

Evaluating NGATS Research Priorities at JPDO

Daniel R. Goldner*

Ventana Systems, Inc., Harvard, MA 01451, USA

Sherry S. Borener†

Joint Planning and Development Office, Washington, D.C., 20005, USA

The Joint Planning and Development Office (JPDO) provides annual air transportation research and development funding guidance to its participating federal agencies, to help coordinate their investments and activities toward the development of the Next Generation Air Transportation System (NGATS). Here we introduce a quantitative framework for helping JPDO evaluate the performance of the research portfolio. The Ventana NGATS Portfolio Simulator estimates the likely aggregate result of NGATS investments on system performance, using heuristics to link investments to performance estimates from more detailed simulations. We observe that even if most investments succeed, total system improvement may still be limited by any remaining constraints. Accordingly, the risk-reducing portfolio design explores a number of independent strategies for addressing each constraint. We will continue use the model to incorporate input from JPDO participating agencies and Integrated Product Teams, to help clarify program interdependencies, and to investigate how those dependencies influence the effectiveness of NGATS investments.

I. Introduction

In the VISION 100 - Century of Aviation Reauthorization Act (P.L. 108-176), Congress called for the development of the Next Generation Air Transportation System (NGATS) to meet United States' air transport needs of the year 2025. The same act created the Joint Planning and Development Office (JPDO), supported by the Federal Aviation Administration, NASA, the Departments of Commerce, Defense, Homeland Security, and Transportation, and the White House Office of Science and Technology Policy. The JPDO "serves as a focal point for coordinating the research related to air transportation for all of the participating agencies."¹ In practice, the JPDO provides annual budget guidance for participating agencies to consider as they set their research and development agendas. The goal of the budget guidance is to encourage research and development investments by each agency to increase the likely benefits, and decrease the likely cost and risk, of the ensemble effort. This paper presents an heuristic process developed at JPDO to give quantitative analytical support for research and development budget guidance.

In an investment portfolio, resources are allocated among a set of potential projects. Each candidate project has an outcome, which is unknown at the time of investment, but which can be characterized as a joint probability distribution of, for example, the cost, benefit, and completion time of the project. In the simplest case, the outcome of each project is independent of the outcome of any other project, and the expected result of the whole is just the sum of the parts. In a more complex situation, external factors correlate the outcome of one project to the outcome of other projects. Positive correlations can occur when an external factor influences both project results, for example, if both results are sensitive to fuel price. Negative correlations occur when two projects are competing for the same resources, usually money or talent. Indeed, since a project's outcome often depends on its funding level, a portfolio allocation process itself introduces anti-correlations among project outcomes.

A yet more complex situation combines externally-imposed correlations with causal dependencies among project outcomes. For example, one project might require the output of a second one, as when a new air

*2975 Pinehurst Drive, Las Vegas, NV 89109

†Director, Evaluation and Analysis Division, 1500 K Street NW, Washington, D.C. 20005

traffic management technique requires the successful invention and deployment of a data communications system. Potentially, a few key projects directly or indirectly drive the success of many others, making overall results highly sensitive to the risk profile of those key programs. In an extreme case, the total portfolio performance can be largely determined by its worst-performing elements: spectacular improvement of one component may have little overall effect if the system is still limited by another component’s performance. Kirby and Mavris² discuss these and other important considerations in technology portfolio analysis. The JPDO research portfolio enjoys all of these types of complexity, and the resulting portfolio analysis is highly non-linear.

The outcome of the portfolio is uncertain, and can be described as a probability distribution. The character of the distribution depends on the balance of correlated and uncorrelated results: investing in correlated projects provides the opportunity for larger total benefit, while anti-correlated projects tend to reduce the variance of the overall outcome. The usual recipe for a successful portfolio is to combine many diverse investments, such that the of the probability distribution of the total return is as high as possible, while the variance of the total return, or risk, is as low as possible. An ideal portfolio delivers the maximum expected return for a given risk level (or the minimum risk for a given expected return). Bhadra and Morser³ elucidate the application of these ideas to National Airspace System (NAS) investments.

The National Institute of Aerospace’s Aviation Plan⁴ took on a related analysis for near-term research and development (R&D) investments, focusing on leveraging existing work and NASA capability. Our goal is to provide a framework to consider these issues, capable of absorbing information as it is being developed at JPDO. We have created an aggregate, heuristic model to estimate how R&D leads to changes in NAS performance over time as NGATS is implemented, which we use to explore R&D portfolio design. In the next section, we describe the model and its inputs. From there we describe some simulations and results, then conclude with implications for R&D budget guidance and for further research.

II. Model

The model is the Ventana NGATS Portfolio Simulator, which was designed to take advantage of NGATS descriptions and studies that already exist or are in development at JPDO. The design of NGATS is a work in progress, and will be for some time; not all of the specifications required for this exercise are available. Our approach uses heuristic placeholder assumptions, described below. As we accumulate more information, each new detail will replace its placeholder surrogate. In the meantime, the heuristics have been chosen carefully to represent general features of portfolio performance we believe are likely to be true regardless of program specifics.

This section presents the model in detail, starting with a description of NGATS *tasks*: R&D, enabling infrastructure, and Operational Improvements (OIs). We then describe the NAS in terms of capacity constraints, and discuss how completing NGATS tasks changes those constraints. Next we cover how overall system capacity is estimated as a function of the constraints. Finally, we describe two auxiliary outputs—consumer surplus and net present value—used to value the portfolio. Figure 1 presents a schematic of the model.

