Operations Research
Vol. 60, No. 5, September–October 2012, pp. 1035–1049
ISSN 0030-364X (print) | ISSN 1526-5463 (online)

http://dx.doi.org/10.1287/opre.1120.1085

Optimizing Boat Resources at the U.S. Coast Guard: Deterministic and Stochastic Models

Michael R. Wagner
Michael G. Foster School of Business, University of Washington, Seattle, Washington 98195, mrwagner@uw.edu

Zinovy Radovilsky
California State University, East Bay, Hayward, California 94542, zinovy.radovilsky@csueastbay.edu

The United States Coast Guard (USCG), a part of the U.S. Department of Homeland Security, is the nation's leading agency in maritime security, safety, and stewardship. One of the primary USCG resources is a fleet of boats (maritime vessels less than 65 feet in length) of various types that must be allocated to USCG stations nationwide. This paper describes the academic-industry collaboration between the authors and the USCG, which resulted in the development of an integer linear programming model that optimally matches supplies of various types of boats to station demands. The paper also introduces a model for the optimal sharing of scarce boat resources. In addition, we generalize our model, using value-at-risk and robust optimization ideas, to manage the risk of boat shortages. The paper reports on the USCG implementation process and discusses internal resistance issues and eventual adoption. We describe USCG modifications to the model recommendations due to practicalities not captured by our model. Finally, we present the significant improvements to USCG quantitative performance metrics that resulted from our model's recommendations. These include a considerable reduction of excess capacity and boat shortages at the stations, a decrease in the overall fleet size with a simultaneous increase in boat utilization, and overall reduction of the fleet operating cost. We also discuss in depth how our model effected these improvements.

Subject classifications: resource allocation; applications; military; risk management.
Area of review: OR Practice.
History: Received June 2010; revisions received July 2010, October 2010, May 2011, December 2011, March 2012; accepted April 2012.

1. Introduction

The United States Coast Guard (USCG), a part of the U.S. Department of Homeland Security, is the nation's leading agency in maritime security. According to the Department of Homeland Security (2010), the USCG are the first responders to any emergency in U.S. ports and waterways; hence its mission is critical to public safety and international commerce. According to the Commandant of the Coast Guard, Admiral Thad Allen (2009a) (the highest ranking officer of the USCG), "Over the past several years, the Coast Guard has faced increasing demands for our services, a deteriorating fleet of operational resources and the need to streamline, simplify and integrate our command and control and mission support structures." Allen (2009b) shows that the USCG is responsible for the safety and security of more than 300 ports, 3,700 marine terminals, 25,000 miles of coastal waterway, and 95,000 miles of combined coastline belonging to the United States. The agency responds to some 50,000 distress calls a year, saving many lives.

The USCG missions are carried out by three main forces: (1) cutters—vessels with a length of more than 65 feet, (2) aircraft (airplanes and helicopters), and (3) boats—vessels under 65 feet in length. In this paper, we concentrate on the USCG boats and their allocation among the USCG stations. The boats operate near shore and on inland waterways and are organized under the supervision of the Office of Boat Forces (OBF) into districts in the Atlantic and Pacific areas of the coastal United States. Each district is divided into sectors, which cover a total of 178 stations. The USCG maintains approximately 800 boats of 11 different types at those stations.

Each station requires a certain amount of boat-coverage hours. The individual station demands can range widely, from 250 to over 5,000 annual hours, with an average station demanding approximately 2,200 hours. In addition to the overall demand for hours, the majority of stations have various demands for specific mission hours such as heavy weather, tactical, pursuit, shallow water, and ice rescue hours.

To fulfill the stations' overall and specific demands for mission hours, boats of different types are allocated to the stations. Each boat is budgeted with a standard annual supply (capacity) of hours. For example, the motor lifeboat (MLB) and response boat–medium (RB-M) boat types have a standard supply of 600 hours each annually, whereas the response boat–small (RB-S) boat type is budgeted for only 500 hours; see Table 1. The total supply of hours for boats

Table 1. Names, abbreviations, and key parameters for the BAT model (SPC = special purpose craft).

Boat name	Abbreviation	Index	Number of boats B_t	Default hours per boat d_t	Fixed cost per boat f_t ($)	Variable cost of one hour v_t ($)
Motor lifeboat	MLB	0	106	600	36,951	120
SPC—Nearshore lifeboat	SPC-NLB	1	2	350	15,000	60
SPC—Heavy weather	SPC-HWX	2	4	350	15,000	120
Response boat–medium	RB-M	3	167	600	36,000	120
Response boat–small	RB-S	4	335	500	5,657	47
Response boat–small auxiliary	RBS-AUX	5	13	500	5,657	47
SPC—Law enforcement	SPC-LE	6	41	1,000	9,217	87
SPC—Shallow water	SPC-SW	7	47	500	4,390	63
SPC—Air	SPC-AIR	8	8	100	2,000	45
SPC—Ice	SPC-ICE	9	24	100	1,000	15
SPC—Skiff	SPC-SKF	10	56	100	500	15

assigned to a station would need to match the station's demand for mission hours and thus provide normal operational capabilities of the boat station to meet its mission requirements. However, starting from the early 2000s, the USCG recognized an existing disparity between the stations' demand hours and actual supply of hours provided by the boats at those stations. Moreover, the growing requirement to protect U.S. coastal areas after the September 11 attacks, and the subsequent increase in the number of boats at the stations, did not resolve this divergence.

We have developed, and the USCG Office of Boat Forces has applied, an integer linear programming model and Excel-based software called the boat allocation tool (BAT), which identifies optimal allocations of boat resources for the USCG boat stations. This paper presents an overview of the model, focusing on its unique features and meaningful generalizations. In addition, we discuss in depth the USCG adoption and application of the BAT model and its beneficial impact on real boat fleet performance.

1.1. Literature Review

The model and application we are discussing in this paper belong to a class of resource allocation models. We position our paper by reviewing three representative groups of literature sources: (a) general OR papers, (b) military research articles, and (c) U.S. Coast Guard research papers. We also discuss sources that apply the value-at-risk technique in conjunction with resource allocation models.

We identified several research papers in the general OR literature that employ analogous ideas and formulations to that which we describe in our paper. Gol'stejn and Dempe (2002) formulate a resource allocation problem using linear programming, with an objective function that minimizes the deviation of realized cost from a predetermined target cost; this objective is a simpler version of the objective we develop in our paper. Zhang et al. (2009) present a finite-horizon integer programming model for allocating operating room capacity, which minimizes inpatients' cost and the mismatch of supply and demand of operating room hours. In our paper, we also apply integer programming to minimize the mismatch of supply and demand of boat hours along with the total fleet operating cost.

A couple of military resource allocation applications are similar to the USCG project described in this paper. Billings (2005) formulates and solves an integer linear program that prescribes a minimum-cost load-and-unload schedule for U.S. Navy ships homeported in San Diego, subject to constraints on ship availability and port capabilities. These constraints are similar to the constraints in our model that describe boat availability and mission classification of stations. Zarybnisky (2003) allocates aircraft resources among sites that conduct reconnaissance, take strike actions, and gather battle damage information; our model is similar in the sense that stations' missions strongly influence which boats are assigned to those stations.

The U.S. Coast Guard has also applied resource allocation models in a variety of contexts. Brown et al. (1996) apply an integer programming model to schedule district cutter assignments for weekly patrol missions; the objective function either minimizes total assignment costs or minimizes the change from the previous scheduling period. An important aspect of their model formulation is that its constraints represent a variety of USCG Business Rules used for allocating cutters, such as assurance of patrol coverage, enforcement of equitable distribution of patrols, and restrictions on consecutive cutter assignments. We have applied a similar approach in our paper to ensure that the USCG Business Rules for boat allocations are satisfied. Deshpande et al. (2006) and Everingham et al. (2008) discuss the design and implementation, respectively, of models to improve the USCG aircraft service parts supply chain. These researchers formulate and implement an assignment model that merges two databases representing the supply of aircraft parts and the demand for maintenance. The model's objective is to minimize the delay between the times aircraft maintenance was requested and received. We have designed an objective function that is structurally similar to that chosen by these authors. It is also noteworthy that Everingham et al. (2008) utilize a large-scale LP solver (from Frontline Systems) to optimize their model in Microsoft Excel, an implementation choice also employed in our paper.

Narrowing the scope of our discussion, to the best of our knowledge there is only one paper, Radovilsky and Koermer (2007), that directly addresses the optimal allocation of small boat resources among USCG boat stations. The authors present an integer programming model that minimizes the deviation of supply and demand hours at boat stations, subject to two sets of constraints: (1) A minimum number of boats is required at each station, to ensure readiness, and (2) a maximum number of boat types is allowed at stations, motivated by minimizing stations' maintenance costs. The authors, using boat and station data from the USCG Pacific districts, demonstrated that the formulated model was capable of reducing excess capacity of boats at stations by 74%, lowering shortages of boats by 92%, decreasing the fleet size by more than 15%, and lowering the fleet cost by approximately 21%. Despite the fact that this model exhibits notable results, it also contains substantial limitations: (1) The model defines boats (not hours) as a primitive variable, which limits the model's ability to match boat supply and station demand hours; (2) it includes only a subset of the existing USCG boat types; (3) it does not incorporate a variety of critically important USCG Business Rules required in managing boat allocations; and (4) the model does not allow sharing limited boat resources between neighboring stations. In our paper we build upon this model to specifically address these weaknesses.

