

US Department of Defense
U.S. Air Force Research

University of Nebraska - Lincoln

Year 2005

Decision support for the career field
selection process at the US Air Force
Academy

Andrew P. Armacost
US Air Force Academy

James K. Lowe
US Air Force Academy



Decision Aiding

Decision support for the career field selection process at the US Air Force Academy

Andrew P. Armacost^a, James K. Lowe^{b,*}

^a *Department of Management, 2354 Fairchild Drive, Suite 6H128, US Air Force Academy, CO 80840, USA*

^b *Department of Management, 2354 Fairchild Drive, Suite 6H242, US Air Force Academy, CO 80840, USA*

Received 21 June 2002; accepted 22 November 2002

Available online 14 November 2003

Abstract

Each year, the US Air Force Academy graduates nearly 1000 young men and women. To support the decision of which cadets will be classified into which career fields, we describe a linear programming formulation with appealing computational properties that enable it as the core of a decision support tool. We explore methods for measuring and balancing cadets' class standing, Air Force career field requirements, and cadets' career field preferences in the context of this model. Our computational experiments demonstrate the improvement of this method over previous classification approaches, yielding more than 10% increase in the number of cadets assigned to their top career field choice and yielding nearly a 100% reduction in the number of cadets not receiving any of their career field choices. We also explore alternative methods for measuring cadets' career field preferences and demonstrate the positive effect of the new measurement scheme on the overall classification. Because of the short running time of this model, it will serve as a flexible, real-time component of the Academy's classification process.

© 2003 Elsevier B.V. All rights reserved.

Keywords: Manpower planning; Human resources; Military; Linear programming; Network flows

1. Introduction

Each year, the United States Air Force Academy in Colorado Springs, CO graduates and commissions roughly one thousand young men and women as Air Force officers. Following their commissioning, each member of the graduating class begins his or her service in one of nearly three

dozen career fields. In past years, the process by which each was assigned to a career field was manually intensive and often resulted in a significant percentage of graduates assigned to fields in which they had little prior interest.

In this paper, we describe an optimization-based methodology for assisting the classification process, which is the process of matching graduating officers to career fields. The model balances the personnel needs of the Air Force with individuals' career field preferences and their performance during four years at the Academy. The assignments suggested by the model provide an

* Corresponding author.

E-mail addresses: andy.armacost@usafa.af.mil (A.P. Armacost), jim.lowe@usafa.af.mil (J.K. Lowe).

initial starting point for the Academy's personnel office and the Air Force Personnel Center (AFPC) in San Antonio, TX to make the final Air Force Specialty Code (AFSC) assignments. The speed of the model and its ability to identify the marginal value of alternate career field allocations make it an attractive real-time tool to support personnel planners and analysts.

We make the following significant contributions:

- model and implement this decision problem as a flexible and easily-solved network flow model;
- demonstrate the performance of the model using data from the Academy's graduating classes of 2001 and 2002; and
- demonstrate the benefits of collecting more detailed information regarding individual job preferences.

This paper is organized as follows. We describe the career field classification process and create a network flow model to aid in this process. We then illustrate several methods for establishing trade-offs between cadet preferences and their class rank. To better match cadets with their top job choices, we propose a change to the manner in which cadet job preferences are measured. Finally, we examine these methods when applied the graduating classes

of 2001 and 2002. Our work with the Class of 2001 provided a baseline for this multi-year project, and our work with the Class of 2002 served as a prototype and proof-of-concept of a complete classification system. This work led to a fully automated system that was adopted and implemented by the Air Force for initial use with the Class of 2003.

2. The Air Force career field classification process

Each fall, the Academy initiates the process of collecting cadet AFSC preferences, ranking the cadets relative to their peers, and assigning cadets to career fields based on rankings, preferences, and quotas (summarized in Fig. 1). We provide detail on the major elements of this process.