NGATS tasks

The main input to the model is a list of tasks that must be performed to implement NGATS. Tasks are of three types. First, *Operational Improvements* are the changes that will be made in the NAS as NGATS is implemented. The OIs are under continuous revision as the JPDO, its member agencies, and aviation stakeholders develop their plans and ideas. The current description of the OIs is maintained in the JPDO Operational Improvements Roadmap,⁵ which contains 206 OIs as of this writing. Second, enabling systems or infrastructure such as ADS-B⁶ must be fielded. In this paper we will refer to these as *platforms*; the platforms required for each OI are listed in the Roadmap. Some platforms, but not all, require on-board equipage, which is treated as a requirement separate from the implementation of the platform’s infrastructure. Third, *Research and Development* is required for the implementation of the OIs. The list of required R&D is being developed at JPDO; for this exercise we have used a list of 152 R&D topics provided by JPDO’s Portfolio Management Division.

The model starts with an input *schedule* of the time span over which each task is to be completed. For OI’s, the schedule is provided by the Roadmap, which assigns each OI to one of eight overlapping, four-year

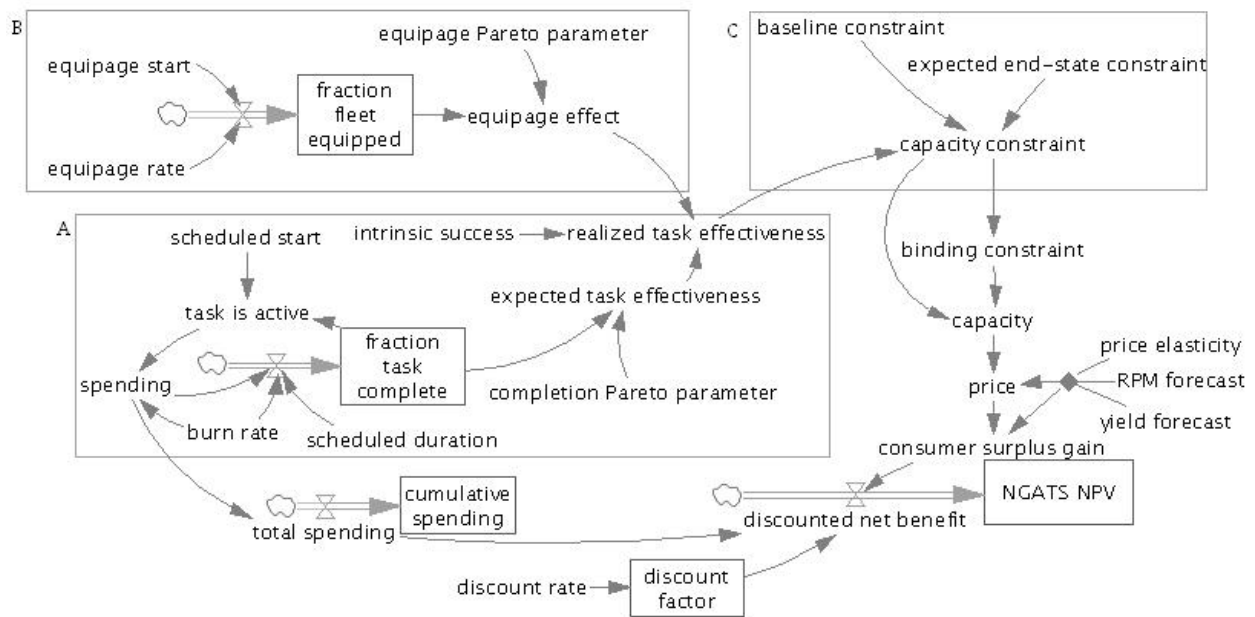


Figure 1. Model schematic. Variable names are connected by arrows indicating functional dependence, e.g., equipage effect = $f(\text{fraction fleet equipped}, \text{equipage Pareto parameter})$. Grey, single-variable boxes indicate time integrals, that is, accumulations. Variables in the large box marked “A” are arrays with one value for each R&D task, platform and OI. Variables in box B are arrays with one value for each platform requiring on board equipage. Variables in box C are arrays with one value for each system constraint. All other variables are scalar.

implementation windows called *segments*. The segments commence every two years, starting with segment 0 (2006-2009) and ending with segment 7 (2020-2023). Each of the R&D tasks have been scheduled by the Portfolio Management Division into a 2- to 12-year window depending on the expected magnitude of the effort. For each platform, we determine the timing required by the schedule of the OIs it supports, and assign implementation schedules consistent with those needs.

A second input is the estimated cost to do the R&D, field the platform or implement the OI. For this analysis we have assumed a placeholder cost for each OI, R&D task and platform of \$10 million per year.^a That implies each OI costs \$40 million over the four year implementation segment, and each R&D activity and platform installation costs between \$20 and \$120 million depending on planned duration. The spending in each year is totaled to give an annual cost of that year’s activity.

By default, we assume that money is spent on each task according to its cost and schedule, and that the percent of the task completed is equal to the percent of the budget spent. This means we are also assuming that every task starts and completes on time and on budget. The model permits relaxing this assumption, allowing true cost to differ from budgeted and also allowing schedule slip, but we have not exercised those features in this paper.

For those platforms requiring aircraft equipage, a third input specifies the fraction of the fleet that will be equipped over time. As a placeholder, we assume that equipage of the fleet begins in the same year as implementation of the platform’s ground infrastructure, and that for all such platforms, one-fifth of the unequipped fleet becomes equipped each year. This assumption results in 50% of the fleet equipped in four years, 75% in seven years and 99% in twenty years. An *equipage Pareto* specifies the percent of the fleet that must be equipped to achieve 80% of the expected benefits of the platform. Our placeholder assumption is linear: 80% of the platform’s benefit is achieved when 80% of the fleet is equipped.

