/APU pairs. Two additional reliability entries had zero delta because those pairs were already Strong/GP and were excluded from success counts. As shown in **Table 2** and **Figure 3a**, the corresponding universal no-shift rates ranged from 62.5% to 87.5%. Even the broadest lever therefore left most pairs unchanged.

TABLE 2 Universal Intervention Performance and Mean Positive Threshold

Airport-Controlled Lever	Positive Shifts Among Weak/APU Pairs	Success Rate	Universal No-Shift Rate	Mean Positive Threshold
Coordination quality	6	37.5%	62.5%	0.133
Policy pressure	5	31.3%	68.8%	0.170
Equipment reliability	2	12.5%	87.5%	0.100
Timely performance	4	25.0%	75.0%	0.125
Monitoring transparency	4	25.0%	75.0%	0.125




Figure 3 (a) Universal no-shift performance, (b) intervention breadth versus mean positive threshold and (c) dominant bottleneck classification.

These universal no-shift rates provide the clearest argument for targeted information. Additional evidence has value when it helps the airport avoid applying a lever that does not change the airline's modeled response. A no-shift outcome does not prove that the lever was incorrect; it may also indicate that

the tested magnitude was insufficient. The metric therefore describes performance within the tested range rather than an intervention-selection error.

If coordination improvement were applied universally, 10 of 16 pairs would absorb some combination of process redesign, training or staff effort without reaching the intended equilibrium shift. For policy pressure, 11 pairs would remain unchanged; for timing and monitoring, 12 would remain unchanged; and for reliability, 14 would not produce a positive shift from Weak/APU within the tested range. A complete economic model would attach different implementation costs and consequences to these outcomes. Even without those monetary values, the high no-shift shares demonstrate why identifying the bottleneck before implementation can have value.

Intervention breadth and positive threshold produced different rankings. Coordination shifted six Weak/APU pairs and had a mean positive threshold of 0.133. Policy pressure shifted five pairs but required the largest mean positive threshold, 0.170. Timely performance and monitoring transparency each shifted four pairs with mean positive thresholds of 0.125. Equipment reliability produced positive shifts in two Weak/APU pairs and had the smallest mean positive threshold, 0.100. Two additional reliability entries with zero delta represented pairs already Strong/GP and were excluded from both success counts and threshold averages. **Figure 3b** should therefore be read as a comparison of breadth and local intervention intensity, not as a monetary cost ranking.

Pair-specific examples clarify the practical meaning. AA at A1 and AA at B2 required no further action. AA at B5 and DL at B2 shifted with a 0.10 reliability improvement, while AS at B5 and UA at B5 required larger coordination changes of 0.25 and 0.35. These outcomes represent three distinct management states: no action, a low-threshold targeted action and a larger intervention that may justify further feasibility assessment. Treating them equally would waste both monitoring and implementation resources.

The threshold results also help organize implementation sequencing. Low-threshold pairs can be treated as early opportunities for pilot deployment because a modest, observable improvement may be sufficient to establish operational confidence. Higher-threshold pairs should not automatically be abandoned, but they may require bundled actions or a stronger business case. For example, a large coordination threshold may indicate that training alone is insufficient without responsibility clarification, milestone integration or rapid-response support. The framework therefore supports a phased portfolio rather than an all-or-nothing airport-wide program.

The broader interpretation is that APU reliance can be rational without implying resistance to sustainability. The APU is an onboard fallback controlled by the aircraft and airline. GP/PCA depends on alignment among infrastructure, maintenance, personnel, communication and timing. If any element is perceived to threaten departure readiness, the airline can assign a high penalty to uncertainty and remain in Weak/APU. The airport must therefore change the condition that makes the APU strategically safer rather than simply restating the environmental advantage of GP/PCA.

This point explains why monitoring, incentives and engineering should be considered together. A dashboard can make poor performance visible, but visibility changes behavior only when responsibility and response are credible. An incentive can improve the airline payoff, but its effect may be overwhelmed by uncertainty if equipment or timing remains weak. Reliability investment can reduce uncertainty, yet its benefit may be lost if crews are not present or if disconnection procedures conflict with departure readiness. The model's value is not in declaring one category universally superior; it is in identifying which combination is most likely to change the best response for a specific pair.

