



The Mitchell Forum

From Platform Costs to Campaign Effects: Embedding Cost-Per-Effect in Force Design

by Maj Erik Schuh, USAF

About the Forum

The Forum presents innovative concepts and thought-provoking insight from air and space experts here in the United States and across the globe. In order to afford publishing opportunities for thoughtful perspectives, Mitchell Institute's Forum provides high visibility to writing efforts on issues spanning technology and operational concepts, defense policy and strategy, and unique interpretations of changing geopolitical relationships.

The opinions, conclusions, and recommendations expressed or implied in this article are those of the author and do not necessarily reflect the official policy or position of the U.S. Government, Department of Defense/War, the U.S. Air Force, the U.S. Space Force, or the Mitchell Institute for Aerospace Studies.

Abstract

Every year, the Department of Defense/War (DOD/W) allocates resources across competing investments without a consistent way to compare their value in a campaign. As a result, defense debates continue to rely heavily on platform-level acquisition metrics that provide limited insight into how forces actually perform in combat. Although a cost-per-effect methodology was intended to address this problem, the DOD/W still lacks an analytic framework capable of applying it systematically to force structure decisions at scale.

Improving force design decisions requires shifting cost comparisons from individual platforms to campaign-level outcomes. The Pentagon needs a standardized comparative cost framework and a tiered modeling architecture that links exploratory trade-space analysis with campaign-level validation to achieve this shift. Institutionalizing a cost-per-effect framework within the defense planning process would allow senior leaders to evaluate investment trade-offs more consistently, shape the campaign scenarios used to guide force planning, and better align resourcing decisions with the national defense strategy.

Figure 1: F-117 Nighthawk aircraft provide an instructive example of cost-per-effect assessment. In terms of flyaway unit costs alone, older model F/A-18s might appear cheaper than stealthy F-117s. But in Desert Storm, a four-ship of F/A-18s required a total strike package of 41 aircraft to pursue one target with three aimpoints. A strike package of 20 F-117s, meanwhile, was able to pursue 28 targets with 38 total aimpoints. The differential in the effect makes clear that unit cost alone is not a sufficient metric for building a budget to meet a strategy.

Credit: U.S. Air Force Photo. See also: *Resolving America's Defense Strategy-Resource Mismatch: The Case for Cost-Per-Effect Analysis*, p. 12.



Why Platform-Level Cost Debates Fail Force Design

When analysts talk about the cost of a capability, it turns out they are not all speaking the same language. The most commonly cited costs—especially those that appear in headlines—tend to be those like average procurement unit cost or the program acquisition unit cost. These metrics may be useful for budgeting and oversight, but they do not capture the true cost to generate military effects in combat.

Consider the U.S. Navy's use of SM-2 and SM-6 missiles to shoot down Houthi drones in the Red Sea.¹ With unit costs of roughly \$2.2 million and \$4.3 million, respectively, these interceptors are widely portrayed as disproportionately expensive responses to relatively cheap threats.² But the missile itself represents only one component of the cost required to produce the defensive effect. The destroyer that fires the interceptor, the crew that operates the system, the fuel and maintenance required to keep it at sea, and its logistics network all contribute to the outcome. While no one should allocate all these costs to a single engagement, ignoring them entirely risks dramatically underestimating what it truly costs to achieve even a simple tactical effect like intercepting a drone.

Part of the difficulty of costing the delivery of effects stems from the absence of a standardized approach across the DOD/W for comparing the operational costs of different capabilities.³ Each service applies its own methodology, and, even then, those cost estimates vary widely. Multiple offices can be considered authoritative sources for cost estimates on the same capability. While independent cost agencies, such as the Office of Cost Assessment and Program Evaluation (OSD CAPE) and the Air Force Cost Analysis Agency, help address this problem by producing Non-Advocate Cost Assessments, these estimates are typically optimized for

budgetary decisions rather than campaign-level effects comparisons.

The deeper issue is not simply the variation of individual cost estimates, nor how they are captured. It is that most cost discussions occur at the platform level, while force design decisions are made at the campaign level. Comparing the price of one missile to the price of another may be useful for understanding procurement trade-offs in the budget, but it reveals little about which combinations of capabilities are the most effective in achieving campaign objectives. As a result, platform-level comparisons often create the appearance of analytic clarity, but they may actually obscure the campaign interactions that ultimately determine military success.