We define the *effectiveness* of a task on a 0–1 scale, zero indicating the task has no effect, and 1 indicating that the task is complete and is producing 100% of its intended effect. We further distinguish between the *expected* effectiveness and the *realized* effectiveness. Expected effectiveness is that expected according to the completion status of the project. As tasks progress from zero to 100% complete, their expected effectiveness rises from zero to 1. In between, however, as with equipage, the effectiveness of the task may be higher or

^aAll monetary values in constant 2005 dollars.

lower than the percent complete, depending on the nature of the task. Some tasks may generate much of their benefit early, perhaps through being implemented at the busiest airports first. Other tasks may not produce benefits until the task is nearly complete, for example, a system that must be implemented everywhere before it can be used anywhere. To account for such differences, we use a *completion Pareto* parameter to specify the percentage completion a task must reach to achieve 80% of its effectiveness. Here, too, the placeholder assumption is linear, *i.e.*, 80% completion gives 80% of the benefit. Together, the percent complete and the completion Pareto parameter set the expected effectiveness of a task at any point in time.

Many factors can cause the *realized* effectiveness to be less than the expected effectiveness. To represent this, an *intrinsic success* parameter is assigned for each task. When set to its default value of 1, this parameter has no effect, but values less than one are used to represent cases in which an OI, R&D task, or platform achieves less than expected. A probability distribution for this parameter can be used to describe the difficulty or risk of a task, with riskier tasks given a distribution shifted toward low values. This parameter could also be used to describe funding risk, by setting its value as a function of the funding level, however we have not done so here.

A task’s realized effectiveness depends not only on its intrinsic success, but also on the realized effectiveness of all its prerequisites. Platforms and R&D are assumed not to have any prerequisites. The success of an OI, on the other hand, may depend on R&D tasks, platforms, or other OIs. An analysis was performed at JPDO to synthesize the information available in the Roadmap and the R&D list.⁷ This analysis provides our placeholder assumptions about each OI’s prerequisites. Thus, a task’s realized effectiveness can be lower than its expected effectiveness for any of three reasons: low intrinsic success, low equipage levels, or low realized effectiveness of prerequisite tasks. The model estimates realized effectiveness for each task by combining these factors in the first of three main heuristics used in this model:

Heuristic 1 *The realized effectiveness of a task is,*

- *for an R&D task, the product of its expected effectiveness and its intrinsic success;*
- *for a platform, the product of its expected effectiveness, its intrinsic success, and (when applicable) its equipage effect; and*
- *for an OI, the product of its expected effectiveness and the minimum realized effectiveness of its prerequisites.*

This “weak link” heuristic reflects the belief that early funding gaps or shortfalls in task execution will reduce the realized effectiveness of the OIs. Figure 2 shows the realized effectiveness of OIs related to aircraft noise for an example simulation.




Figure 2. Realized effectiveness for OIs related to aircraft noise from the default simulation. Numbers refer to the the OI identification number from the Roadmap; see Figure 3 for descriptions.

Aspects of NGATS performance

R&D tasks and platforms are assumed not to provide any direct system improvement, but act only as prerequisites to enable OIs. Likewise, some OIs, such as standards definitions or plan development, have no direct effect on NGATS performance, but are necessary precursors to other OIs that do. Most OIs, however, are intended to directly improve the performance of one or more aspects of the NAS. We have identified fourteen system aspects affected by the OIs, chosen for this exercise to meet two requirements:

that that the OIs could be mapped to them, and that as far as possible, other models, studies and techniques available or expected at JPDO could be used to estimate the performance improvement produced by the OIs. The system aspects selected on this basis are: (i) safety; (ii) security; (iii) capacity for certification of new procedures or equipment; (iv) operating and maintenance cost of the system; (v) aircraft noise; (vi) aircraft emissions; (vii) capacity for land side movement of people, baggage and cargo; (viii) oceanic capacity; (ix) good-weather en route capacity; (x) good-weather terminal airspace capacity; (xi) good-weather runway capacity; (xii) good-weather capacity for surface movement of aircraft and support vehicles; (xiii) robustness of capacity to bad weather; and (xiv) the skill of the system for efficient traffic flow management. Each OI in the Roadmap that directly affects the system does so by affecting one or more of these items.^b Figure 3 shows the OIs related to aircraft noise and their prerequisite tasks.




Figure 3. Dependencies among R&D tasks and OIs related to aircraft noise.

Defining capacity

Many of the performance aspects affected by the OIs relate to NAS capacity. For the purpose of this paper, we define capacity as the number of revenue passenger miles (RPMs) the system can provide, with *delays*, *noise* and *emissions* below defined thresholds, assuming flight schedules with time, space, and aircraft size distributions similar to today's. This choice is pragmatic, as it is a measure that is available from more granular simulations performed at JPDO. These simulations involve multiple models and many details, but briefly, the procedure is as follows. The FAA's Terminal Area Forecast⁸ is extrapolated to a future year, *e.g.*, 2025. The resulting traffic growth at each airport is used as an input to a schedule-generating algorithm that scales up today's schedule, but non-uniformly in order to respect differences among airports' forecast growth rates.^c The schedule is flown in a simulated NAS, with airport and sector acceptance rates set either to today's, or to assumed future values, according to scenario. Resulting delays, noise and emissions are calculated and, based on these, the fraction of demand that could be accommodated with delay, noise and emissions below the specified thresholds is estimated. That level of demand is the the overall capacity estimate. In short, we define capacity is the *delay-feasible, environmentally-tolerable* flights, under the assumptions described. The resulting numerical values depend, of course, on the choice of delay and environmental thresholds. However, any reasonable threshold can serve as a useful standard against which to measure performance, so long as the threshold is applied consistently across scenarios.