3.2 Bottleneck Diagnosis and Candidate Evidence

Demonstration B included 60 operation-style records across the 16 pairs. Average GP utilization was 42.6% and average modeled APU time was 58.1 minutes per operation. Weak direct economic incentive relative to operational risk was classified as the dominant bottleneck for 10 pairs, or 62.5%. Timing uncertainty and equipment reliability weakness each affected two pairs, while airline delay risk and coordination gap each affected one pair. **Figure 3c** summarizes the distribution, while **Table 3** links each diagnosis with candidate confirmatory evidence and a mapped airport response.

TABLE 3 Bottleneck Classification, Candidate Confirmatory Evidence and Response Mapping

Dominant Bottleneck	Pairs (share)	Evidence needed to confirm diagnosis	Candidate airport response
Weak economic incentive relative to operational risk	10 (62.5%)	Fuel/electricity costs, APU maintenance, incentive structure and delay-risk weights	Pair-specific pricing, incentive support or accountability; threshold not estimated in Demonstration A
Timing uncertainty	2 (12.5%)	Connection/disconnection times, turnaround milestones, forecasts and schedule slack	Predictive timing support and departure-readiness-safe procedures
Equipment reliability weakness	2 (12.5%)	Failure incidents, outage duration, failed connections and repair-response time	Preventive maintenance and rapid-response protocols
Airline delay risk	1 (6.25%)	Delay propagation, pilot/crew confidence and prior GP/PCA experience	Operational assurance and risk mitigation; threshold not estimated in Demonstration A
Coordination gap	1 (6.25%)	Crew response, responsibility assignment, communication and process alignment	Coordination protocols and responsibility clarification

Demonstration A and Demonstration B answer different questions. Demonstration A tests how selected airport-controlled levers change equilibrium under assigned conditions, whereas Demonstration B identifies the modeled binding constraint when structured inputs are available. Accordingly, the two category sets overlap only where a bottleneck has a directly corresponding airport lever. Reliability, timing and coordination are directly linked. Economic incentive and airline delay risk require candidate responses that were mapped but not threshold-tested in this paper. Monitoring transparency and policy pressure remain enabling levers that may support more than one bottleneck. Demonstration B should therefore be interpreted as diagnosis-and-response mapping, not as a complete threshold analysis for every diagnosed category.

The prevalence of the economic category should not be interpreted as evidence that airlines prefer consuming fuel. It indicates that the direct fuel or maintenance advantage of GP/PCA was not strong enough to overcome assigned uncertainty and delay exposure. In practice, cost interacts with reliability, timing and operational confidence. A nominally attractive GP/PCA option may still lose in the payoff comparison when the perceived consequence of a late departure is large.

The combined results reveal three different forms of strategic importance. Bottleneck prevalence identifies how often a constraint appears. Intervention breadth identifies how many pairs respond to an intervention. The intervention threshold indicates how much change is required before the equilibrium shifts. The most frequent underlying strategic constraint was economic, the broadest intervention was

coordination and the lowest positive pair-specific threshold was reliability. These are not contradictory findings. They show why airports should not allocate resources only according to the most common category.

For a resource-constrained airport, the distinction is consequential. A frequency-based response may prioritize pricing or accountability because the economic bottleneck appears most often. A decision-based response first asks whether a smaller reliability or coordination improvement can shift a specific pair. Maintenance and coordination evidence may therefore have high value even when another bottleneck is more prevalent. This is the central value-of-information principle in the framework: evidence is valuable when it changes the preferred action or prevents a consequential mismatch.

3.3 Proposed Adaptive Information Pathway and Stopping Rule

The current analyses motivate but do not empirically verify a unique stopping package for each airline-gate pair. The proposed sequence begins with D_0 , which screens existing operational records and identifies pairs already aligned with Strong/GP. D_1 adds equipment incidents, failed connections, maintenance response, crew timing and turnaround milestones to distinguish reliability, timing and coordination constraints. D_2 adds perceived delay risk, operational confidence, stakeholder weights, pricing, monitoring and accountability when D_1 does not stabilize the diagnosis. D_3 serves as the full-information benchmark for unresolved cases and future validation.