2.1. Ranking cadets

Each fall, a central board of senior Academy leaders convenes to provide an overall assessment of each cadet's performance during his or her four years at the Academy. This assessment takes into consideration a cadet's academic performance, athletic performance, and leadership involvement. Their assessment is a combination of objective and subjective measures that results in a score between 30 and 50 points for each cadet.




Fig. 1. The Academy's process of ranking its graduates and assigning them to initial career fields.

This “AFSC board” consists of ten members divided equally into two separate panels. The composition of each panel ensures representation from each of the Academy’s mission areas: academics, athletics, and leadership training. Each panel evaluates between 450 and 500 cadets, depending on the actual graduating class size. For a given panel, each reviewer rates only the members from that panel’s batch of cadets. Each panel member assigns each cadet a score between 6 and 10, and the overall rating is a sum of the five panel members’ individual scores. The results from each panel are standardized and combined into a single rank-ordered list of all cadets in the graduating class.

2.2. Classifying cadets

Cadets will enter a career field in one of two general categories: *rated* and *non-rated*. Rated cadets are those who will enter a flying career field (i.e., pilot or navigator) and non-rated cadets are those entering a non-flying career field. Each year, AFPC officials determine the number of rated slots available to the Academy. Pilot-qualified cadets compete for those slots according to the rank they receive from the AFSC board. In recent years, nearly every physically qualified graduate received a pilot slot. This report focuses upon the process of classifying non-rated cadets into career fields.

All cadets are asked in the fall of the senior year to pick their top six non-rated AFSCs and to specify their preference on a 1–6 scale, with 1 indicating their top choice. In past years, this classification followed a greedy matching process. Beginning with the top ranked cadet (based upon AFSC Board results), each cadet selected their AFSC from the set of available AFSCs. In general, this method assured the top of the rank ordering of getting their top AFSC choice. The lower ranked cadets were less likely to receive their top choices and were increasingly likely to face a set of AFSCs that were not among their original six choices. Because of the limitations of this greedy classification process, we proposed a new optimization-based process for classifying cadets to AFSCs. This process meets Air Force needs, rewards the Academy’s top performers, and recognizes the preferences of all cadets.

3. Optimization-based classification model

Given the described process, we offer the following model to assist this classification decision. Let set I represent all non-rated cadets considered in the selection process and set J denote all possible jobs (e.g., career fields). The value of assigning cadet $i \in I$ to job $j \in J$ is denoted by c_{ij} . We defer the important discussion of how we arrive at this value to later sections. AFPC provides target quotas for each AFSC. More generally, we allow AFPC to specify lower and upper limits for each career field, denoted by l_j and u_j .

Our primary decision is to determine the *best* assignment of cadets to jobs (the definition of *best* remains). We model this with the binary decision variable, y_{ij} , which equals 1 if cadet i is assigned to job j and zero otherwise. Through the use of slack variables, we allow deviation from the upper and lower limits for each job. The variable s_j^+ is the amount by which the number of cadets assigned to job j exceeds the upper limit, u_j . Similarly, s_j^- is the amount by which we under-fill job j . The cost of violating the upper and lower limits are given by d_j^+ and d_j^- , respectively.

With the goal of maximizing the value of assignments made for the cadets, offset by the cost of missing the Air Force’s objectives for each career field, we introduce the following optimization model:

$$\max \sum_{i \in I} \sum_{j \in J} c_{ij} y_{ij} - \sum_{j \in J} (d_j^+ s_j^+ + d_j^- s_j^-)$$

$$\text{subject to: } \sum_{j \in J} y_{ij} = 1 \quad i \in I, \quad (1)$$

$$\sum_{i \in I} y_{ij} - s_j^+ \leq u_j \quad j \in J, \quad (2)$$

$$\sum_{i \in I} y_{ij} + s_j^- \geq l_j \quad j \in J, \quad (3)$$

$$y_{ij} \in \{0, 1\} \quad i \in I, \quad j \in J, \quad (4)$$

$$s_j^+, s_j^- \geq 0 \quad \text{and integer } j \in J. \quad (5)$$

Constraint set (1) forces the assignment of each cadet to exactly one AFSC. Constraint set (2) limits the number of cadets assigned to an AFSC to the upper limit plus any deviation from that limit. Constraint set (3) forces the number of cadets assigned to an AFSC to the lower limit minus

any deviation from that limit. Constraint sets (4) and (5) enforce the assignment variables to be binary and the AFSC deviation to be integral.