Capacity constraints

We conceive of the NAS as a system whose capacity, as defined above, is limited by multiple capacity constraints. Many of the OIs are intended to improve capacity by relieving one or more of these constraints. For this paper we focus on four capacity constraints: noise, emissions, runway capacity, and en route capacity. Consistent with our capacity definition, noise and emissions are included as constraints because community

^bGenerally the effect is positive, but in principle there could be trades where an OI would improve one aspect but degrade another.

^cOne consequence of this method is that it keeps the ratio of RPMs to flights nearly the same as today's, as the distributions of stage lengths, load factors and aircraft size are shifted only through the effects of differential growth rates among airports. This simplification can be explicitly relaxed, but for the present, changes in the ratio of RPMs to flights can be considered to be an element in the demand uncertainty discussed in Section III.

tolerance for these effects can limit the number of permitted flights in some situations. Reductions in aircraft noise and pollution can thus enable more flights.

Table 1. Assumed baseline (no-NGATS) and expected NGATS constrained capacity by individual constraint, in billions of RPMs per fiscal year. For comparison, 540 billion RPMs were flown in FY2004.⁹

Component	Baseline	Expected
Noise	900	1800
Emissions	980	1800
Runway	1070	1800
En route	1330	1800

To quantify the effect of each constraint, we ask, What would the system capacity be if this constraint were the *only* constraint on the system? For each constraint, the model requires two numerical inputs: a baseline (no-NGATS) capacity, and an expected end-state performance value, *i.e.*, the expected capacity if all the related OIs were implemented with perfect effectiveness. A Constraints Analysis¹⁰ was performed in the spring of 2006 to determine capacity, by the above definition, in different scenarios *without* NGATS. These scenarios assumed today’s system, augmented by future runways planned in the OEP,¹¹ but no other capacity improvements. From these simulations we estimated the Baseline performance values shown in Table 1. The runway capacity value is the capacity of the system with infinite en route capacity, and no noise or emissions restrictions. Similarly, the en route-constrained

capacity figure assumes infinite runway capacity, again with no noise or emissions restrictions. For noise and emissions, the study assumed community tolerance thresholds above which additional traffic would not be permitted. The resulting noise-limited permissible traffic assumes infinite tolerance for emissions, and vice versa. The Appendix contains more information on the derivation of the Baseline values.

These values provide us with an operational definition of constraint severity, which can be estimated through detailed simulation both for the no-NGATS scenario, as has been done, and for various stages of NGATS completion, for which work is currently underway. In the meantime, for this paper we assume the OIs lead to high performance, choosing 1,800 billion RPMs/year as a placeholder values for the expected end-state performance each of the four components. This is the value shown in the “Expected” column of Table 1. Thus, for the present, we are assuming that if all OIs are implemented successfully, our four constraints will be significantly improved.

But what if only some OIs are implemented? What about the transition period, when they are not all complete? What if not all are successful? The performance improvements are generally a complicated function of the status of the OIs. Some improvements may best be described by a weak-link analogy, whereby no performance gains are achieved unless all related OIs are up and running. Others aspects of performance may be incrementally improved by independent activities. Eventually the NGATS plan will be sufficiently articulated that we will be able to express these differences correctly. In the meantime, we fill in with the second major heuristic of the model:

Heuristic 2 *The improvement in a constraint is estimated by the average realized effectiveness of the contributing OIs.*

That is, for each constraint in year t , the capacity limited by that constraint alone $C(t)$ is estimated by

$$C(t) = C_b + (C_e - C_b) \frac{1}{n} \sum_{i=1}^n r_i(t) \quad (1)$$

where C_b and C_e are the constraint’s baseline and expected end-state values (as in Table 1), i is an index over the n OIs affecting the constraint, and r_i is the realized effectiveness of each such OI. Figure 4 presents example output showing the improvement in constraints over time.

System capacity

With the constraint changes calculated, we next estimate the capacity of the NAS as a whole. It would be simplest to assume system capacity equals the minimum of the capacities associated with each constraint. However, we recognize that constraints do not affect the system uniformly. For example, even if noise is,

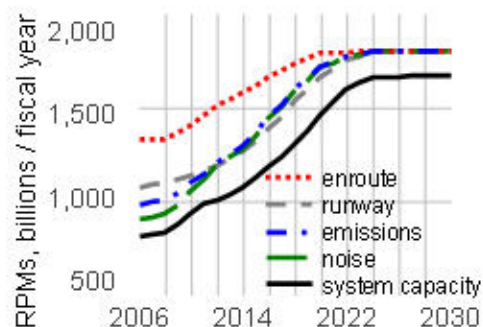


Figure 4. Constrained capacity by individual constraint, and overall system capacity, in billions of RPMs per fiscal year. As NGATS is implemented, individual constraints are lifted, raising system capacity.

overall, the most binding constraint, there are routes and schedules for which emissions or runway or en route capacities are more binding than noise. The overall capacity will therefore be somewhat less than if the most-binding constraint were the only constraint. We reflect this in the third and final heuristic:

Heuristic 3 *Total system capacity is estimated as the capacity of the binding constraint, less a further amount that depends on the severity of the less-binding constraints.*

The functional form used to apply this heuristic—including the definition of “a further amount”—is explained in the Appendix. The effect of the heuristic is illustrated in Figure 4, where the system capacity at any given time is slightly less than the most-binding constraint.