For this reason, improving force design decisions requires shifting the discussion away from the cost of individual systems and toward the cost of generating campaign-level effects. Doing so begins with a clearer understanding of what “cost” and “effect” mean in the context of future conflict.

Defining Cost & Effect at the Campaign Level

Shifting cost discussions from platforms to campaigns requires first clarifying what analysts mean by both cost and effect. These terms appear straightforward, but in practice they are often defined too narrowly to support force design decisions.

In most defense debates, cost is treated primarily as an acquisition metric. While procurement cost metrics provide insight into how expensive a system is to buy, they reveal little or nothing about what it costs to generate the intended operational effect. A capability that appears inexpensive in isolation may require tankers, forward basing, and fighter escorts to generate the intended effect. Understanding cost at the campaign level requires accounting for these supporting

requirements and dependencies to get the real cost of employing that capability.

The 1991 Gulf War offers a clear illustration. On a per-unit basis, legacy fighters looked far cheaper than the F-117 stealth aircraft. However, it took a strike package of 41 non-stealth aircraft to hit eight targets, whereas 20 F-117 stealth fighters were able to strike 28 targets in the same amount of time. If analysts only considered aircraft numbers, the non-stealth systems would seem more economical. But when comparing these systems as part of a force package, those non-stealth aircraft required escorts, jammers, and other suppression platforms to survive the mission. After adding up all these additional costs, the more expensive F-117 required a much smaller logistics footprint and proved to be more cost-effective.⁴

In the example above, the desired outcome was to neutralize enemy threats. But effects can range from destroying individual tactical units to the cumulative results of a multi-domain campaign. Because broader effects are harder to measure, discussions often collapse down to one-on-one comparisons between systems. While these comparisons make performance easier to quantify, they say little about what determines success in combat. A missile that is 10 percent more lethal in a tactical fight might yield no meaningful improvement operationally if the real constraint is an aircraft's sensors rather than the missile itself. Defining effect is clearly a unique challenge to conducting a cost-per-effect assessment.

This means that cost-per-effect analysis must be connected to campaign-level outcomes, where the full set of operational interactions and constraints determine the full cost or impact of delivering effects with a specific capability. Even then, cost-per-effect assessment cannot account for all variables. Its purpose is to identify which force structures

are most effective in a future conflict, not to dictate how those forces are built, sustained, or employed over time. Moreover, as effects become more ambiguous—such as deterring or coercing adversaries' military actions and responses—the ability to measure them degrades rapidly. In these cases, quantitative analysis must be complemented by human judgment and qualitative assessment of how force structure operates as a larger system. At the same time, major defense acquisitions remain overwhelmingly focused on individual warfighting capabilities, so cost-per-effect remains a powerful and relevant analytic tool.⁵

What the “Cost” in Cost-Per-Effect Should Measure

Unit costs are frequently invoked in debates over defense budgeting decisions, but not every dollar spent by the DOD/W is relevant to evaluating campaign performance. The Department must move from looking at acquisition costs to comparative costs in order to understand the true resource drain of an operation. Using cost-per-effect as a meaningful analytic tool, therefore, requires a standardized categorization of dollars that should not change the outcome of a decision. These categories include:⁶

- **Sunk Costs:** Costs that have already been incurred and cannot be recovered. One persistent cost assessment shortfall in the Pentagon is the inclusion of sunk costs irrespective of whether a program continues to be valuable. Missiles already in the inventory have already passed their research and development phase, so those sunk costs shouldn't be included in the cost comparison. Including them risks making already-fielded capabilities appear artificially expensive relative to alternatives whose development costs have not yet been incurred.

- **Common Costs:** Costs that are identical in both timing and magnitude for all alternatives. Because common costs affect multiple capabilities equally, they do not meaningfully influence comparisons between them. For example, when multiple aircraft rely on the same infrastructure—such as staging from Andersen Air Force Base—those shared costs should not drive force-structure comparisons.
- **Indirect Costs:** Costs that cannot be attributed directly to the system. Direct costs—such as sensors, refueling aircraft, and logistics support required to employ a capability—must be incorporated because they determine whether that capability can generate effects in contested environments. Indirect costs, however, are often too difficult to attribute consistently across alternatives. For example, determining how much of a training base’s support structure should be assigned to a specific aircraft in a campaign scenario quickly becomes analytically unstable and rarely changes the outcome of the comparison.
- **Negligible Costs:** Costs that are so small they will not affect the results. Negligible costs typically represent less than 1 percent of total unit cost. This helps analysts determine how far back in the pipeline of direct costs to consider. Eventually the costs become so insignificant that the effort to obtain those numbers doesn’t match the change in value to make them worth collecting. For example, there might be a minute, if not immeasurable, difference in cost needed to train or keep operators certified in the capability.