While the model allows deviations both above the upper AFSC limit and below the lower AFSC limit, AFPC typically specifies a single target quota for each AFSC and expects no deviation. As such, we set the upper and lower limits to the target quota and increase the costs of deviating (d_j^+ and d_j^-) from the target to be large so that the model will satisfy the target quotas exactly.

To solve this 0–1 integer program, we only need to solve its linear programming relaxation. This model can be mapped directly to a network flow problem, which has the property that any solution to the linear programming relaxation will be integral and, therefore, optimal (see [1]). Because of the mapping, the formulation also has the integrality property and we need not consider methods like branch and bound to find the optimal integer solution. Moreover, because this is a network flow problem, specialized algorithms (i.e., the network simplex method) can be used to solve this model within seconds. As such, the model's fast run times (i.e., seconds) make it ideal for supporting the classification process in real-time; reducing classification iterations between the Academy and AFPC from weeks to potentially hours.

Models of similar form have been used for military personnel allocation, including decisions to assign US Navy personnel to vacant jobs with the goal of minimizing moving cost [9]. The details of the implementation of the Navy's assignment system, including issues such as user acceptance, are described in Blanco and Hillery [2] and Whisman et al. [12]. Krass et al [8] describe a non-linear network optimization problem to assist the Navy in determining assignments of personnel to combat duty. Reeves and Reid [10] introduces an interactive, preference-driven manpower planning model for the US Army Reserves. This model balances five objectives for the assignments within a 100 person unit over the course of a one-year planning horizon. Weigel and Wilcox [11] describe a hierarchy of models that support the Army's enlisted personnel system. The models are used to support a variety of decisions, from policy-level analysis to detailed, unit-level planning issues. Fi-

nally, optimization methods have been applied to decisions of manpower scheduling, including the assignment of service department workers with specialized skills to jobs (see [4]) and the determination of a maintenance workforce specialization structure and the assignment of specific tasks to workers (see Dietz and Rosenshine [6]).

The problem we consider differs from these previous works in two key respects. First, we are not considering one-to-one assignments of persons to jobs; rather, we are assigning persons to career fields, which have a required target for the number of people. Second, our objective function does not deal with explicit measures such as travel costs. Ultimately, we are concerned with giving the decision-maker, not the model, the final say regarding the delicate trade-off between cadet preferences and their standing in the class.

3.1. The objective function

The output of the AFSC board is an aggregate subjective score and subsequent rank order of each cadet. The output of the cadet's career field selection is a ordering, from 1 to 6, of their job preferences. These two outputs combine to form the objective function coefficients (c_{ij}) of classifying the j th job preference to the i th ranked cadet. In this section, we discuss alternatives for combining the inputs (i.e., cadet rank/score and job preference) into a single value that captures the appropriate trade-off between the two inputs. For example, we might wish to make the top-ranked cadet's second choice ($c_{1,2}$) equal to the middle cadet's top choice ($c_{163,1}$). The exact trade-off depends on the decision-maker's preference. In fact, decision-makers are more concerned with the net effect of the objective function (i.e., the impact upon the overall classification of cadets) than the function itself. In this section, we first describe methods to represent the value associated with a cadet's class rank, $v(i)$, that will give us a flexible method for capturing this trade-off. We then show how this value is combined with each cadet's AFSC preferences, $\text{pref}(i, j)$, to yield the model's objective function coefficients:

$$c_{ij} = f[v(i), \text{pref}(i, j)]. \quad (6)$$

In Section 3.2 we describe alternatives for representing $v(i)$ and in Section 3.3 we describe methods for representing $\text{pref}(i, j)$. In Section 3.3, we will present methods to better measure job preferences and build them into the objective function.