Auxiliary outputs

One measure of the value of capacity is the additional consumer surplus it generates.^{12,13} To compute the increase in consumer surplus provided by NGATS *vs.* the No-NGATS case requires an input demand curve for each year, which we specify using the FAA’s combined Mainline and Regional Commuter forecasts for RPMs and Domestic Yield.⁹ For years beyond the forecast horizon, we extrapolate the trend. The forecast provides an estimate of the quantity demanded at a specific price in each year. To estimate the quantity demanded at any other price for that year, we assume a constant-elasticity demand curve $q = q_0 (p/p_0)^e$, where q is the quantity of RPMs demanded, p is the average price, q_0 and p_0 are the forecast RPMs and forecast average price, respectively, and e is the elasticity. For this exercise we have set the elasticity to -1 . The demand curves for each future year are shown in Figure 5. Average price in each year is estimated by the intersection of the capacity for that year with the demand curve for that year. Given two capacities in year t , one with NGATS $C_N(t)$ and a baseline case $C_B(t)$ without NGATS, we determine two corresponding prices $p_N(t)$ and $p_B(t)$, and the consumer surplus gain for year t is $\int_{p_N(t)}^{p_B(t)} q(p(t))dp$.




Figure 6. Annual spending (solid), annual consumer surplus gained (dashed) and net present value of the two through each year to 2025 (dotted), for the default scenario in which all tasks are carried out successfully.

the no-NGATS system around 2015; after that, the value of the NGATS-enabled capacity grows with demand over time, reaching \$50 billion per year by 2025. Under our highly simplified assumptions, the net present value in 2006 of consumer surplus less NGATS spending through 2025 is about \$60 billion.

To summarize, we have assembled a simple model to simulate how completion of R&D and implementation of the OIs results in increased system capacity. The main inputs to the model are (i) definitions, schedules, costs and dependencies among NGATS tasks, (ii) a mapping of those tasks to the system constraints they aim to improve, and (iii) estimates of how much each constraint will improve if the OIs are successful. The main outputs are estimates of overall system capacity, and valuation measures of that capacity. The model




Figure 5. Demand curves based on FAA forecast⁹ (extrapolated) and assumed price elasticity $e = -1$.

To give a summary indication of the overall value of the portfolio, we compare the the consumer surplus to the total spending in each year. The difference between them is the value contribution for that year. Discounting at 7% per year and cumulating from 2006 to any future year gives the net present value of NGATS through that future year—or at least, the present value of those aspects of NGATS included here. The result for the reduced set of constraints considered in this paper appears in Figure 6. Early spending eventually pays off as the NGATS investments keep capacity ahead of demand. Under the placeholder assumptions of the simulation, the cumulative spending through 2025 is approximately \$20 billion. Consumer surplus benefits do not begin to accrue until the forecast demand exceeds the estimated capacity of

is intended for two uses: as a platform for evaluating the robustness and completeness of existing NGATS plans, and as a capture tool to collect and synthesize understanding of how components relate to each other and to the whole. The model’s three heuristics fill in placeholder assumptions until more specific articulation is obtained. In the meantime, they provide a first-order estimate how any particular set of investments in R&D, platforms and OIs affects overall system capacity for flights at particular delay and environmental thresholds.

The model will be modified over time to remain consistent with new information as it becomes available. Though the heuristics will be incrementally replaced, and may be modified, they are unlikely ever to disappear completely from the model. The space covered by all potential scenarios is vast, and heuristics will always be required in some form to provide approximate estimates of system performance for the many scenarios that can not be explored in detail.

III. Analysis

We aim to use the model to help assess the value of the R&D portfolio, and to explore alternative portfolio designs that increase the value. More specifically, we wish to understand how accounting for uncertainty affects considerations for portfolio composition. By uncertainty we mean the possibility that input assumptions are incorrect, for example, the unpredictability of economic growth rates (and thus air transportation demand), or the unpredictability of a program’s completion date or technical success. Before exploring uncertainty, we prepare the ground by examining the character of the deterministic calculation.

In the simulation described in Section II, there is no uncertainty: the schedule and technical success of the tasks, as well as demand and equipage levels, are assumed to come true as specified. Capacity is set primarily (though not entirely) by the most binding constraint at any given time. The capacity gained by removing the most binding constraint is limited by the next-most-binding constraint. On the other hand, if the binding constraint is not removed, improving the second-most-binding constraint has little effect. The situation is analogous to improving throughput in a production process, where the value of loosening the current bottleneck is limited by the next bottleneck, and the value of loosening anything else is zero. The potential value of R&D related to any one constraint therefore depends strongly on the outcome of R&D related to other constraints. This will have important implications for portfolio design.

How does uncertainty affect the portfolio value? Two major sources of uncertainty affect the model: project risks, and external uncertainty. *Project risks* refer to all the ways the future may differ from the plan. What if some tasks are underfunded, or unfunded? What if, for whatever reason, some tasks do not deliver the expected results? What if tasks take longer than expected, or cost more, or both? What if R&D concludes that a particular OI is not practically achievable? *External uncertainty* refers to uncertainty in the exogenous inputs to the model: demand levels and equipage rates.

Here we will explore technical risk, leaving task schedules and costs, equipage rates and demand levels as originally specified. We represent technical risk by Monte Carlo simulation in which, for each of 1000 realizations,^d the intrinsic success parameter for each task is determined by a random draw from the triangular distribution shown in Figure 7. We have assumed a moderately optimistic distribution (mode success = 0.85) in which most tasks achieve most of their objective.