One of the main contributions of our paper is the value-at-risk (VAR) measurement and analysis of the effects of demand uncertainty on USCG allocations. Therefore, we briefly discuss literature sources on the VAR technique. In general, VAR is a popular tool for measuring a risk of loss on a financial portfolio of assets in a given period of time. As a measure of financial risk, VAR has been extensively discussed in the financial and banking research literature (e.g., see Hull 2000). However, the VAR approach is also employed for measuring and analyzing risk management decisions in nonfinancial applications. Wu and Olson (2010) use the VAR technique in assessing supply risks as a part of an enterprise risk management system. Kauffman and Sougstad (2007) propose a VAR-based model to advise vendors on setting optimal parameters for IT service contracts subject to acceptable levels of risk. Ravindran et al. (2010) develop a supplier selection model that describes operational disruption risk using VAR constructs. To the best of our knowledge, our application of VAR, in conjunction with robust optimization ideas, has not been applied in any USCG-related resource allocation model nor in any existing USCG risk management technique. More generally, Birge and Louveaux (2011) is a good reference for the probabilistic constraints that we utilize to model the VAR ideas, especially §3.2 (probabilistic or chance constraints).

1.2. Contributions

On the practical side, we have designed an integer linear programming model, called the boat allocation tool (BAT) model, for the USCG allocation problem of matching the supply of boat hours with the demand of stations. The BAT model is the first USCG boat allocation model that applies optimization techniques and incorporates new operational constraints that arose after the tragic events of September 11, 2001. The BAT model is currently being used for allocating the entire USCG fleet of boats nationwide, and in this paper we describe in detail the USCG's process of adopting and using our model and its significant performance results. On the modeling side, we have also provided a number of contributions. We have introduced a model for the optimal *geographical sharing* of boat resources. We have defined boat *hours*, rather than boats, as a primitive decision variable in our model, which adds operational flexibility in the sense that each boat can supply different types and amounts of supply hours and each station can demand different types and amounts of demand hours. Finally, we introduce risk management of boat shortages into our model, by allowing stations to have uncertain demand, modeled as a random variable with known summary statistics but unknown distribution. Applying value-at-risk (VAR) and robust optimization concepts, we derive *linear* inequality constraints for managing risk so that our generalized model remains a linear (rather than quadratic) one. To the best of our knowledge, the unique combination of VAR and robust optimization ideas has not been applied in any resource allocation model nor in any existing USCG risk management technique, and their application is one of the main contributions of our paper.

The paper is organized as follows. In §2 we provide a complete description of the BAT model, including a generalized version that models uncertain demand. Section 3 reports on the USCG adoption of the BAT model, the modification and implementation of the model recommendations, and the substantial improvements of USCG performance metrics. Section 4 presents three parametric studies, analyzing (1) the effect of changing a critical sharing parameter, (2) the trade-offs between the different objectives of the model, and (3) the trade-off between risk management and optimality. We provide concluding thoughts in §5.

2. The Boat Allocation Tool (BAT) Model

2.1. Model Formulation

The U.S. Coast Guard application described in this paper is a resource allocation problem. The task is to allocate a fleet of approximately 800 boats of 11 different types to 178 Coast Guard stations. The BAT model, an integer linear program, uses the following sets:

- $t \in T$: Set of boat types.
- $s \in S$: Set of stations.
- $m \in \mathcal{M}$: Set of specialized station missions.
- $s \in S_m$: Set of stations that are assigned mission $m \in \mathcal{M}$.
- $t \in T_m$: Set of boat types that are appropriate for mission $m \in \mathcal{M}$.

- $t \in T_s$: Set of boat types allowed at station $s \in S$.
- $t \in T_c$: Set of critical boat types, whose presence requires the presence of another boat type.

The BAT model parameters are as follows:

- B_t: Available number of boats of type t.
- d_t: Yearly default capacity, in hours, of a type t boat.
- H_s: Yearly demand, in hours, of station s.
- f_t: Yearly fixed cost of utilizing one boat of type t.
- v_t: Variable cost of utilizing one boat of type t for one hour.
- b_m: Minimum number of boats required to satisfy mission m at a station.
- H_{As}: Yearly demand for a class $A \subseteq T$ of boats at station s, in hours.
- m_t: Multiplier (for d_t) to provide minimum allowable hours assigned to a boat of type t.
- M_t: Multiplier (for d_t) to provide maximum allowable hours assigned to a boat of type t.
- $d_{s,s'}$: Distance between stations s and s'.
- γ: Distance threshold to allow MLB sharing.
- $R = \{(s, s') \in S \times S: s < s' \wedge d_{s,s'} \leqslant \gamma\}$: Set of pairs of stations eligible to share MLB boats.

Table 1 summarizes the boat types and key parameters used in our model. We briefly elaborate on how certain parameters are computed. The yearly demands H_s and H_{As} of station s (and boat class A) are calculated using five-year (2005–2009) averages of historical data; however, if a change in workload is anticipated, a BAT model user can easily update the relevant parameter. The yearly fixed cost f_t of using a single boat of type t is calculated by averaging five years of station-level maintenance costs, electronic navigational chart (map) updates, and support overhead. The hourly cost v_t of operating boat type t is calculated by averaging five years of parts and materials cost, fire/flooding/collision repairs, maintenance services transportation costs, and a centralized outboard engine overhaul program. For each boat type, the default annual boat capacity d_t, measured in hours, is derived from the expected lifespan of boats (e.g., 25 years for the MLB) and their depreciation over the lifespan, which is mandated by the USCG budget.

The BAT model's variables are as follows:

- x_{st}: Integer number of boats of type t allocated to station s, $\forall s, t$.
- y_{st}: Binary variable indicating whether or not boat type t is utilized at station s, $\forall s, t$.
- h_{st}: Number of hours of boat type t assigned to station s, $\forall s, t$.
- $q_{s,s'}$: Binary variable indicating whether or not stations s and s' share an MLB boat that is hosted by station s' (station s has no MLBs), $\forall (s, s') \in R$.
- $r_{s',s}$: Binary variable indicating whether or not stations s' and s share an MLB boat that is hosted by station s' (station s has no MLBs), $\forall (s', s) \in R$.

Finally, we present the BAT model and describe the formulation in detail.

$$\min w_1 \sum_{s\in S} \left| \sum_{t\in T} h_{st} - H_s \right| + w_2 \sum_{s\in S}\sum_{t\in T} y_{st} + w_3 \sum_{s\in S}\sum_{t\in T}(f_t x_{st} + v_t h_{st}), \tag{1}$$

$$\sum_{s\in S} x_{st} \leqslant B_t, \quad \forall t \in T, \tag{2}$$

$$x_{st} \leqslant B_t y_{st}, \quad \forall s \in S,\ t \in T, \tag{3}$$

$$\sum_{t\in T_m} x_{st} \geqslant b_m, \quad \forall s \in S_m,\ m \in M, \tag{4}$$

$$\sum_{t\in T\setminus T_s} x_{st} = 0, \quad \forall s \in S, \tag{5}$$

$$\sum_{t\in T} x_{st} \geqslant 2, \quad \forall s \in S, \tag{6}$$

$$\sum_{t\in T\setminus T_c} x_{st} \geqslant x_{s\tau}/B_\tau, \quad \forall s \in S,\ \tau \in T_c, \tag{7}$$

$$h_{st} \leqslant d_t B_t x_{st}, \quad \forall s \in S,\ t \in T, \quad \sum_{s\in S} h_{st} \leqslant d_t B_t, \quad \forall t \in T, \tag{8}$$

$$m_t d_t x_{st} \leqslant h_{st} \leqslant M_t d_t x_{st}, \quad \forall s \in S,\ t \in T, \tag{9}$$

$$\sum_{t\in A_i} h_{st} \geqslant H_{A_i s}, \quad \forall s \in S,\ i = 1, \ldots, 6, \tag{10}$$

$$x_{s0} \geqslant 1 - \sum_{s':(s,s')\in R} q_{s,s'} - \sum_{s':(s',s)\in R} r_{s',s}, \quad s \in S, \tag{11}$$

$$\sum_{s':(s,s')\in R} (q_{s,s'} + r_{s,s'}) + \sum_{s':(s',s)\in R} (q_{s',s} + r_{s',s}) \leqslant 1, \quad s \in S, \tag{12}$$

$$x_{st} \in \mathbb{N},\ y_{st} \in \{0, 1\},\ h_{st} \geqslant 0, \quad \forall s \in S,\ t \in T. \tag{13}$$

The multiobjective function presented in Expression (1) contains three terms. The first term represents the deviation of the overall supply of hours from overall demand of hours. The second term measures the number of types of boats at each station, as fewer boat types at a station translate directly into lower boat maintenance costs. The third term represents total fleet operating cost. To provide flexibility in assigning importance to the three objectives, a BAT model user can define inputs $w_1, w_2, w_3 \geqslant 0$, where $w_1 + w_2 + w_3 = 1$; §4.2 reports the results of parametric experiments that provide guidance for choosing these weights. Finally, the absolute value term is linearized by the standard technique of adding supplemental variables.

Constraints (2)–(3) model the fact that the number of boats utilized can be at most the supply of boats and establishes the relationship between the number of boats x and respective binary variables y, respectively. The remaining constraints model USCG Business Rules, which represent long-standing USCG policy.

Constraint (4) makes sure that stations $s \in S_m$ that are assigned mission $m \in M$ are given enough boats b_m, of qualified types T_m, to satisfy the mission. For example, the USCG defines certain stations to have a tactical, or law enforcement, mission. These stations must have at least two boats capable of performing law enforcement activities, which are the MLB, RB-M, RB-S, and SPC-LE. Other mission classifications include pursuit missions, which require fast boats, and cold weather rescue missions. Note that missions, of which there are 15, are not necessarily mutually exclusive, and many stations are responsible for multiple missions. A related set of business rules state that certain boat types are not allowed at certain stations. Constraint (5) makes sure that a station only gets boats that are allowed at the station.