The pathway has three conceptual stopping conditions. First, a pair already classified as Strong/GP stops at D_0 without additional intervention. Second, a clear low-threshold technical or coordination lever can justify action after D_1 . Third, when multiple explanations remain plausible, D_2 evidence is required before a larger policy, incentive or monitoring response is recommended. Resource efficiency therefore means advancing only when the next package can plausibly change the decision, not claiming that the fewest variables are always sufficient.

In field implementation, the rule would be applied iteratively. Following a targeted intervention, the airport would continue monitoring GP/APU outcomes and update the pair-level game. A successful shift would provide evidence that the diagnosed bottleneck was actionable, while a failed shift could indicate an omitted constraint, insufficient intervention magnitude or changing behavior. The pair could then advance to the next information package rather than triggering comprehensive data collection for the entire airport.

Operationally, airports can screen existing GP/PCA, APU, airline, gate and turnaround records, link only underperforming pairs to maintenance and milestone systems and reserve interviews, surveys or new observation for unresolved cases. D_3 should be used for validation rather than as the default data requirement. This sequence avoids both comprehensive monitoring before decision relevance is established and costly intervention based only on descriptive GP/APU indicators.

3.4 Implementation Priorities for Airport Practice

A practical pilot should begin with a limited set of airline-gate pairs rather than an airport-wide mandate. The airport can select one already-aligned pair as a reference, one low-threshold technical case and one coordination-sensitive case. This structure allows managers to compare no action, a focused engineering response and a process intervention while using the same operational screen. The pilot should define in advance the decision that each new data source is expected to support, the condition for moving to the next information package and the period over which the intervention will be evaluated.

Implementation should also separate corrective action from accountability. Equipment incidents and failed connections require rapid technical response, while repeated avoidance of reliable GP/PCA may justify stronger monitoring or incentive measures. Combining these conditions too early can reduce stakeholder trust because airlines may perceive enforcement as unfair when the airport has not first demonstrated dependable service. A staged approach therefore gives reliability and process assurance priority before larger pricing or policy changes, unless the available evidence clearly identifies an economic or accountability bottleneck.

Finally, the airport should document both successful and unsuccessful shifts. Success provides an empirical basis for stopping additional data acquisition and potentially transferring the intervention to similar pairs. Failure identifies the value of the next information package. This learning structure turns implementation into a sequence of testable decisions rather than a one-time infrastructure or policy program.

3.5 Interpretation Boundaries

All equilibrium frequencies, intervention counts, bottleneck distributions and thresholds are outputs of controlled synthetic assignments and are not field estimates for a specific airport. Information packages are ordered by relative burden rather than monetary cost, and no exhaustive testing of all information-package combinations was conducted. The paper therefore specifies a decision-sufficiency criterion and proposed stopping rule but does not claim an empirically verified lowest-burden package for each pair. Economic-incentive and airline-delay-risk responses were mapped as candidate actions but were not threshold-perturbed in Demonstration A. In addition, one-at-a-time perturbation does not capture interactions among bundled levers.

The two-player model also compresses a turnaround system involving pilots, ramp crews, maintenance staff, airport operations and regulators into an airport-airline interaction. Future calibration should distinguish airline organizational incentives from pilot and ground-crew confidence, including on-time-departure risk and trust following successful or failed GP/PCA experiences.

Within these boundaries, the contribution is a structured sequence: monitoring identifies where APU dependence occurs, game-theoretic diagnosis identifies why the condition can remain stable and burden-ordered evidence mapping specifies what should be examined before action.

4. CONCLUSIONS

This study developed a resource-efficient game-theoretic framework for airport gate-electrification decisions under incomplete information. Persistent APU use was treated as a strategic condition in which airport support and airline perceptions of reliability, coordination, timing and delay risk jointly determine whether GP/PCA becomes the preferred operational choice. Additional evidence was treated as an intermediate decision input rather than an end in itself. Fourteen of the 16 synthetic airline-gate pairs were classified as Weak/APU and two as Strong/GP under the assigned baseline conditions. Coordination produced the greatest intervention breadth, while equipment reliability had the smallest mean positive threshold, 0.100, among shifted Weak/APU cases. No universal lever shifted more than 37.5% of pairs and universal no-shift rates ranged from 62.5% to 87.5%. Demonstration B separately showed that bottleneck prevalence, intervention breadth and threshold are different decision criteria and should not be collapsed into one airport-wide ranking.