Despite the relatively high rate of success of individual OIs, the resulting capacity gain (Figure 8) is only $38 \pm 5\%$ of the capacity increase seen when all tasks enjoyed perfect success. The reduced capacity gain is caused by the network of prerequisites among the R&D, platforms and OIs: if one task performs poorly, the tasks depending on it are disadvantaged, as are the tasks that depend on *those*. By itself, this effect would not necessarily be so adverse: there are many independent threads of activity in the portfolio, so poor performance in one tends to affect only a few others. However, even though each imperfection does not propagate through the whole portfolio, many of them do propagate through most of the tasks affecting a single constraint. Thus in any realization, there is usually at least one constraint on which little progress is made. Even though all the other constraints might be successfully improved, the overall system performance remains limited by the unmoved constraint.



Figure 7. Probability density function of each task’s intrinsic success.

^dOne thousand realizations can be simulated in about 10 seconds on an ordinary PC.

The value of capacity depends on the demand, and therefore, uncertainty in future demand growth drives uncertainty in the portfolio value. If the average annual demand growth rate from 2006 to 2032 is assumed to have a normally distributed uncertainty of $\pm 0.5\%$, this uncertainty *combined* with technical risk in NGATS implementation leads to a wide range of potential scenarios for demand *vs.* capacity (Figure 9, left). As a result, the NPV shows wide variation (Figure 9, right) depending on the outcomes for both capacity and demand. In all cases, however, the NPV through 2025 is substantially positive.

What is the value of the R&D portfolio? Investment in NGATS-related research is investment in a real option:¹⁴ unless the R&D is carried out, the nation will not have the option to implement the OIs. The available alternative, of course, is to postpone the initiation of NGATS research. Under the assumptions of this analysis, this is not an advisable strategy. Even the most pessimistic growth scenarios (*e.g.*, Figure 9, left) require action to avoid significant capacity shortfalls in the next two decades. Figure 9 (right) indicates the effect of postponing NGATS for 10 years: while the uncertainty in value is somewhat reduced, the expected value of NGATS (*i.e.*, the mean of the probability distribution) is significantly lower. Without beginning the R&D now, we lose the option to complete NGATS in the next two decades, when it is quite likely to be needed even under modest growth assumptions. (Option theory can also be applied to individual task investment decisions; this is discussed further in Section IV.)




Figure 8. Probability distribution of capacity, based on 1000 Monte Carlo realizations in which the intrinsic success parameter for each task is drawn from the distribution shown in Figure 7. Shown here are the mean value (thin line), 67% confidence interval (dark shaded region), 95% confidence interval (combined dark and lighter-shaded regions), capacity if intrinsic success = 1.0 for all tasks (thick solid line, same as the “system capacity” line in Figure 4), and FAA forecast demand,⁹ extrapolated (dashed line).

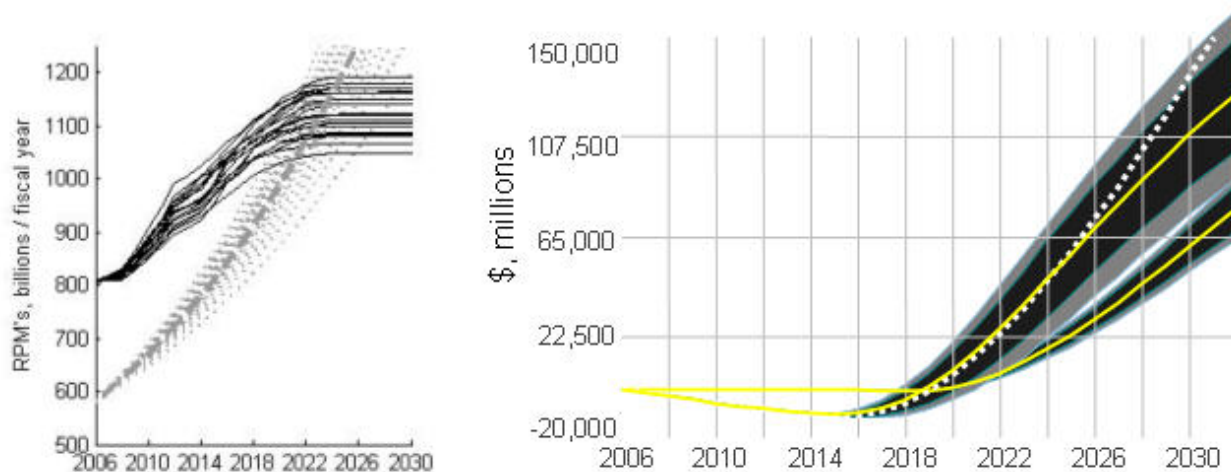


Figure 9. Left: 20 Monte Carlo realizations of demand (dotted) based on stochastic growth rate, and NGATS capacity (solid) based on stochastic intrinsic success. Forecast demand (dashed) shown for reference. Right: Probability distributions of NPV, based on 1000 Monte Carlo realizations of intrinsic success and RPM demand, showing mean value (thin line), 50% confidence interval (dark shaded region), and 67% confidence interval (combined shaded regions). Upper distribution is for NGATS as currently scheduled; dotted line shows NPV with perfect task success and forecast demand. Lower distribution is the same, but with all NGATS tasks deferred for 10 years.

Given the results above, what features of the R&D portfolio design will tend to increase expected return? How can risk be reduced? The results and discussion above suggest three important considerations:

1. *Address every constraint.* Failure to improve even one major constraint will significantly diminish the value of successful work on other constraints.
2. *Explore multiple strategies for each constraint.* If the strategies being considered pose technical, eco-

conomic or political challenges, the increased risk must be balanced by exploring more avenues toward the same goal, to increase the odds that at least one of them will succeed. If all the OIs addressing a constraint depend on the same systems or techniques, those systems and techniques must be well-supported and very likely to succeed.