Next, we present the business rules' approach to the management of various risks. Any station must have at least two boats, of any type, as represented in Constraint (6). This business rule is motivated by the fact that not all boats are operational 100% of the time, and having at least two boats increases the chances that the station has at least one operational boat at all times. The USCG also introduced business rules designed to diversify the boat allocation, especially around critical boat types. For example, if an MLB is used at a station, then another boat type must also be present at the station. Constraint (7) makes sure that if a boat of type $t \in T_C$ is present, then another boat type $t \notin T_c$ must also be present.

The first part of Constraint (8) ensures that the BAT allocation assigns only h_{st} hours of type t to station s if there are boats of that type at that station. The second part of Constraint (8) makes sure the BAT model does not assign more hours overall than there are available, with respect to the default hours d_t. Constraint (9), required by the business rules, establishes limits regarding how many hours can be assigned per boat of type t, as a function of the default hours per boat d_t. The hours per boat of type t at station s, assuming $x_{st} > 0$, is h_{st}/x_{st}, which might be much less or much more than d_t. For each boat type t, the USCG provided two multipliers $0 < m_t < 1$ and $M_t > 1$ that preclude extreme fluctuations. The USCG determined these multipliers by examining historical boat usage data. For example, historically, an MLB boat would always use at least half its allocated hours and, with special OBF permission and additional maintenance, would use at most 50% more than its allocated hours. Thus, the MLB boat type was assigned the values $m_0 = 0.5$ and $M_0 = 1.5$.

The BAT model also eliminates supply shortages for critical demands. For all stations s, the USCG provided H_{As}, the number of hours demanded for a class $A \subseteq T$ of boats. In particular, there are six possible values for A:

$A_1 =$ "Big boats" $=$ {MLB, RB-M, SPC-NLB},

$A_2 =$ "Tactical" $=$ {RB-M, RB-S, MLB, SPC-LE},

$A_3 =$ "Pursuit" $=$ {RB-M, SPC-LE},

$A_4 =$ "Shallow Water" $=$ {SPC-SW, SPC-SKF},

$A_5 =$ "Ice Rescue Long Haul" $=$ {SPC-AIR},

$A_6 =$ "Ice Rescue Short Haul" $=$ {SPC-ICE},

where boat supply shortages are precluded, as modeled in Constraint (10). Note that these special groupings of demand hours are not collectively exhaustive and shortages at any station are still possible.

Finally, we describe an intrinsic supply shortage problem and our novel modeling solution of optimally sharing resources. The U.S. Coast Guard project revealed that MLB boats, a critical boat type required by many stations, must be shared, because the supply was not sufficient to meet demand. The USCG stipulated that if two stations s_1 and s_2 each require an MLB boat, then they are eligible to *share* one boat, but only if they are close enough to each other geographically. In particular, if the distance, defined as the shortest path length that a boat can traverse (in miles) between two stations s_1 and s_2, is less than a user-defined threshold γ, then station s_1's MLB boat can be used to satisfy station s_2's MLB coverage needs (or vice versa). The USCG determines the distance between two adjacent stations by using a search-and-rescue map to manually calculate the shortest distance by water between the two stations. Constraints (11)–(12) represent our modeling solution.

As the threshold γ increases, the responsiveness of the boats decreases, because they have to travel a longer distance on average. Therefore, the minimum value of γ that allows a feasible boat allocation was of interest to the USCG. It is straightforward to verify that the feasible region of the BAT model increases with γ, and we were able to apply a simple binary search to find that $\gamma^* = 28$ miles is the smallest value that still allows a feasible allocation. This concept of *optimally sharing* the supply of a scarce resource (MLB boats) was a new and valuable planning tool for the USCG. Previous to the BAT model, sharing was not a part of the boat allocation process, and the business rules were consistently violated (specifically, the rule to have two big boats at stations that experience heavy weather).

2.1.1. Post Processing. After the BAT model is optimized, supply hours are then evenly assigned to boats as follows: assuming $x_{st} > 0$, each boat of type t at station s will get h_{st}/x_{st} supply hours. Shared MLB boats are allocated hours only at their assigned stations, and an MLB boat operating at a station that is sharing it will use hours from another boat at that station.

2.2. BAT Model Generalization: Risk Management via Value-at-Risk

One main drawback of the BAT model, as requested by the USCG, is the omission of any measure of demand variability. The risk of a mismatch between supply and demand, especially boat shortages, resulting from an inaccurate demand estimate is not explicitly captured. The business rules do incorporate risk management indirectly in two ways. First, supply variability is addressed by requiring each station to have at least two boats (Constraint (6) in the BAT model formulation), motivated by the fact that boats are not available 100% of the time. Indeed, each boat type has a capability rate, essentially a likelihood of being operational, ranging from 76% to 85%; requiring at least two boats results in a station having an operational boat at least 94.2% ($= 1 - (1 - 0.76)^2$) of the time. Second, high demand variability exists at a small number of stations (to remain unnamed at USCG request), and the USCG compensates by increasing the point estimate of demand to induce a "safety stock" of boat hours at the station. More specifically, the USCG will update the demand estimate H_s according to a trend analysis; i.e., if historical demand is trending strongly, the demand estimate (an average) is updated according to the trend analysis. However, in contrast to the USCG's approach to modeling supply uncertainty, this trending method does not utilize any quantitative measure of uncertainty and this section presents our generalization of the BAT model to incorporate demand uncertainty, as measured by a standard deviation, more rigorously.

Our approach combines the financial concept of value-at-risk (VAR) and techniques of robust optimization. Traditionally, VAR is utilized in financial applications and is defined as the minimum value θ, where the probability of a loss greater than θ is at most a user-defined confidence level ϵ; see Hull (2000) for further financial details. While VAR constraints are equivalent to chance constraints (see Sarykalin et al. 2008), the VAR parameters θ and ϵ allow us to conveniently describe user-defined risk tolerances, in the form of an acceptable yearly boat shortage θ at a station and a corresponding acceptable likelihood ϵ that the shortage limit is breached. For example, the manager at a boat station is willing to accept only at most a probability $\epsilon = 0.05$ of a shortage of at least 500 hours ($\theta = 500$). Robust optimization allows us to avoid dependence on specific distributions of demand, and our user-defined risk tolerances must be satisfied for all distributions with a given mean and standard deviation. Finally, the combination of value-at-risk and robust optimization allows users to incorporate demand uncertainty without increasing the complexity of the BAT model—it remains a linear integer program, as opposed to a quadratic one.

2.2.1. Modeling Approach. Let $S_{\text{uncertain}}$ denote the set of stations that do not have an accurate point forecast. The choice of $S_{\text{uncertain}}$ can reflect actual limitations of forecasting ability or can instead depend on the importance of avoiding shortages at certain important stations (i.e., near highly populated cities at risk of terrorist attacks).

We model each H_s, $s \in S_{\text{uncertain}}$, as a random variable with a mean of μ_s and a standard deviation of σ_s. However, we assume that the distribution of H_s is not known. Indeed, insufficient data at the USCG precluded an accurate estimation of the demand distribution for all stations. A benefit of our modeling approach is that demand forecasting can be accomplished via only two numbers (μ_s, σ_s) rather than a complete probability distribution, the determination of which is a more daunting task in practice.

The risk of a boat shortage forms the core of USCG risk management. The shortage (random variable) at station $s \in S_{\text{uncertain}}$ is defined as $H_s - \sum_{t\in T} h_{st}$. To better manage these shortages, our modeling approach utilizes two user-defined parameters per station. The threshold parameter $\theta_s > 0$ represents an upper limit of acceptable shortage hours, and $\epsilon_s \in (0, 1)$ represents the acceptable likelihood that the upper limit is breached; in probabilistic terms, $\mathbf{P}(H_s - \sum_{t\in T} h_{st} \geqslant \theta_s) \leqslant \epsilon_s$. Because a distribution is not prescribed, we require that the definitions of θ_s and ϵ_s hold for all distributions F of H_s that have the given mean and standard deviation. Mathematically,

$$\max_{F\sim(\mu_s,\,\sigma_s)} \mathbf{P}\left(H_s - \sum_{t\in T} h_{st} \geqslant \theta_s\right) \leqslant \epsilon_s \quad \forall\, s \in S_{\text{uncertain}}, \tag{14}$$

where $\max_{F\sim(\mu_s,\,\sigma_s)}$ represents the maximum over all distributions F with a mean of μ_s and a standard deviation of σ_s. In other words, a planner can guarantee that the shortage at station s will be at most θ_s with at least $(1-\epsilon_s)$ probability, for any valid distribution F. These are robust value-at-risk (VAR) constraints, where the "value" is the shortage of supply at a given station and the "risk" is the probability of a shortage exceeding a predetermined threshold.

The inherent conservatism of this approach is indeed desirable in today's society, which is vulnerable to low-probability, high-impact events. Furthermore, there is no assumption of independence among the H_s random variables. Indeed, the worst-case distribution is determined by considering all distributions with the given mean and standard deviation, including those marginal distributions that are dependent on other stations' random variables.

Finally, note that there are a number of ways to generate the parameters discussed in this section. The two-point forecast (μ_s, σ_s) can be generated using historical data. If data are limited, μ_s can represent a qualitative estimate of demand and σ_s is chosen to represent the confidence in that estimate (i.e., low values of σ_s for high confidence). The parameter θ_s will clearly depend on the station s and its assigned missions. Critical missions, such as law enforcement, would require a lower threshold θ_s, whereas less critical missions, such as general patrol, might allow a higher threshold. The ϵ_s parameters depend on the risk tolerances of a decision maker but are generally equal to 1%–5%; e.g., see Pearson (2002).