The principal contribution is a proposed burden-ordered decision pathway, not a monetary least-cost estimate or a universally dominant intervention. Airports can begin with existing operational records, add technical or strategic evidence when the current package does not stabilize the diagnosis and reserve full-information collection for validation. The present paper defines that decision-sufficiency criterion and demonstrates the equilibrium and bottleneck logic, while pair-level testing of alternative information packages and monetary cost optimization remain for empirical validation.

Accordingly, the practical proposition is methodological: move from monitoring to pair-specific diagnosis, then to targeted evidence acquisition and intervention, rather than treating all observed underuse as the same problem. This sequence is intended to reduce unnecessary monitoring and intervention mismatch while preserving operational reliability.

The results validate internal framework behavior rather than airport-specific magnitudes. The scenarios were synthetic, information-package burden was ordinal rather than monetary and exhaustive portfolio ablation was not performed. Economic and airline-delay-risk responses were mapped but not threshold-tested, one-at-a-time perturbation did not capture interactions among bundled levers and the two-player representation simplified the roles of pilots, ramp crews, maintenance staff and airport operations.

Future research should calibrate the framework using operation-level GP/PCA and APU records, equipment-reliability logs, turnaround milestones and targeted stakeholder interviews. A systematic information-package comparison should evaluate decisions based on D_0, D_1 and D_2 against a full-information benchmark to estimate agreement, regret, diminishing information value and the empirically lowest-burden package. Monetary estimates of integration, instrumentation, staff time, survey burden and governance would then support a cost-regret frontier.

Multi-actor and repeated-game extensions should distinguish airport decisions, airline incentives, pilot confidence and ground-crew coordination and examine how trust changes after successful or failed GP/PCA experiences. Validation across multiple airports would identify which information packages are transferable and which require local calibration.

5. ACKNOWLEDGMENTS

The authors gratefully acknowledge Pietro Achatz Antonelli for prior modeling work that provided an important foundation for this study. OpenAI ChatGPT was used for grammar editing, language clarification, and code debugging. All model-assisted outputs were reviewed and verified by the authors for consistency and contextual accuracy.

This work was supported by the National Center of Excellence for Aviation Operations Research and the Airport Design Studio at the University of California, Berkeley.

9. REFERENCES

- Abbas, A. E., and G. B. Hazen. 2025. "On the Value of Information Across Decision Problems." *Decision Analysis* 22 (1): 1-13. <https://doi.org/10.1287/deca.2024.0187>.
- Achatz Antonelli, P. 2023. A Framework for Monitoring, Modeling, and Management of Airport Ground Power Systems During Aircraft Turnaround Operations. Doctoral dissertation, University of California, Berkeley. <https://escholarship.org/uc/item/0dw9171k>.
- Cortes-Berruero, L. E., C. Gershenson, and C. R. Stephens. 2016. "Traffic Games: Modeling Freeway Traffic with Game Theory." *PLOS ONE* 11 (11): e0165381. <https://doi.org/10.1371/journal.pone.0165381>.
- Dixit, A., T.-M. Choi, P. Kumar, and S. K. Jakhar. 2024. "Roles of Reciprocity and Fairness Concerns in Airline-Airport Systems with Environmental Considerations." *European Journal of Operational Research* 312 (3): 1011–1023. <https://doi.org/10.1016/j.ejor.2023.07.016>.
- Federal Aviation Administration. 2021. *United States Aviation Climate Action Plan*. Washington, DC: U.S. Department of Transportation.
- Fudenberg, D., and J. Tirole. 1991. *Game Theory*. Cambridge, MA: MIT Press.
- Greer, F., J. Rakas, and A. Horvath. 2021. "Reduce Aviation's Greenhouse Gas Emissions through Immediately Feasible and Affordable Gate Electrification." *Environmental Research Letters* 16 (5): 054039. <https://doi.org/10.1088/1748-9326/abf7f1>.
- Hamby, D. M. 1994. "A Review of Techniques for Parameter Sensitivity Analysis of Environmental Models." *Environmental Monitoring and Assessment* 32: 135-154. <https://doi.org/10.1007/BF00547132>.
- Howard, R. A. 1966. "Information Value Theory." *IEEE Transactions on Systems Science and Cybernetics* 2 (1): 22-26. <https://doi.org/10.1109/TSSC.1966.300074>.
- International Civil Aviation Organization. 2022. *Long-Term Aspirational Goal for International Aviation CO2 Emission Reductions*. Resolution A41-21. Montreal: ICAO.
- Kinsey, J. S., M. T. Timko, S. C. Herndon, E. C. Wood, Z. Yu, R. C. Miake-Lye, P. Lobo, et al. 2012. "Determination of the Emissions from an Aircraft Auxiliary Power Unit (APU) during the Alternative Aviation Fuel Experiment (AAFEX)." *Journal of the Air & Waste Management Association* 62 (4): 420–430. <https://doi.org/10.1080/10473289.2012.655884>.
- Levinson, D. 2005. "Micro-Foundations of Congestion and Pricing: A Game Theory Perspective." *Transportation Research Part A: Policy and Practice* 39 (7-9): 691–704. <https://doi.org/10.1016/j.tra.2005.02.021>.
- Nash, J. F. 1950. "Equilibrium Points in N-Person Games." *Proceedings of the National Academy of Sciences* 36 (1): 48-49. <https://doi.org/10.1073/pnas.36.1.48>.
- National Academies of Sciences, Engineering, and Medicine. 2019. *Optimizing the Use of Electric Preconditioned Air (PCA) and Ground Power Systems for Airports*. ACRP Research Report 207. Washington, DC: The National Academies Press. <https://doi.org/10.17226/25623>.