3. *Don't wait.* Postponing the work reduces the expected value of the portfolio.

IV. Future Work

Mainstream work

The main objectives of upcoming work are (i) to capture the description of NGATS tasks as they are developed by JPDO agencies and Integrated Product Teams (IPTs), and (ii) to identify portfolio priorities, based on that input, that most enhance the expected value *vs.* risk characteristics of the portfolio.

This work will help to refine the developing NGATS specification by requiring unique and non-overlapping descriptions of tasks, and by making dependency assumptions explicit so they may be reviewed and revised. Over time, ongoing engagement with the particular agencies and IPTs responsible for a constraint or subset of OIs will allow us to replace the heuristics and placeholders used here with more specific information about schedule, cost, Pareto parameters, and interdependency. We also will be able to collect and apply expertise on reasonable ranges to assume for uncertain parameters, and provide greater articulation of sources of risk, such as distinguishing funding risk from technical risk.

At the same time, we will be combining this input with performance assessment results from more detailed simulation models. Performance assessments for constraints not discussed in this paper, such as surface movement of aircraft, will be treated. Similarly, system performance attributes other than capacity, such as operating cost and a quantification of safety, will be made explicit in the model.

By bringing the NGATS dependency description together with estimates of resulting system performance, we will be able to help recommend investment priorities for the best balance of benefit to risk in the overall portfolio.

Optimization and adaptive behavior

In addition to the mainstream work, this effort affords opportunities for experimenting with numerical optimization techniques, and for considering the portfolio value in light of adaptive behavior by the agencies and by NAS users.

Most portfolio optimization algorithms are designed for portfolios in which different components, while possibly correlated, do not explicitly depend on one another's result to determine their own contribution. In the NGATS case, where overall performance is determined largely by the constraint on which the least progress is made, selecting the optimal portfolio for a constrained budget is a highly non-linear program. The search space is also large, and no feasible closed form or even complete-search solution exists. The performance of incomplete-search algorithms such as simulated annealing or a genetic algorithm may depend strongly on the quality of the initial selection. Given the observations in this paper, successful heuristics for choosing an initial subset of proposed R&D activities will likely favor more binding constraints over less (since solving the latter is only valuable if the former are also solved), and topics with potential benefits for multiple constraints over those targeted at a single constraint.

To do a proper job of valuing the portfolio, the option value of R&D programs, enabling platforms and OIs should be considered individually. In the Monte Carlo simulations of Section III, those tasks that drew low success numbers were nevertheless carried out to completion. Sometimes an ill-fated project is not recognized as such until near its expected end date. However, preliminary research or investment analysis often makes it clear that the project is unlikely to succeed, and the agency can adapt by canceling the effort. This could be reflected with a stopping criterion in the model, a rule that causes a task to shut down if it fails to meet some threshold. A Monte Carlo simulation with stopping rules that depend on project status would yield the expected NPV of the portfolio *including* the effects of decision-making along the way—which is the option value of the portfolio. However, defining an appropriate stopping rule is an interesting problem in itself. Should some high-risk programs (*i.e.*, those unlikely to succeed) be protected if their potential value is large? What stopping rule gives the best option value for the portfolio? Should the stopping rule be prescriptive, or descriptive of observed government practice?

Another step in the same direction would be to model research outcomes more explicitly. In such a scheme, the cost, implementation duration, and technical difficulty of an OI would be represented by a probability distribution around a “true” value. The effect of R&D would be to narrow the uncertainty so that the likely outcome of the OI would be more clear; the resulting value would influence project initiation or stopping criteria.

A different extension would represent the adaptive behavior of system users. Currently, demand and equipage rates are stochastic inputs to the model. However, increasing system capacity reduces delays, lowering operator costs and increasing the attractiveness of air travel. These effects tend to increase both supply and demand for air travel services. This has a secondary effect of accelerating fleet growth, increasing the fraction of new aircraft in the fleet, potentially increasing equipage rates and accelerating capacity gains. Models such as the NAS Strategy Simulator¹⁵ estimate system demand endogenously via those supply and demand effects, taking capacity as an input, while the NGATS Portfolio Simulator estimates system capacity endogenously, and takes demand as an input. Connecting the two models could explore NGATS capacity in the context of *reactive* system demand, and the implications for consumer surplus and for system costs and revenues.

NGATS is a large mosaic, of which these simulations have placed only a few of the pieces. We look forward to continued progress toward understanding the dependencies and feedback among its many aspects.

Appendix

Considering the NAS as a collection of constraints raises an issue, because the constraints are not uniformly binding. In this study, when we make a statement such as “noise is more binding than emissions”, we are saying a system constrained *only* by noise tolerance would permit fewer flights than a system constrained *only* by emissions tolerance. But when both constraints are present, even if noise is the dominant limiting factor, there will still be some routes at some times for which emissions is the constraint. Thus, even if noise alone is more binding than emissions alone, noise and emissions together are even more binding in total.

A similar effect will be true for any combination of constraints, and this is observed in the detailed simulation results used as inputs to this study. To maintain consistency with those inputs, we need a way to approximate this additional penalty, that is, the capacity reduction beyond that imposed by the tightest constraint which is due to the times and places where the other constraints are limiting.

Table 2. Estimated capacities for combinations of en route (E), runway (R), emissions (X) and noise (N) constraints, developed from Constraints Analysis¹⁰ results (CA) or estimated from those results using Equation 2 (*Est.*) as described in the text.