2.2.2. Linearizing the Constraints. We next show that Constraint (14) can be represented using linear inequality constraints, which facilitate their incorporation into the BAT model. We employ the following one-sided version of the Chebyshev inequality for a random variable Y with mean μ, variance σ^2, and constant $\delta > 0$:

$$\mathbb{P}(Y \geqslant (1+\delta)\mu) \leqslant \frac{\sigma^2}{\sigma^2 + \mu^2\delta^2}.$$

Note that this inequality is tight in the sense that for any $\delta > 0$, there exists a distribution F for Y, with the given mean μ and standard deviation σ, such that $\mathbb{P}(Y \geqslant (1+\delta)\mu) = \sigma^2/(\sigma^2 + \mu^2\delta^2)$. We apply the Chebyshev inequality to the probability expression in the left-hand side of Constraint (14) with $\delta = (\theta_s + \sum_{t\in T} h_{st} - \mu_s)/\mu_s$ to see that

$$\mathbb{P}\left(H_s \geqslant \theta_s + \sum_{t\in T} h_{st}\right) \leqslant \frac{\sigma_s^2}{\sigma_s^2 + (\theta_s + \sum_{t\in T} h_{st} - \mu_s)^2}, \quad s \in S_{\text{uncertain}}. \tag{15}$$

Because the Chebyshev inequality is tight, requiring that the right-hand side of Constraint (14), namely ϵ_s, be greater than or equal to the right-hand side of Inequality (15) gives us

$$\theta_s + \sum_{t\in T} h_{st} \geqslant \mu_s + \sigma_s \sqrt{\left(\frac{1}{\epsilon_s} - 1\right)}, \quad \forall\, s \in S_{\text{uncertain}}, \tag{16}$$

a linear inequality that is equivalent to Constraint (14). Intuitively, this constraint states that the available supply hours (sum of acceptable shortage and assigned hours) must be sufficient to cover at least a user-defined standard deviation multiple above the mean demand.

2.2.3. Incorporating Value-at-Risk Constraints Into the BAT Model. To add the value-at-risk constraints, some changes to the BAT model are required. First, objective function (1) must be modified. For $s \in S_{\text{uncertain}}$, the first component of the original objective function is $|\sum_{t\in T} h_{st} - H_s|$, a random variable. In these expressions, we replace the random variable H_s with its mean μ_s, which results in a new objective function:

$$w_1\left(\sum_{s\in S\setminus S_{\text{uncertain}}} \left|\sum_{t\in T} h_{st} - H_s\right| + \sum_{s\in S_{\text{uncertain}}} \left|\sum_{t\in T} h_{st} - \mu_s\right|\right) + w_2 \sum_{s\in S}\sum_{t\in T} y_{st} + w_3 \sum_{s\in S}\sum_{t\in T}\left(f_t x_{st} + v_t h_{st}\right). \tag{17}$$

Note that the H_s parameters appear nowhere else in the deterministic model previously described; the specific demands in Constraint (10) are still known, as they are absolute lower bounds on demand. Second, we add the VAR Constraint (16) to the BAT model; all other Constraints (2)–(13) in the formulation remain the same. The BAT model that incorporates both these changes is denoted the BAT-VAR model.

3. Implementation and Impact of the BAT Model

In this section we report on the practical implementation and impact of the model presented in §2.1. The model generalization in §2.2 is the subject of a current ongoing project and has not yet been implemented.

As mentioned previously, the USCG requested that we implement the BAT model in Excel, which is the Coast Guard's standard tool for managing and planning boat resources and their allocations. In particular, we utilized Frontline Systems' Premium Solver Platform for Excel V9.5 and Standard Large-Scale LP Solver Engine V9.0 Windows as the optimization engine. To streamline and ease the implementation and utilization of the BAT model with these software packages, we developed an Excel-based decision support system (DSS) .

The DSS allows a BAT user to modify all parameters in Table 1, as well as general and specific demand hours (H_s and H_{As}). However, for the duration of the project, these parameters remained fixed at the values listed in Table 1. The user can also modify the importance of each optimization's criteria—minimum deviation of demanded and supplied hours, minimum number of boat types per station, and total allocation cost—by varying their respective weights in the BAT model objective. Once the BAT model is optimized, the DSS will display the optimal number of boats of various types at each station, respective amount of boat hours to be allocated to each boat and performance metrics.

For training the USCG users on how to apply the BAT model and its Excel-based DSS, we have developed a technical manual—the BAT User's Guide. This guide provides a detailed description of the BAT model, its input requirements, and output results. The guide also presents a step-by-step implementation process and offers in-depth instruction on using the DSS. With the help of the BAT User's Guide, operations research specialists were able to quickly implement the BAT model and start utilizing it.

3.1. USCG Implementation

The BAT model described in §2.1 is currently being utilized by the platform division (PD) of the USCG's Office of Boat Forces (OBF) for optimizing boat allocations among the boat stations. A PD group, consisting of four civilian and military personnel with operations research training, is the primary BAT model user and program manager for all USCG boats. This PD group directly communicates with district boat managers, who are in charge of boat allocations in their own districts. The PD is also responsible for coordinating the rearrangement of boats, which is accomplished through the issuance of formal military messages to the districts. Finally, each station has a boat office that communicates directly with its district manager.

The USCG-defined benchmark for evaluating the BAT model is the original USCG boat allocation for 2010–2015, which we denote as the original allocation. This allocation was created in September 2009 to incorporate the replacement of UTB boats, which are over 40 years old, with new RB-M boats. The USCG has a contract to have 30 new RB-M boats delivered each year from 2010 through 2015, which motivates the 2010–2015 time frame. From a modeling perspective, the final 2015 number of RB-M boats and all other boat numbers are essentially fixed, barring another September 11th type event (the USCG significantly increased their boat numbers after 9/11).

After optimizing the BAT model, we created the recommendations for the OBF on boat allocations, which we denote as the BAT allocation. These recommendations were utilized by the OBF to create a new boat allocation, denoted the implemented allocation, to replace the original allocation. The reasons for the differences between the BAT and implemented allocations are discussed below in §3.4. Boat re-assignments, prescribed by the implemented allocation, started in 2010. The USCG is planning a six-year implementation, a domino process dictated by the deliveries of new RB-M boats. In 2010, the USCG received 30 new RB-M boats, which were delivered to stations according to the implemented allocation; when a station receives a new RB-M, the station's remaining boat allocation is updated according to the implemented allocation. The USCG expects 80% of the implemented allocation to be completed by the end of 2012, 90% by the end of 2013 and 100% by 2015. The large jump in completion percentage in 2012 is due to the default hours per boat d_t being relaxed to implement the variable assignment hours h_{st}. Finally, the major boat reallocation due to the BAT model, consisting of at least 30 boat changes per year (i.e., delivery of RB-M boats and corresponding changes), is much larger than historical reallocations, which occurred due to station officers' requests with the OBF and consisted of six changes per year, on average.

3.2. Internal Resistance

Within the USCG, there was some resistance to the changes brought about by the BAT model. Occasionally, a station officer or district manager will desire to allocate boats differently than the implemented allocation. This situation usually occurs due to the conflict between centralized and decentralized decision making. For example, a station manager will want the best boat X to meet his mission while the central OBF authority states that boat Y, which might not be the best, is still sufficient for the mission. As another example, some station officers wanted to have both an MLB and RB-M at the same station, which is precluded by the official USCG Business Rules. Another issue is that some district managers and station officers would prefer to keep standard hours defined for each boat type as opposed to the more rigorous variable hours suggested by the BAT model. According to the OBF, conservatism, resistance to change and additional workload (i.e., tracking hourly boat-specific assignments) drive this reluctance. Finally, the OBF stated that only a small number of stations attempted to diverge from the official results, and most comply with the recommendations.

The OBF considers it a responsibility to educate the district and station personnel that the allocations must follow the business rules, as embodied in the BAT model. The OBF is also willing to listen to these managers and officers, and incorporate their suggestions, if possible, into the model. For example, some station managers will argue for a larger boat type than that assigned to them, because of "trailer" requirements, which are not incorporated into the BAT Model. These requirements mean that a small boat will be hitched as a trailer to a larger boat for transportation across long distances. Consequently, if some resistance is well reasoned, the OBF will make minor changes to the boat allocation.

3.3. Performance Metrics

We developed the following 10 performance metrics to quantify and compare the boat allocation improvements in the BAT and implemented allocations over the original allocation. We utilize the notation $\mathbf{1}\{E\}$ to indicate whether an expression E is true or not (i.e., $\mathbf{1}\{E\}=1$ if E is true and $\mathbf{1}\{E\}=0$ if E is false). Also, recall that there are a total of 178 USCG stations. Finally, note that these metrics are generally designed to measure how well the supply of boats, as embodied by the x_{st} and h_{st} variables, match the stations' demand of hours H_s, which are USCG-supplied parameters exogenous to the BAT model. Furthermore, these metrics mirror traditional metrics that the USCG utilized to evaluate their past boat allocations and were created at the USCG's request.