- Osborne, M. J., and A. Rubinstein. 1994. *A Course in Game Theory*. Cambridge, MA: MIT Press.
- Pan, D., Z. Xue, and M. S. Shahriar. "Identifying Freeway Traffic Instability from PeMS Data Using Rolling-Window Speed and Occupancy Variability." UC Berkeley: Institute of Transportation Studies at UC Berkeley, 2026. <https://escholarship.org/uc/item/2wn1n8b8>.
- Padhra, A. 2018. "Emissions from Auxiliary Power Units and Ground Power Units during Intraday Aircraft Turnarounds at European Airports." *Transportation Research Part D: Transport and Environment* 63: 433-444. <https://doi.org/10.1016/j.trd.2018.06.015>.
- Rakas, J., P. Achatz Antonelli, C. Walia, P. Rouzbahani, and G. Gikas. 2025. *Reducing Emissions through Monitoring and Predictive Modeling of Gate Operations of Idle Aircraft: A Case Study on San Francisco International Airport*. Berkeley: Institute of Transportation Studies, University of California. <https://escholarship.org/uc/item/71q5d7x3>.
- Sadati, S. M. S., and K. S. Cetin. 2020. "Data-Driven Method to Study the Impact of Utilizing Electric Ground Power Systems on Airport Electricity Demand Profile." *Transportation Research Record* 2674 (4): 261–271. <https://doi.org/10.1177/0361198120911051>.
- Saltelli, A., M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, and S. Tarantola. 2008. *Global Sensitivity Analysis: The Primer*. Chichester: John Wiley and Sons. <https://doi.org/10.1002/9780470725184>.
- Sargent, R. G. 2013. "Verification and Validation of Simulation Models." *Journal of Simulation* 7 (1): 12-24. <https://doi.org/10.1057/jos.2012.20>.
- Shahriar, M. S. (2026). Compositional Co-Design of Traffic Law Reform: Integrating Citizen-Sourced Data with Strategic Game-Theoretic Analysis. *Transportation Research Record: Journal of the Transportation Research Board*. <https://doi.org/10.1177/03611981261473514>
- Shahriar, M. S., P. Luengjitvatana, M. S. A. Baig, S. Albreiki, and A. Luther. 2026. "Distributional equity of replacing fuel taxes with vehicle miles traveled fees: Evidence from five San Francisco Bay Area cities." UC Berkeley: Institute of Transportation Studies at UC Berkeley. <https://escholarship.org/uc/item/6j93d459>
- Schmidt, M. 2017. "A Review of Aircraft Turnaround Operations and Simulations." *Progress in Aerospace Sciences* 92: 25-38. <https://doi.org/10.1016/j.paerosci.2017.05.002>.