Distinguishing among these categories helps ensure that cost-per-effect analysis captures the costs that matter most for campaign performance without becoming

analytically unmanageable. Sunk, common, indirect, and negligible costs like historical expenditures or support costs that do not materially affect campaign outcomes should not influence future force design decisions. This shift transforms program acquisition unit cost (PAUC) from a budgeting metric into a comparative decision tool aligned with campaign-level analysis.

Defined this way, capability costs reflect the additional resources required to generate an intended effect rather than the total historical expenditure associated with a program. Focusing on remaining quantities allows analysts to compare alternative force structures based on what additional investment is required—not what has already been spent.

Aggregating these real costs of capabilities required to deliver the intended effect now creates a structured method for comparing alternative force structures against campaign objectives. Moreover, it provides a way to assess how investments in survivability, magazine depth, and logistics support translate into operational advantage over time.

Despite its usefulness, however, cost-per-effect has not become a standard framework for evaluating force-structure alternatives across the DOD/W. The challenge is not conceptual disagreement about its value, but the absence of an analytic architecture to compare campaign-level effects across alternative force structures.

The Missing Analytic Layer for Comparing Campaign-Level Effects

Every few years, the DOD/W develops defense planning scenarios to anticipate future conflict and guide force development.⁷ These scenarios help the Department evaluate which force structures are most likely to succeed. The Office of the Secretary of War and the Joint Staff—working with services’ futures organizations and analytic communities—must be able

to explore an enormous range of force combinations quickly and systematically to effectively accomplish this task.

Looking at the Gulf War example from earlier, integrating all the capabilities and associated costs with one strike mission takes careful consideration. A single operation can have hundreds of missions, and a campaign can have thousands, all going on simultaneously across air, land, sea, space, and cyberspace. So while options for the single attack mission are relatively easy to compare, the number of interdependencies at the campaign level explodes to the point where no human can compare all the interactions properly. These questions quickly outstrip human capabilities and demand sophisticated modeling.

The problem is that today's analytic tools sit at two extremes. High-fidelity campaign simulations capable of capturing complex interactions across domains comprise one end of the spectrum, but they are limited in their ability to evaluate more than a small number of force-structure alternatives at the same time. At the other end, exploratory trade space tools can evaluate large numbers of alternatives quickly, but they cannot capture the operational dependencies that determine campaign success. Requirements and programming decisions frequently occur between these two analytic regimes, and there is no consistent mechanism linking exploratory analysis to campaign-level validation.

This mismatch between analytic speed and analytic realism limits the Department's ability to compare competing investments at scale. Force design decisions are rarely choices between two clearly defined alternatives. Instead, they involve evaluating large combinations of capabilities across multiple timelines and theaters under uncertainty about future threats. Without a framework capable of linking cost and campaign outcomes, cost-per-effect remains difficult to apply as a

practical tool for evaluating alternatives.

The result is that cost analysis, operational analysis, and programming decisions often proceed on parallel tracks rather than within a shared analytic structure. Cost estimates may be available, and campaign simulations may exist, but the connective layer required to compare alternative force structures systematically across future operating environments remains incomplete. Addressing this gap requires a modeling approach that balances speed and realism while supporting comparisons across multiple campaign timelines.

A Tiered Modeling Framework for Campaign-Level Cost-Per-Effect

A tiered modeling framework provides a practical way to connect exploratory analysis with campaign-level validation and preserve the analytic speed required for force structure decision-making. Rather than relying on a single modeling approach to answer all questions, a tiered system allows different analytic tools to contribute at different stages of the decision process.

Tier 1: Rapid Exploration

Rapid exploration is designed to quickly search the force design trade space. Planning organizations—either through their internal analytic shops or in partnership with organizations like OSD CAPE or the Joint Staff's Force Structure Resources, and Assessment Directorate (J8)—would use fast exploratory models to examine millions of force combinations. The U.S. Air Force's Combat Forces Assessment Model (CFAM) is one such model. Instead of simulating every engagement, the model uses optimization techniques to explore millions of possible force structures, allowing analysts to identify promising force mixes. The output of rapid exploration is not a decision but a ranked set of feasible force structures.