P1: Total size of utilized boat fleet

$$\triangleq \sum_{s\in S}\sum_{t} x_{st}.$$

P2: Percentage of stations with excess hours

$$\triangleq \frac{\sum_{s\in S}\mathbf{1}\{\sum_{t\in T} h_{st} > H_s\}}{178}.$$

P3: Percentage of stations with a shortage of hours

$$\triangleq \frac{\sum_{s\in S}\mathbf{1}\{H_s > \sum_{t\in T} h_{st}\}}{178}.$$

P4: Average excess hours per station with an excess

$$\triangleq \frac{\sum_{s\in S}(\sum_{t\in T} h_{st} - H_s)\mathbf{1}\{\sum_{t\in T} h_{st} > H_s\}}{\sum_{s\in S}\mathbf{1}\{\sum_{t\in T} h_{st} > H_s\}}.$$

P5: Average shortage hours per station with an shortage

$$\triangleq \frac{\sum_{s\in S}(H_s - \sum_{t\in T} h_{st})\mathbf{1}\{H_s > \sum_{t\in T} h_{st}\}}{\sum_{s\in S}\mathbf{1}\{H_s > \sum_{t\in T} h_{st}\}}.$$

P6: Percentage of stations with more than two boat types

$$\triangleq \frac{\sum_{s\in S}\mathbf{1}\{\sum_{t\in T} y_{st} > 2\}}{178}.$$

P7: Average number of boat types per station

$$\triangleq \frac{\sum_{s\in S}\sum_{t\in T} y_{st}}{178}.$$

Table 2. Differences between original, BAT, and implemented allocations.

Boat type	Original allocation	BAT allocation	Implemented allocation
MLB	106	102	102
SPC-NLB	3	2	3
SPC-HWX	4	0	4
RB-M	166	166	158
RB-S	360	208	318
RBS-AUX	10	10	10
SPC-LE	33	26	20
SPC-SW	47	47	47
SPC-AIR	8	8	12
SPC-ICE	0	24	0
SPC-SKF	67	29	42
Totals	804	622	716

P8: Fleet operating cost

$$\triangleq \sum_{s\in S}\sum_{t\in T}(f_t x_{st} + v_t h_{st}).$$

P9: Capacity utilization $\triangleq$(supply utilized)/(supply available):
—Supply available $\triangleq \sum_{t\in T} d_t B_t$.
—Supply utilized $\triangleq$ supply available − excess hours:
∗Original USCG allocation, excess hours
$\triangleq \sum_{s\in S}(\sum_{t\in T} d_t x_{st} - H_s)\mathbf{1}\{\sum_{t\in T} d_t x_{st} > H_s\}$.
∗BAT allocation, excess hours
$\triangleq \sum_{s\in S}(\sum_{t\in T} h_{st} - H_s)\mathbf{1}\{\sum_{t\in T} h_{st} > H_s\}$.

P10: Demand shortfall rate $\triangleq$ (shortage hours)/(total demand hours):
—Total demand hours $\triangleq \sum_{s\in S} H_s$.
∗Original USCG allocation, shortage hours
$\triangleq \sum_{s\in S}(H_s - \sum_{t\in T} d_t x_{st})\mathbf{1}\{H_s > \sum_{t\in T} d_t x_{st}\}$.
∗BAT allocation, shortage hours
$\triangleq \sum_{s\in S}(H_s - \sum_{t\in T} h_{st})\mathbf{1}\{H_s > \sum_{t\in T} h_{st}\}$.

3.4. Discussion of the BAT Model's Impact

The practical impact of the model is significant. In Table 2, we summarize the composition of the boat fleets in the original allocation, the BAT allocation, using the USCG defined weights $w_1 = 0.95$ and $w_2 = w_3 = 0.025$, and the implemented allocation.

Focusing on the original and implemented allocation columns of Table 2, we see a significant decrease in the number of many types of boats (e.g., RB-S, SPC-LE, SPC-SKF) utilized in the field. The reason for these decreases is primarily the BAT model's ability to assign optimal supply hours to boats (h_{st}) rather than using the default hours (d_t). This greater planning flexibility allows many boats to satisfy more demand hours, reducing the required number of boats. Therefore, the hourly assignment variables h_{st} not only allowed a better match of supply and demand, but they also resulted in a reduction of the number of boats required. These improvements clearly depend on the values of m_t and M_t, $t \in T$, which determine how far a boat's allocation of hours differs from the default hours d_t. Decreasing m_t and/or increasing M_t will result in a better match of supply and demand and an even smaller fleet size. Therefore, m_t should be as small as possible and M_t as large as possible to maximize the impact of the BAT model. Extending these parameters beyond their derivation from boat-specific historical data (c.f., §2.1), which describes suboptimal boat behavior, will increase planning flexibility and result in a better match of supply and demand, with fewer boats.

Next, we discuss the discrepancies between the BAT and implemented allocations in Table 2. The USCG modified the BAT allocation for approximately 25% of the stations, which increased the necessary fleet size by 15%. Many of the changes in the implemented allocation were minor and resulted from station-specific conditions that the BAT model did not incorporate. The reasons for many of these changes were related to the *degree* of heavy weather or sea roughness a station receives. In the BAT model, a station that experiences heavy weather is assigned a "big boat," such as an MLB or RB-M (c.f., definition of set A_1, in §2.1). However, the BAT model does not differentiate between the degrees of heavy weather a station experiences, or the different capabilities of the boat types to handle these degrees. As a result, the USCG replaced MLBs at certain stations with RB-Ms, and at other stations the opposite occurred. The overall reduction in RB-M boats is also a result of the "degree" of heavy weather a station experiences, where the expensive RB-M is replaced with a cheaper RB-S, at eligible stations. We suggested to the USCG to create a new business rule that would address the degree of heavy weather or sea roughness, incorporate this rule into the BAT model, and then re-optimize the boat allocations. However, rather than increase the complexity of the BAT model, the USCG preferred to tweak the final allocation using these considerations.

Two major differences between the BAT and implemented allocations are substantial increases in the number of RB-S boats (53% increase) and SPC-SKF boats (45% increase). The increase in RB-S boats is mainly due to a readiness rule that was implemented by the USCG after the BAT model was delivered. In particular, each RB-S boat is historically operational 80% of the time (the 20% reflecting maintenance time), but the USCG stipulated that a large class of stations must have an operational RB-S boat available 99.0% of the time. Since this new RB-S constraint was not incorporated into the BAT model, additional RB-S boats were required at a multitude of stations, resulting in this discrepancy between the BAT and implemented allocations. SPC-SKF boats, on the other hand, are primarily used to respond to flood events, which are unpredictable. The modification of SPC-SKF assignments reflects real-time information about flood potential at different areas; indeed, the allocation of SPC-SKF boats can change at any time. In fact, due to the unpredictable nature of flooding, the USCG debated about the inclusion of the SPC-SKF

Table 3. Performance metrics for BAT model.

	Original allocation	BAT allocation	Implemented allocation
P1: Total size of utilized boat fleet	804	622	716
P2: Percentage of stations with excess hours (%)	61.2	1.7	41.6
P3: Percentage of stations with a shortage of hours (%)	38.8	0.0	1.1
P4: Average excess hours per station with an excess	556.3	102.0	209.8
P5: Average shortage hours per stations with a shortage	563.1	0.0	70.0
P6: Percentage of stations with more than two boat types (%)	37.6	30.9	30.9
P7: Average number of boat types per station	3.1	2.3	2.2
P8: Fleet operating cost ($)	45,648,887	43,379,851	43,541,610
P9: Capacity utilization (%)	85.3	99.0	96.2
P10: Demand shortfall rate (%)	9.90	0.00	0.04

boat type into the BAT model, ultimately favoring inclusion. Finally, the USCG did make one major change to the BAT model recommendations: the SPC-ICE boat type was eliminated from the implemented allocation, as the "ice rescue short haul" missions (set A_6) are now being satisfied by nonboat resources. Overall, to avoid increasing the complexity of the BAT model, the USCG decided to adjust the boat allocations rather than incorporate the new changes into the model and then re-optimize it.

In Table 3, we summarize the performance metrics P1–P10 obtained from the original, BAT, and implemented allocations. We first compare the implemented and original allocations. Most notably, the implemented allocation, based heavily on the BAT allocation, substantially improves the performance metrics when compared with the original allocation. Because the total required fleet size decreased by 88 boats in the implemented allocation versus the original allocation, the performance metrics show that the boats are now being utilized more effectively, with respect to the three objectives of the BAT model. First, the mismatch between supply and demand is significantly improved. Indeed, there is a substantial reduction in the percentage of stations with shortage or excess hours, and the average amount of shortage or excess hours per station is also extensively reduced; the greatest improvements are for the shortage Metrics P3 and P5. Second, the percentage of stations with more than two boat types and the average number of boat types per station have a healthy reduction in the implemented allocation as compared with the original allocation (Metrics P6 and P7). Third, the cost savings are 4.6% (Metric P8), albeit lower than that of the fleet reduction (10.9%). This was because the fleet operating cost's decrease is associated with a reduction in the total fixed cost $\sum_{s\in S,\, t\in T} f_t x_{st}$, which depends on the number of boats. At the same time, the fleet operating cost's reduction is not affected by the total variable cost $\sum_{s\in S,\, t\in T} v_t h_{st}$, which depends only on the assigned number of supply hours h_{st}, and not the fleet size changes.

The increase in capacity utilization (Metric P9) signifies that the implemented allocation is more effectively using a smaller fleet of boats; in other words, waste is significantly reduced. The drastic reduction in the demand shortfall rate (Metric P10) substantially increases the USCG's service level. Almost all, as opposed to 90 out of 100, USCG customers will be served. As a result of the BAT model, there is substantially less need to cancel patrols or operate a boat longer than its maintenance schedule prescribes, and there will almost always be a boat ready to assist distressed swimmers and boaters.