3.2. Valuing a cadet's rank, $v(i)$

There are many ways to represent the value awarded when cadets receive their top AFSC choice. This value is one component of the overall function described in (6). In general, we assume that the value curve, $v(i)$, should be strictly decreasing with respect to class rank. The general shape of the curve will have a drastic effect upon how the optimization model makes trade-offs between the AFSC classification for persons with different class standing. In this section, we explore three such curves: (1) $v(i)$ displays a constant marginal reduction as class rank increases, (2) $v(i)$ displays a diminishing marginal reduction as class rank increases, and (3) $v(i)$ is determined by the actual AFSC board score.

3.2.1. Constant marginal difference

In general, we would like the value awarded for matching cadets with their top choice to be a linear, diminishing function of board rank. That is, the first cadet receives the greatest value. The difference between the first and second cadets will be the same as the second and third, which will be the same as the third and fourth, and so on.

Thus, the value received due to cadet rank is given by

$$v(i) = N - \text{rank}(i), \quad (7)$$

where N is the number of graduates. If we care only about the cadet's standing relative to their non-rated peers, N is the number of non-rated cadets and $\text{rank}(i)$ is cadet i 's rank within the non-rated group. Note that this is a linear function of rank, as shown in Fig. 2, where N is 327 (the number of non-rated cadets in 2001). Again, this curve represents the value awarded whenever a cadet receives his or her top choice AFSC.




Fig. 2. Value received for first choice as a linear function of board rank.

3.2.2. Decreasing marginal difference

Again, we would like the value awarded for matching a cadet with their top choice to be a decreasing function of class rank. In this section, marginal change in this value, $v(i)$, should increase rapidly as we move toward the top of the class. This function is consistent with the philosophy that “top” cadets have truly distinguished themselves in their four years at the Academy. The marginal difference decreases as we move lower in the class rank.

To capture this philosophy, we employ a learning curve function (for example, see [3]). Assume N is the total number of cadets, and α is a number between 0 and 1 that represents a rate of learning (or the slope of the learning curve). Let i be the cadet's rank. Then the value awarded for assigning each cadet is

$$v(i) = N \cdot i^{\log_2 \alpha}. \quad (8)$$

When α is, say 0.95, the learning curve reduces the value by 5% each time cadet rank is doubled. In our application, the value of assigning cadet #2's top choice returns a value that is 5% less than cadet #1's. Assigning a top choice to cadet #4 has a value that is 5% less than cadet #2's, and so forth. This function provides marginal changes consistent with the philosophy of rewarding top performers. Fig. 3 shows the 95% learning curve for $N = 327$.




Fig. 3. Value received for the first choice is a decreasing function of board rank, with decreasing marginal difference.

3.2.3. Board score approach

Our final curve once again assumes that the value function, $v(i)$, decreases with class rank. In this section we introduce a function that provides both increasing marginal rewards for top performers and increasing marginal penalties for low performers. In order to entice the model to give more top choices to the top ranked cadets, it must be increasingly difficult for the lowest ranked cadets to get their top choices. Thus, in the lower tier of class rank, the marginal penalty increases as we near the bottom. Interestingly, plotting AFSC board score (see Fig. 4) gives the desired distribution of values, $v(i)$. Moreover, this method uses direct output from the AFSC board as an input to the model.




Fig. 4. Value received based on board ranking.

3.3. Accounting for career field preferences, $pref(i, j)$

In the previous sections, we described three methods for representing a baseline value, $v(i)$, for assigning the top choice to cadets based upon their board rank/score. We now create the objective function, c_{ij} , by combining the rank value, $v(i)$, with the preferences, $pref(i, j)$, that each cadet assigns to each AFSC. The current method relies solely upon the 1–6 preference scale provided by cadets. We also consider a second method that captures the relative preference (or indifference) between their six career field choices.