Tier 2: Validation

Once a few force structures have been identified, those analytic organizations like OSD CAPE are needed to develop the most realistic modeling scenario the U.S. military has to validate these force structures. The gold standard for campaign analysis is a simulation called Synthetic Theater Operations Research Model (STORM). STORM is a high-fidelity campaign simulation—essentially a massive computer model that can simulate a conflict—with detailed force interactions across domains. But STORM scenarios are enormously time-consuming to build and run. That constraint means this model must be used sparingly to validate the most promising options identified by faster exploratory models.

Tier 3: Force Structure Innovation

Models used in rapid exploration and validation cannot search the entire decision space of modern war. A force structure innovation tier focuses on expanding that frontier. The Defense Advanced Research Projects Agency's Strategic Chaos Engine for Planning, Tactics, Experimentation and Resiliency (SCEPTER) is an early example. The idea is to have artificial intelligence assistants as battle managers that can help human commanders evaluate courses of action in a dynamic, complex battlespace—something far beyond the ability of today's models.⁸ These models should be pushed to all planning organizations across the military to explore the trade spaces across all levels of war.

Together, these three tiers form a coherent, scalable cost-per-effect ecosystem. Rather than attempting to replace existing models, this framework connects them into a decision-support structure that allows exploratory analysis and campaign validation to inform each other. In doing so, it enables decision-makers to evaluate how alternative force structures translate resources into operational outcomes across future conflicts.

Institutional Implications for Operationalizing Cost-Per-Effect

Applying cost-per-effect consistently across force design decisions requires closer integration between exploratory modeling, campaign-level simulation, and programming analysis than currently exists across the DOD/W. While each of these analytic activities already occurs, they rarely operate within a shared framework capable of supporting campaign-level comparisons across large sets of alternatives.

Improving this integration does not require replacing existing organizations or acquiring new models. Instead, it requires integrating a structured modeling architecture into the defense planning process that links exploratory analysis, campaign simulation, and programming decisions within a single analytic framework. Establishing this shared analytic structure would enable cost and operational analysis to reinforce one another rather than proceed on parallel tracks.

Such an approach would also support comparisons across future operating timelines consistent with strategic guidance. Force design decisions necessarily involve trade-offs between near-term readiness and long-term modernization. A modeling architecture capable of evaluating campaign-level cost-per-effect across multiple time horizons would allow decision-makers to assess these trade-offs more transparently and consistently than is currently possible.

Ultimately, cost-per-effect is not intended to replace expert judgment in force design decisions. Rather, it provides a structured method for comparing how alternative investments translate resources into operational advantage. Building the analytic architecture required to support these comparisons is therefore a prerequisite for using cost-per-effect as a practical tool for force design.

Institutionalizing Cost-Per-Effect in Defense Planning

Identifying how investments can become advantages in the battlespace requires more than improved analytic techniques—it requires institutionalizing cost-per-effect as a standard component of the defense planning process. Today, cost analysis, campaign modeling, and programming decisions are often conducted within separate communities using different assumptions and timelines. Integrating cost-per-effect across these activities would provide senior leaders with a consistent basis for evaluating how potential investments could support campaign outcomes. The DOD/W can take three practical steps to establish this framework across the force design enterprise.

1. **Standardize a Comparative Cost Framework**

Across the DOD/W: The Secretary of Defense/War should direct OSD CAPE, in coordination with the Joint Staff J-8 and the service cost agencies, to establish a standardized comparative cost methodology and supporting database. A centrally maintained cost baseline would allow analytic organizations across the DOD/W to evaluate alternative force structures using consistent assumptions about the supporting requirements needed to generate operational effects. It is time for services and program offices to move away from their own cost numbers and toward a shared comparative cost framework that aligns with campaign-level decision-making.

2. **Establish a Tiered Cost-Per-Effect Modeling**

Architecture: The Secretary of Defense/War should direct OSD CAPE, in coordination with the Joint Staff J-8, to establish a standardized three-tier modeling architecture that links rapid exploratory trade space models with high-fidelity campaign simulations and innovative AI models. Today, these

analytic activities largely occur in isolation, leaving force design decisions up to modeling based on speed or realism but rarely both. A tiered modeling framework would allow analysts to identify promising force structures at scale and validate them under campaign conditions earlier in the force design process. Connecting these modeling layers would provide senior leaders with a consistent process for comparing how alternative investments translate into campaign-level outcomes before committing resources.