Comparing the performance metrics of the implemented and BAT allocations, we need to point out that the USCG modification of the BAT allocation, that substantially increased RB-S and SPC-SKF boat numbers, resulted in considerably higher excesses of boat hours (c.f., Metrics P2 and P4). However, the remainder of the metrics stayed relatively unchanged. Therefore, the USCG changes did not deteriorate the BAT model's impact on the boat allocations.

Practically speaking, the USCG can easily reduce the fleet size by selling unneeded boats through auctions, or as gifts to other governmental agencies; older boats are simply turned into scrap. For the few boat types where an increase is recommended (e.g., SPC-AIR), the USCG does have funds to purchase additional boats, but must follow specific procedures in the Department of Homeland Security's Acquisition Manual and the USCG's Major System Acquisition Manual, which implement the U.S. Government's Federal Acquisition Regulations.

Finally, the USCG regards the BAT model as a very important source of information because it (1) formally models a set of official USCG Business Rules that dictate feasible boat allocations, (2) establishes the assignment of variable hours to boats, especially the ones in high demand and low supply, (3) provides the optimal number of boats of each type at each station, and (4) optimally shares the limited supply of MLB boats. Indeed, comments from station officers have indicated that the partial allocation implemented in 2010–2011 has already resulted in the ability to "achieve the same results using less resources," as well as an observed reduction in maintenance costs. Furthermore, despite the modification of the BAT allocation, the USCG stated "the spirit of the model is being implemented." Consequently, the practical impact of this project closely mirrors the impact of the BAT allocation solution.

Table 4. Effect of varying γ.

	$\gamma^* = 28$	$\gamma = 40$	$\gamma = 50$	$\gamma = 75$
$\sum_{s\in S}\left\|\sum_{t\in T} h_{st} - H_s\right\|$	306	306	308	308
$\sum_{s\in S}\sum_{t\in T} y_{st}$	409	408	409	408
$\sum_{s\in S}\sum_{t\in T}(f_t x_{st} + v_t h_{st})$ (\$)	43,379,851	43,360,108	43,330,395	43,315,781

4. Computational Experiments

In this section, we discuss the results of a series of computational experiments to better understand the BAT model and its BAT-VAR generalization.

4.1. Parametric Analysis of the Effect of γ

We begin by studying the effect of γ, the threshold in miles for allowing MLB boats to be shared between two stations, on the BAT model's recommendations. As previously noted, the USCG wanted to know the smallest value of γ that would still allow a feasible allocation, since a small value of γ allows for high responsiveness of the shared boats. Recall that $\gamma^* = 28$ miles, the minimum value for which there exists a feasible boat allocation. We now vary γ to study its effect on the BAT allocation. We consider $\gamma \in \{40, 50, 75\}$ in addition to γ^*; higher values are not realistic as the response time of a shared MLB boat would be unacceptably high. We utilize the USCG prescribed weights of $w_1 = 0.95$, $w_2 = w_3 = 0.025$ and report, in Table 4, on the values of the three terms of objective function (1): (1) deviation of the overall supply of hours from overall demand of hours, (2) the number of types of boats at each station, and (3) total fleet operating cost.

Despite augmenting the feasible region of the BAT model, increasing values of γ do not significantly affect any of the three objective values. Not surprisingly, the boat allocations are essentially the same regardless of the value of γ. Intuitively, MLB sharing achieves model feasibility, by making sure an MLB is accessible when needed but does not substantially contribute to optimizing the objective function, which is primarily driven by the hourly assignments h_{st}. Therefore, increasing γ has little quantifiable value, and thus it is well justified that the USCG is interested in finding the minimum value of γ (for which a feasible solution exists) in order to increase responsiveness.

Next, we study how the optimized parameter γ^* behaves as we change the available number B_0 of the MLB boats. Recall that the motivation for sharing MLB boats, and the definition of γ, stemmed from an intrinsic shortage of MLB boats. In Table 5 we present values of γ^*, determined using a simple binary search, for increasing values of B_0; the current BAT model values are in bold.

Table 5. Effect of available number of MLB boats B_0 on γ^*.

MLB boats B_0	96	101	**106**	111	116	121	126
Sharing threshold γ^*	Infeasible	31	**28**	27	21	14	0

Note that in the vicinity of the currently available number of MLB boats $B_0 = 106$, i.e., $101 \leqslant B_0 \leqslant 111$, the value of γ^* is rather insensitive to changes in B_0. However, as B_0 is increased beyond 116, the threshold γ^* drops very quickly. While small changes (i.e., ± 5 boats) in the available number of MLB boats do not substantially change γ^*, and consequently the BAT model results, larger changes in B_0 require a new approach to the boat allocation. Indeed, an increase of 20 MLB boats eliminates the need for sharing, while a decrease of 10 boats renders sharing infeasible. Overall, the USCG found the results presented in Table 5 to be quite useful for identifying the range of MLB boat numbers that allow feasible boat sharing.

USCG policy, which dictates maximum tolerable response times, should be discussed with these results in mind. A target response time for MLB boats implicitly determines a target value of γ^*, and our results indicate the number of MLB boats needed. For example, a 25% reduction in γ^* is achieved by adding 10 more MLB boats to the existing 106. Note that, for 5 more boats, or a total of 15 additional MLB boats, the response time is reduced 50%. The nonlinear relationship between the available number of MLB boats B_0 and threshold γ^*, namely increasing marginal returns, should be included in any discussion of USCG policy on MLB response times.

4.2. Parametric Analysis of Objective Function Components

Up to this point, we have presented the official improvements with the weights set at $w_1 = 0.95$ and $w_2 = w_3 = 0.025$. In this section we quantify the trade-offs in the different objective function components as the weights change. In particular, we compare the components pairwise. This analysis was partially motivated by a USCG request to demonstrate the flexibility of the BAT model. In addition, this analysis was utilized to identify the weights ($w_1 = 0.95$ and $w_2 = w_3 = 0.025$) used in previous sections.

In Figure 1, we eliminate the hours $\sum_{s\in S}|\sum_{t\in T} h_{st} - H_s|$ objective by setting $w_1 = 0$ and quantify the trade-off between the boat-type $\sum_{s\in S}\sum_{t\in T} y_{st}$ and cost $\sum_{s\in S}\sum_{t\in T}(f_t x_{st} + v_t h_{st})$ objectives. In particular, we vary

Figure 1. Trade-off between the boat-types objective $\sum_{s\in S}\sum_{t\in T} y_{st}$ and the cost objective $\sum_{s\in S}\sum_{t\in T}(f_t x_{st} + v_t h_{st})$.

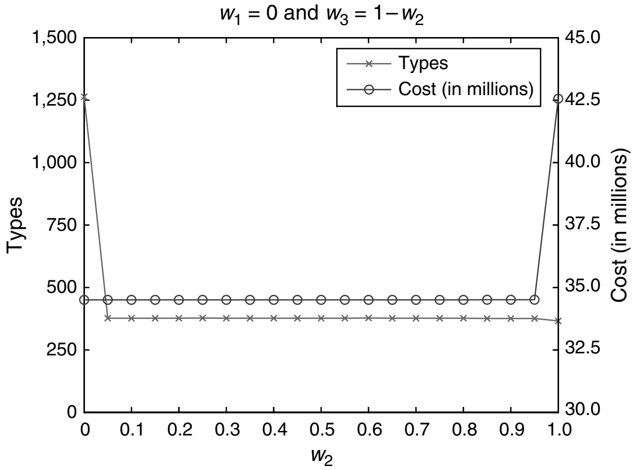

w_2 and w_3, subject to $w_2 + w_3 = 1$, in increments of 0.05. We see that there is a rather binary relationship between the boat-type and cost objectives. In particular, for all combinations of weights where $w_2 > 0$ and $w_3 > 0$, these objectives are approximately 377 total types and 34.5 million dollars, respectively. If $w_2 = 0$, then the boat-type objective increases to 1,264 and the cost remains practically unchanged. If $w_2 = 1$, the boat-type objective decreases to 366 while the cost objective increases to 42.5 million dollars.

In Figure 2, we quantify the trade-off between the hours $\sum_{s\in S}|\sum_{t\in T} h_{st} - H_s|$ and cost $\sum_{s\in S}\sum_{t\in T}(f_t x_{st} + v_t h_{st})$ objectives, by varying w_1 and w_3. Here we see a more uniform trade-off between the two objectives, which can be summarized quite practically: You have to pay to better match supply and demand of hours. However, there is still a somewhat binary relationship; there is essentially no change in the two objectives for $0.15 \leqslant w_1 \leqslant 0.95$.

Figure 2. Trade-off between the hours objective $\sum_{s\in S}|\sum_{t\in T} h_{st} - H_s|$ and the cost objective $\sum_{s\in S}\sum_{t\in T}(f_t x_{st} + v_t h_{st})$.

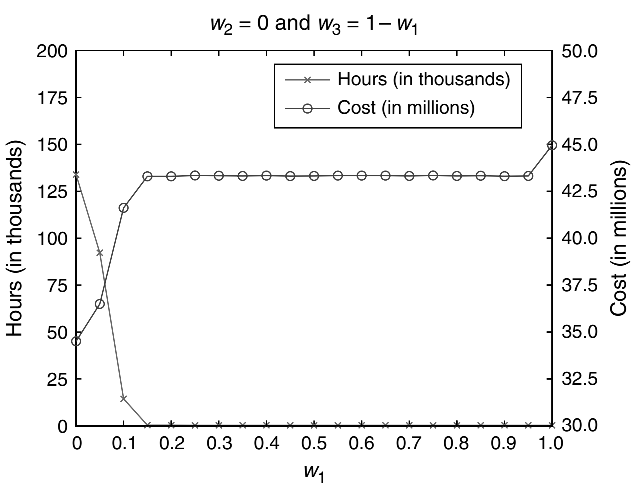

Figure 3. Trade-off between the hours objective $\sum_{s\in S}|\sum_{t\in T} h_{st} - H_s|$ and the boat-type objective $\sum_{s\in S}\sum_{t\in T} y_{st}$.