3.3.1. Ordinal preferences

We denote cadet i 's preference for job j by $pref(i, j)$, which equals 1 for their top choice and 6 for their 6th choice. We let $pref(i, j) = 0$ when job j is not one of the cadet i 's top six choices. The objective function coefficient that combines both rank and preference is given by

$$c_{ij} = \begin{cases} \frac{v(i)}{pref(i, j)}, & \text{when } pref(i, j) > 0, \\ -v(i), & \text{otherwise.} \end{cases} \quad (9)$$

Fig. 5 displays the board score (see Section 3.2.3) family of functions. The top curve in the figure represents the value of assigning cadets their top choice. The second highest curve shows the value of assigning cadets' to their second AFSC choice. The third through sixth curves represent assigning cadets to their third through sixth AFSC. The bottom curve is the negative of the top choice and applies whenever a cadet is assigned to an AFSC not listed as one of their top six.

Note that the bottom ranked cadet receives the same value (i.e., 0) for all career fields. The model will be indifferent about which career field to assign this cadet. We can add an additional parameter to this method that helps control the extent to which the lower ranked cadets are sacrificed for the benefit of the others. The objective function coefficients determined in (9) become

$$c_{ij} = \begin{cases} \frac{v(i) + k}{pref(i, j)}, & \text{when } pref(i, j) > 0, \\ -(v(i) + k), & \text{otherwise.} \end{cases}$$




Fig. 5. Value received based on class rank and career field preference; receiving a non-preferred career field is penalized.

The effect of this change is that the curves in Fig. 5 become increasingly separated as we increase the constant, k . This increase causes the model to avoid assigning cadets to non-preferred jobs. Rather than selecting an a priori value for k , we vary this parameter to provide decision-makers with a set of possible classifications from which to choose.

3.3.2. Relative preferences

Our primary measure of the model's performance was the number of cadets receiving their top preference(s). Given this, we readdress how preferences are measured. When ranking jobs on a scale of 1–6, we have no information regarding the relative preference between job selections. A model that uses the objective function coefficients described in (9) assigns significantly different values to a cadet's first and second AFSC. However, it might be the case that the cadet is indifferent between the two AFSCs. In this situation, awarding the cadet's second choice (based upon the forced rank ordering) may free a slot that enables a lower ranked cadet to receive a higher preference.

Modifying our approach for creating the objective function coefficients is straightforward. We define $r(i, j)$ as the relative satisfaction cadet i has for job j . A cadet's top preference(s) will receive a relative satisfaction of 1, and all other preferences

have a value less than one. For each cadet $i \in I$ and job $j \in J$ the objective function coefficients are given by

$$c_{ij} = \begin{cases} r(i, j) \cdot v(i) & \text{if } r(i, j) > 0, \\ -v(i) & \text{if } r(i, j) = 0, \end{cases} \quad (10)$$

where $v(i)$ is determined by methods such as the three discussed earlier.

A number of mechanisms can be used to solicit relative preferences from the cadets. But the overall idea is the same: to gain some recognition of how each cadet feels about their six job choices. Capturing individual preferences is superior to any of the arbitrary preference schemes inherent in the 1–6 ranking process. The bottom line is the impact upon classifications: which cadets receive which jobs. Because individual preferences and AFSC quotas greatly affect the classification of each graduating class, we do not anticipate developing a "correct set" of value functions. Instead, we develop a process that enables decision-makers the ability to easily, and quickly, explore the effects of alternative functions.

4. Computational results

In this section, we offer empirical results that demonstrate the success of this optimization-based

classification model. We begin with computational results from the Class of 2001 and demonstrate the solution quality and its improvement over the traditional greedy approach described in Section 2. We then illustrate the effects of objective function changes upon the model's solution. We next turn to the Class of 2002, with whom we examine the effect of using individually specified preferences {see (10)} rather than the 1–6 ranking of a cadet's six AFSC choices {as prescribed in (9)}.