Constraints Applied	Capacity <i>B RPM/FY</i>	Source
E	1330	CA
R	1070	CA
ER	1050	CA
X	980	<i>Est.</i>
ERX	940	CA
N	900	<i>Est.</i>
ERN	870	CA
ERXN	840	CA
ERXN	830	<i>Est.</i>

The Constraints Analysis¹⁰ from which we derived no-NGATS constraint levels simulated several combinations of constraints: en route capacity only; runway capacity only; en route and runway; en route, runway and emissions; en route, runway, and noise; and finally all four. Table 2 presents the results. The effects of heterogeneity described in the previous paragraph can be seen, for example, in the en route- and runway-constrained capacity (ER) being slightly less than the runway-alone case (R), and in the fully-constrained capacity (ERXN) being slightly less than the en route-, runway- and noise-constrained capacity (ERN).

In NGATS Portfolio simulations, the combinations and severity of constraints changes in each time step of each realization, so we need to estimate heuristically the effects on overall capacity of different constraint combinations. In a suitable heuristic, the impact of a secondary constraint on system capacity should increase with the severity of the constraint, reaching a maximum as the secondary constraint approaches the severity of the most-binding constraint. Conversely, the impact of the sec-

ondary constraint should decrease to zero as its severity decreases. Finally, adding more constraints should decrease the system capacity. We instantiate these requirements in the following heuristic for system capacity

$S(t)$:

$$S(t) = C_{\min} - k \left[1 - \exp - \left(\sum_{i=1}^n \exp \frac{-(C_i - C_{\min})}{k} - 1 \right) \right]. \quad (2)$$

Here C_i is the constrained capacity under a single constraint i as defined in Equation 1, and $C_{\min} = \min C_i$. This is not derived from first principles; rather, it is devised to possess the desired properties with a minimum of free parameters. The constant k is the only such parameter, setting the scale for the impact of the non-binding constraints on system capacity. In the limit as k approaches zero, the system capacity approaches the capacity set by the most binding constraint.

The Constraints Analysis did not provide estimates for emissions- and noise-constrained capacities directly, only in combination with en route and runway capacities. These two capacities and the parameter k make three unknowns, whose values were determined by substituting them into three applications of Equation 2, along with the ER, ERX, and ERN data points from Table 2. The solution of this system gave estimates of $k = 125$ billion RPMs/year, emissions-constrained capacity $X = 980$ billion RPMs/year, and noise-constrained capacity $N = 900$ billion RPMs/year. As a check, these values were substituted into (2) along with the Constraints-Analysis values for E and R to estimate the system capacity for all four constraints. The resulting value was 830 billion RPMs/year, adequately close to the value of 840 derived from the Constraints Analysis.

In short, we have developed an expression for the combined effect of multiple constraints which meets common sense objectives and which, when calibrated to three data points provided by detailed simulations, comes reasonably close to predicting a fourth. Whether this same expression will adequately predict a fifth, sixth or seventh data point will not be known until additional simulation results become available. In the meantime, this heuristic serves to smoothly and reasonably interpolate among the few cases that have been analyzed in detail.

Acknowledgments

This work would not have been possible without J. Fearnside's thoughtful analysis of OI dependencies. Thanks also to C. Burg, J.-P. Clarke, N. Fujisaki, T. Fiddaman, and the Portfolio Management Division and Evaluation and Analysis Division of JPDO for comments and suggestions.

References

- ¹Joint Planning and Development Office, "Fact Sheet," http://www.jpdo.aero/site_content/factsheet.html, 2006.
- ²Kirby, M. and Mavris, D., "An approach for the intelligent assessment of future technology portfolios," AIAA 2002-0515.
- ³Bhadra, D. and Morser, F. R., "Analysis of System-wide Investment in the National Airspace: A Portfolio Analytical Framework and an Example," Presented at the 5th Annual ATIO/AIAA Conference, Crystal City, VA.
- ⁴National Institute of Aerospace, "Responding to the Call: Aviation Plan for American Leadership," <http://www.nianet.org/nianews/AviationPlan.php>, 2005.
- ⁵Joint Planning and Development Office, "Operational Improvements Roadmap," <http://techchanger.jpdo.aero>, 2006.
- ⁶Federal Aviation Administration, "ACB-420 ADS-B Home Page," <http://adsb.tc.faa.gov/ADS-B.htm>, 2006.
- ⁷Fearnside, J. and Goldner, D., "Dependencies among NGATS OIs, R&D, and Enabling Infrastructure," Unpublished JPDO working paper.
- ⁸Federal Aviation Administration, "Terminal Area Forecast," <http://www.apo.data.faa.gov/main/taf.asp>.
- ⁹FAA Office of Policy and Plans, *FAA Aerospace Forecasts, Fiscal Years 2006-2017*, Federal Aviation Administration, 2006.
- ¹⁰Borener, S., James, K., and Orme, J., "Initial Constraints Analysis: Capacity and Environmental Impacts on NAS Performance," Presented at JPDO, Mar 2006.
- ¹¹Federal Aviation Administration, "Operational Evolution Plan," <http://www.faa.gov/programs/oep/>.
- ¹²Office of Management and Budget, "Circular No. A-94 (Revised)," <http://www.whitehouse.gov/omb/circulars/a094/a094.html>, Oct 1992.
- ¹³Cavolowski, J. and Olson, L., "FAA-NASA Socio-Economic Demand Forecast Study," Jan 2004.
- ¹⁴Dixit, A. K. and Pindyck, R. S., *Investment Under Uncertainty*, Princeton University Press, Princeton, New Jersey, 1994.
- ¹⁵Ventana Systems, Inc., "NAS Strategy Simulator," Ventana Systems, Inc., 60 Jacob Gates Road, Harvard MA 01451, www.ventanasystems.com.