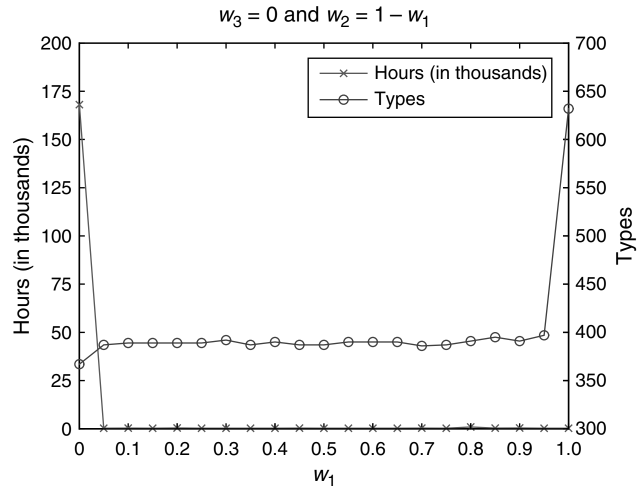

In Figure 3, we quantify the trade-off between the hours $\sum_{s\in S}|\sum_{t\in T} h_{st} - H_s|$ and boat-type $\sum_{s\in S}\sum_{t\in T} y_{st}$ objectives. Similar to the trade-off between boat-types and cost, there is a binary relationship. For all combinations of weights where $w_1 > 0$ and $w_2 > 0$, these objectives are approximately 365 hours and 390 boat types. If $w_1 = 0$, then the values are approximately 168,063 and 367; if $w_1 = 1$, the values are 308 and 632.

These results show that the 20–80 rule is in full effect. By using (at most) 20% of the total possible objective function weight, one gets at least 80% of the possible improvement. Therefore, from a practical point of view, the precise weights input into the model do not matter too much, as long as all weights are positive and greater than some minimum value (e.g., 0.025).

Finally, we mention that boat allocations and objective function values do not change when passing from the USCG preference weights ($w_1 = 0.95$ and $w_2 = w_3 = 0.025$) to a nondominated set of weights ($w_1 = w_2 = w_3 = 1/3$). This is consistent with the parametric studies presented above, which show that as long as extreme weights are not used (i.e., 0 or 1), the boat allocation and objective function values are stable. Indeed, further tests show that as long as all weights are at least 0.025, the boat allocation does not change.

4.3. Parametric Analysis of the BAT-VAR Model

In this section we provide the results of computational experiments that compare the BAT model with the BAR-VAR model. We define $S_{\text{uncertain}} = S$ to study the effect our risk management model has on the BAT allocation. We let μ_s be the original point estimate of demand, provided by the USCG, and model the standard deviation σ_s as a percentage of the original point estimate μ_s. It is convenient to present the following results in terms of a station-independent coefficient of variation $CV = \sigma_s/\mu_s$, $\forall s \in S_{\text{uncertain}}$; we consider $CV \in \{0.10, 0.25\}$. The acceptable shortage threshold θ_s is defined as a station-independent percentage p of the point estimate μ_s: $\theta_s = p\mu_s$ for all $s \in S_{\text{uncertain}}$; we consider $p \in \{0.10, 0.25\}$. Finally, we utilize station-independent confidence $\epsilon \in \{0.01, 0.05\}$ and employ the USCG prescribed weights of $w_1 = 0.95$, $w_2 = w_3 = 0.025$. We report changes in the BAT Allocation, along with changes in performance metrics, for all combinations of CV, p, and ϵ.

The first lesson we learned from our computational experiments is that using the range of (CV, p, ϵ) parameters just defined, the BAT-VAR model was not feasible. Intuitively, more boats are needed to achieve the different levels of hedging against underestimated demand forecasts. Therefore, to learn which types of boats are needed, the available number of all boat types B_i, $i = 0, \ldots, 10$, are doubled (and in two cases, quadrupled). In Table 6 we present the required increases in boat numbers, with respect to the BAT allocation, for all combinations of values of the (CV, p, ϵ) parameters.

General trends are apparent in Table 6. As the user-defined risk tolerances are tightened (i.e., decreasing p and ϵ), the total fleet sizes generally increase. The BAT-VAR Model generally first adds cheaper boats (e.g., SPC-SW, SPC-AIR, SPC-ICE, SPC-SKF), the one exception being the relatively more expensive RB-S boat, because it is the cheapest boat able to handle multiple station missions. However, the largest relative increases in fleet size are in moving from $CV = 0.10$ to $CV = 0.25$. For example, holding $CV = 0.10$ and modifying risk tolerances (p, ϵ) result in fleet size increases, as compared with the BAT allocation, ranging from 47 to 126 boats. In contrast, holding $CV = 0.25$ constant results in fleet size increases ranging from 114 to 768 boats. Therefore, the amount of demand uncertainty, as measured through the CV parameter, is the main driver of any robust allocation, and risk tolerances have second-order effects.

Examining the role of risk preferences on the fleet changes, we found that the confidence level ϵ has a stronger effect than that of the percentage p. When $\epsilon = 0.05$, the cheaper boat types (listed above) must have additional boats; when ϵ is reduced to the stricter 0.01, almost all boat types require additional boats. The p parameter has a similar but weaker effect on the resulting allocation.

In addition to the described trends, we want to point out that not all VAR-related parameters require additional boats. Extending our study to test the combination $CV = 0.10$, $p = 0.50$, and $\epsilon = 0.05$, the BAT-VAR model gave the exact same allocation as the BAT allocation, thus requiring no additional boats; consequently, the original BAT allocation was robust to small amounts of demand uncertainty and relaxed risk preferences. To expand on this observation, we investigate stations' shortage violations, where a violation is defined as a station's shortage $H_s - \sum_{t\in T} h_{st}$ being greater than the station's acceptable shortage threshold θ_s (c.f., Constraint (14)). More precisely, we calculate the probability $\mathbf{P}(H_s - \sum_{t\in T} h_{st} > \theta_s)$ and the expected violation $\mathbf{E}[\max\{H_s - \sum_{t\in T} h_{st} - \theta_s, 0\}]$ for all stations s, when demand is normally distributed with a mean of μ_s, and standard deviation of σ_s, for the BAT allocation. When the variability is low and the shortage threshold is high ($CV = 0.10, p = 0.25$), the probability of violation ranges, over all stations, from 0.00% to 0.62% and the expected violation ranges from 0.00 to 2.33 hours; therefore, the BAT allocation is indeed robust to low levels of variability and relaxed risk preferences. However, at the other extreme of high variability and low shortage threshold ($CV = 0.25, p = 0.10$), the probability ranges from 13.08% to 34.46%, and the expected violation ranges from 4.25 to 670.35 hours. Therefore, the chance of a shortage violation and the expected violation vary considerably, depending on the underlying demand variability and acceptable shortage threshold, which further motivates the need for the BAT-VAR model.

We also perform two rankings of the stations in the BAT allocation, according to the probability of a shortage

Table 6. Required increases in boat numbers to achieve different levels of risk management.

Boat type	$\epsilon = 0.05$ $CV = 0.10$ $p = 0.25$	$\epsilon = 0.05$ $CV = 0.10$ $p = 0.10$	$\epsilon = 0.05$ $CV = 0.25$ $p = 0.25$	$\epsilon = 0.05$ $CV = 0.25$ $p = 0.10$	$\epsilon = 0.01$ $CV = 0.10$ $p = 0.25$	$\epsilon = 0.01$ $CV = 0.10$ $p = 0.10$	$\epsilon = 0.01$ $CV = 0.25$ $p = 0.25$	$\epsilon = 0.01$ $CV = 0.25$ $p = 0.10$
MLB	—	—	—	—	—	—	4	21
SPC-NLB	2	2	2	2	2	2	6	6
SPC-HWX	—	—	—	2	—	1	6	7
RB-M	—	—	—	51	—	12	62	104
RB-S	—	40	41	33	12	47	362	364
RBS-AUX	—	—	—	—	—	—	—	—
SPC-LE	—	—	3	6	3	4	42	45
SPC-SW	5	6	29	26	13	33	66	75
SPC-AIR	3	3	5	5	1	4	17	21
SPC-ICE	21	15	18	19	16	14	48	54
SPC-SKF	16	15	16	21	7	9	72	71
Totals	47	81	114	165	54	126	685	768

violation and the expected violation, respectively. These computational results show that the stations with the "heavy weather" mission assignments have the highest probability of violating the acceptable shortage limit. This observation can be explained by the fact that the heavy-weather stations require the MLB boat, a limited resource; therefore, if demand is variable, it becomes likely that the limited MLB boats will not be sufficient. In contrast, the stations that had the highest expected violations were the stations without any specific mission. Because these stations are not assigned any hours to satisfy the critical demands in Constraint (10), they are also more susceptible to demand variations. Combining these two observations, it is the stations that require limited resources and those without a specific mission that are most vulnerable to demand variability.

Next, we present in Table 7 the Performance Metrics P1–P10, as defined in §3.3, for the different risk management parameters described above. We see that reducing the risk of a boat shortage (c.f., zero values for Metrics P3, P5, and P10) requires an increase in the hours assigned to stations, as seen in increased values of Metrics P2 an P4, driven primarily by the coefficient of variation. For example, with $\epsilon = 0.05$ and $p = 0.25$, and CV increasing from 0.10 to 0.25, the average excess hours per station increases 5.2 times. Fixing CV and p, and decreasing ϵ from 0.05 to 0.01 also resulted in a multiplicative increase in average excess hours, ranging from 2.4 to 4.6 (e.g., column 7's excess hours (Metric P4) divided by column 3's give 4.6).