3. **Require Campaign-Level Cost-Per-Effect Analysis Before Every Program Objective Memorandum (POM) Cycle:**

The DOD/W should require campaign-level cost-per-effect analysis as part of the analytic basis supporting each POM cycle. Tying resourcing decisions to campaign outcomes shifts the process away from platform-level comparisons and toward evaluating how capabilities contribute to campaign performance. Integrating campaign-level cost-per-effect into the POM process would provide senior leaders with a consistent basis for evaluating trad-offs, strengthen the Department's ability to explain to Congress what additional resources are buying in operational terms, and align force design decisions more directly with the National Defense Strategy and defense planning scenarios.

Designing Forces for Campaign Outcomes, Not Platform Comparisons

Force design decisions shape the services' ability to achieve strategic objectives in joint force operations across future conflicts, but many cost debates continue to focus on the price of individual platforms rather than the campaign outcomes those platforms enable. As a result,

discussions that appear analytically precise at the platform level often obscure the operational interactions that determine success in war.

Cost-per-effect provides a way to compare alternative force structures based on their ability to generate campaign-level outcomes rather than individual engagements. When supported by a tiered modeling architecture that links exploratory trade space analysis with campaign-level validation, cost-per-effect becomes a practical method for evaluating how resources translate into operational advantage across multiple future operating environments.

Strategic guidance defines what joint force operations must be capable of achieving, but campaign-level analysis provides the most rigorous method for evaluating how alternative force structures support those objectives over time. Building the analytic architecture required to support campaign-level cost-per-effect comparisons is therefore essential to designing forces capable of sustaining operational advantage in future conflicts. ✪

Endnotes

- 1 Lauren C. Williams, “[Mideast missile duels have cost US Navy nearly \\$1B, secretary says](#),” *Defense One*, April 16, 2024.
- 2 Wes Rumbaugh, “[Cost and Value in Air and Missile Defense Intercepts](#),” *Commentary*, Center for Strategic and International Studies, February 13, 2024.
- 3 Katharina Ley Best et al., *[Beyond Cost-per-Unit](#)* (Santa Monica, CA: RAND Corporation, June 2024).
- 4 David A. Deptula and Douglas A. Birkey, *[Resolving America’s Defense Strategy-Resource Mismatch: The Case for Cost-Per-Effect Analysis](#)* (Arlington, VA: The Mitchell Institute for Aerospace Studies, 2020).
- 5 Secretary of War, “[Transforming the Defense Acquisition System into the Warfighting Acquisition System to Accelerate Fielding of Urgently Needed Capabilities to Our Warriors](#),” Memorandum for senior Pentagon leadership, Commanders of the Combatant Commands, and Defense Agency and DOW Field Activity Directors, November 7, 2025.
- 6 Background Paper on Comparative Cost Analysis, Secretary of the Air Force, Studies & Analysis internal document, 2025.
- 7 Michael J. Mazarr et al., *[The U.S. Department of Defense’s Planning Process](#)* (Santa Monica, CA: RAND Corporation, February 4, 2019).
- 8 Sydney J. Freedberg Jr., “[3 ways DARPA aims to tame ‘strategic chaos’ with AI](#),” *Breaking Defense*, September 5, 2023.

About The Mitchell Institute

The Mitchell Institute educates the general public about aerospace power's contribution to America's global interests, informs policy and budget deliberations, and cultivates the next generation of thought leaders to exploit the advantages of operating in air, space, and cyberspace.

Forum Submissions and Downloads

For more information about submitting papers or ideas to the Forum, or for media inquiries, email our publications team at forum.mitchellaerospacepower@afa.org

Copies of Forum papers can be downloaded under the publications tab on the Mitchell Institute website at <https://www.mitchellaerospacepower.org>

About the Author

Maj Erik Schuh, USAF, is an Air Force operations research analyst currently serving in the Department of War's Chief Digital and Artificial Intelligence Office. He has held operational and analytical assignments across Europe, the Middle East, and the Pentagon. He holds graduate degrees in operations research, international relations, and international security studies and is pursuing his doctorate at the Johns Hopkins School of Advanced International Studies.