To obtain more practical insights, notice that for a fixed value of ϵ, increasing CV or lowering p raises the proportion of stations with more than two boat types, as seen in Metric P6, by 62%–155% over the BAT model. This implies that increased uncertainty and stricter shortage preferences result in higher maintenance costs, because more stations must operate three or more boat types and corresponding maintenance crews must be present. However, the average number of boat types per station, as measured in Metric P7, increases by smaller percentages, namely 13%–48%, over the BAT model. Therefore, we can conclude that the BAT-VAR model increases the variability of the number of boat types at each station, which will complicate the assignment of maintenance crews to stations.

For a fixed ϵ, the fleet operating cost (Metric P8) and fleet utilization (Metric P9) uniformly increase and decrease, respectively, as uncertainty increases or risk tolerances tighten. Very simply, a BAT-VAR model user must pay, both in dollars and utilization, for higher levels of demand uncertainty, higher levels of confidence, and lower acceptable thresholds.

Finally, note that the value-at-risk constraints themselves do not change the value of γ^*. As seen in §4.1, the MLB sharing is based on the number of available MLB boats, while the VAR constraints are based on the allocation of hours to stations. However, any necessary increase in the available number of MLB boats, to make the BAT-VAR model feasible, will reduce the value of γ^*.

These computational studies allow us to make concrete policy recommendations for the USCG to enable risk management through the BAT-VAR Model. Currently, the Office of Boat Forces (OBF) only maintains the average demand H_s of hours for all stations s, which is obtained from the stations. We recommend that the platform division (PD) of the OBF, the BAT model user, calculate the standard deviation of demand at each station for the same time period that the H_s parameters are defined over (i.e., five years), because our studies show that any allocation that incorporates risk management is highly dependent on the level of uncertainty present. If stations do not maintain demand information for the five-year period that underlies the BAT model, the OBF also needs to request that those stations start recording the demand data.

In addition, we recommend introducing new probabilistic performance measures for the BAT-VAR model, e.g., the probability of a shortage or excess of capacity hours at a given station, and the probability of a demand shortfall. These metrics would better measure the model's ability to handle uncertainty than the existing metrics.

We also suggest employing the BAT-VAR model to identify the additional number and types of boats necessary to achieve a desired level of risk management. A cost increase from these additional boats can be, at least partially, offset by

Table 7. Performance metrics for BAT-VAR model.

	BAT model	$\epsilon = 0.05$ $CV = 0.10$ $p = 0.25$	$\epsilon = 0.05$ $CV = 0.10$ $p = 0.10$	$\epsilon = 0.05$ $CV = 0.25$ $p = 0.25$	$\epsilon = 0.05$ $CV = 0.25$ $p = 0.10$	$\epsilon = 0.01$ $CV = 0.10$ $p = 0.25$	$\epsilon = 0.01$ $CV = 0.10$ $p = 0.10$	$\epsilon = 0.01$ $CV = 0.25$ $p = 0.25$	$\epsilon = 0.01$ $CV = 0.25$ $p = 0.10$
P1	622	805	852	902	965	842	922	1,485	1,568
P2 (%)	1.7	92.7	95.5	100.0	100.0	100.0	100.0	100.0	100.0
P3 (%)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
P4	102.0	335.9	672.2	1,753.3	2,084.2	1,544.4	1,875.2	4,737.5	5,068.3
P5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
P6 (%)	30.9	63.5	57.9	57.9	60.7	50.0	53.4	78.7	69.1
P7	2.3	3.0	2.7	2.7	2.8	2.6	2.7	3.4	3.3
P8 ($)	43,379,851	44,741,064	47,784,036	67,046,054	76,287,193	62,461,465	70,342,142	106,666,406	115,720,450
P9 (%)	99.0	91.2	84.0	60.1	53.0	64.6	57.5	46.8	43.3
P10 (%)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

the savings generated by the implemented allocation (which is more than \$2 million; see Table 3, Performance Metric 10, implemented versus original allocations).

These recommendations have been submitted to the OBF for further consideration. The OBF plans to evaluate the BAT-VAR model for possible implementation after the completion of the boat reallocation using the BAT model.

5. Conclusion

In this paper we have described the collaboration between the USCG and the authors. The result is an integer linear programming model, denoted the BAT model, that is currently being applied by the USCG for the allocation of their boat fleet nationwide. The model is a resource allocation one, with additional capabilities of optimally sharing a limited supply of boats, making decisions at the hourly level, and utilizing a high resolution description of supply and demand. Risk-management capabilities, while not yet implemented by the USCG, provide additional modeling and decision making flexibility. The risk-management additions to the BAT model are linear inequality constraints, and the model remains a linear one. The USCG adoption of the BAT model and internal resistance are discussed in detail. We also report on the USCG process of modifying the BAT model recommendations to obtain an allocation that is implemented in the field. We present quantitative performance metrics that show the significant practical impact of the BAT model. This substantial impact is based on a decrease of 87.6% of shortage hours as well as a reduction of 62.3% of the excess hours at the stations, based on the implemented vs. original allocations. At the same time, the size of the boat fleet is reduced by 88 boats, which is 10.9% of the original fleet size, with a simultaneous growth in boat utilization. The fleet operating cost is also lowered by more than \$2 million, i.e., 4.6% of the original operating cost. In addition, we conduct a variety of computational experiments, whose results support a number of recommendations presented in this paper.

A direction of possible future research is to consider dynamic decision making. The model described in this paper, even with the risk-management capabilities, is static. From the USCG perspective, boat allocation decisions are made once a year, or even less frequently. Therefore, incorporating real-time data about realized station demands can result in an even better match of supply and demand.

Acknowledgments

The authors gratefully acknowledge CDR Todd Wiemers, Todd Aikins, Thomas Owens, and CDR Tamara Koemer for their collaboration and for making this project possible. The first author gratefully acknowledges research funds from the Graduate Business Programs at Saint Mary's College of California. Finally, the authors thank the area editor, associate editor, and two referees for their insightful comments, which greatly improved the quality of the paper.

References

Allen T (2009a) Modernization: Building the 21st century Coast Guard. *Sea Tech.* 50(1).

Allen T (2009b) Voyage of renewal. *Surveyor* Fall:2–5.

Billings R (2005) Scheduling ammunition loading and unloading for U.S. Navy ships in San Diego. Technical report, Naval Postgraduate School, Monterey, CA.

Birge J, Louveaux F (2011) *Introduction to Stochastic Programming*, 2nd ed. (Springer, New York).

Brown G, Dell R, Farmer R (1996) Scheduling Coast Guard district cutters. *Interfaces* 26(2):59–72.

Department of Homeland Security (2010) Featured components—Coast Guard. Accessed May 11, 2010, http://www.dhs.gov/index.shtm.

Deshpande V, Iyer A, Cho R (2006) Efficient supply chain management at the U.S. Coast Guard using part-age dependent supply replenishment polices. *Oper. Res.* 54(6):1028–1040.

Everingham K, Polaski G, Reidlin F, Shirk M, Deshpande V, Iyer A (2008) Operations research enhances supply chain management at the U.S. Coast Guard aircraft repair and supply center. *Interfaces* 38(1):61–75.

Gol'stejn E, Dempe S (2002) A minimax resource allocation problem with variable resources. *Eur. J. Oper. Res.* 136:46–56.

Hull J (2000) *Options, Futures, and Other Derivatives*, 4th ed. (Prentice Hall, Englewood Cliffs, NJ).

Kauffman R, Sougstad R (2007) Value-at-risk in IT services contracts. *Proc. 40th Annual Hawaii Internat. Conf. System Sci.* (IEEE Computer Society, Washington, DC).

Pearson N (2002) *Risk Budgeting: Portfolio Problem Solving with Value-at-Risk*, 1st ed. (Wiley, New York).

Radovilsky Z, Koermer T (2007) Allocation of U.S. Coast Guard boats utilizing integer programming. *J. Acad. Bus. Econom.* 7(2):130–135.

Ravindran R, Bilsel U, Wadhwa V, Yang T (2010) Risk adjusted multi-criteria supplier selection models with applications. *Internat. J. Production Res.* 48(2):405–424.

Sarykalin S, Serraino G, Uryasev S (2008) Value-at-risk vs. conditional value-at-risk in risk management and optimization. *TutORrials in Operations Research* (INFORMS, Hanover, MD), 270–294.

Wu D, Olson D (2010) Enterprise risk management: A DEA VAR approach in vendor selection. *Internat. J. Production Res.* 48(16):4919–4932.

Zarybnisky E (2003) Allocation of air resources against an intelligent adversary. Technical report, Air Force Institute of Technology, Wright-Patterson AFB, OH.

Zhang B, Murali P, Dessouky M, Belson D (2009) A mixed integer programming approach for allocating operating room capacity. *J. Oper. Res. Soc.* 60(5):663–673.

Michael R. Wagner is an assistant professor of operations management at the Michael G. Foster School of Business at the University of Washington. His research interests are in decision making under uncertainty (stochastic, online, and robust optimization, as well as hybrids thereof). In particular, he investigates the value of information in a variety of application domains, such as supply chain management, inventory control, resource allocation, and logistics.

Zinovy Radovilsky is a professor of management in the Department of Management at California State University, East Bay. His research interests include supply chain optimization specifically in supply chain planning and scheduling, forecasting, e-commerce business modeling, and enterprise resource planning and management.