Because this decision support environment relies upon Excel spreadsheets and Access databases, we opted to model within a spreadsheet environment on a Pentium III workstation. We implemented the model using both the Large-Scale Solver from Frontline Systems [7] and XPRESS-MP version 13.0 [5]. The Frontline solver is the large-scale version of the Excel Solver. To use XPRESS-MP, we linked the solver to the spreadsheets using ODBC connections and SQL. The fact that this model is a network flow problem means the optimal linear solution is always integer and, hence, solve times are fast. We do not report solution times because the model solves in seconds, which is more than sufficient for the purposes of our model.

4.1. Results for Class of 2001

Of the graduating class of 2001, 327 graduates were classified into non-rated AFSCs. The AFPC initially specified 322 target quotas across 36 career fields, and allowed the Academy to classify the additional 5 cadets into their preferred AFSCs. We penalized deviation from each AFSC quota heavily. Thus, the model matched each AFSC quota exactly (aside from the difference between graduates and total target quotas).

4.1.1. Comparison with greedy classification

Our first goal was to compare the model's solution to that generated by the greedy classification approach. This greedy method begins by awarding the top-ranked cadet their first choice. We simulate the greedy method by decrementing the number of remaining slots in that career field and turn to the second ranked cadet. This cadet selects their highest choice among the remaining available

slots. The process continues until all cadets are classified. If, in our simulated heuristic, a cadet finds no slots among their six choices, we proceed to the next cadet and the process continues. At the end, we randomly assign the cadets who did not receive one of their top six choices to the remaining slots. In practice, a cadet who finds no available openings throughout their original six choices makes a selection from any of the available openings rather than risk a random classification. Therefore, as we compile the statistics on the number getting one of their six choices versus a non-preferred career field, our heuristic *over estimates* the number getting one of their AFSC preferences. Recall that our heuristic does not allow a cadet to "take away" a preferred AFSC of a lower ranked cadet if none of their preferences were available. Thus, the quality of the solution (in practice) can only be worse than the results reported.

Table 1 provides both the results obtained using the optimization versus the greedy heuristic. The "# assigned" columns indicate the number of cadets who received their first choice, second choice, and so forth (proceeding down the rows). Notice the difference between the numbers of first choices given using the optimization versus the existing greedy approach. Moreover, the number of "non-preferred" jobs assigned is reduced from 12 to either 0 or 3, depending upon the objective function used. Thus, cadet preferences are better satisfied using the optimization. Presumably, this will have a positive effect on morale and long-term commitment to the Air Force.

4.1.2. Viewing alternative solutions

Comparing the results from the objective function reported in Table 1, we see a difference that seems to indicate the learning curve objective function (Fig. 3) gives a better aggregate solution. The learning curve model reports more first preferences received and zero "non-preferred" choices received. However, this may be misleading as we examine the differences more closely.

Fig. 6 displays the AFSC preference received of all 327 cadets using the two objective functions. While the learning curve objective function yields a greater number of first choices, notice the large

Table 1
Results for the Class of 2001, comparing the model using two objective functions with the greedy approach to job classification

Job preference	Model results					
	Greedy Heuristic		Learning curve objective		Board score objective	
	# Assigned	Cum%	# Assigned	Cum%	# Assigned	Cum%
1	215	65.7	248	75.8	244	74.6
2	53	82.0	42	88.7	40	86.9
3	18	87.5	14	93.0	10	89.9
4	13	91.4	7	95.1	6	91.7
5	12	95.1	9	95.9	14	96.0
6	4	96.3	7	100.0	10	99.1
No. pref	12	100.0	0	100.0	3	100.0




Fig. 6. Graph of job preference received versus class rank using both the learning curve objective function (a) and the board score objective function (b).

number of cadets in the top quarter of the class who did not get their first choices. Conversely, the *board score* objective function does award three cadets AFSCs for which they had no preference, but these three cadets are at the bottom of the class. The overall effect of the *board score* objective function is that the model sacrifices the preferences of cadets at the bottom of the class to free higher preferences for those near the top of the class.

It is not our goal to specify a single objective function that will dominate all other objective

functions. Given that each year’s data changes dramatically, it is not sensible to expect that an objective function will have the same effect on the classification. We emphasize that the model does not establish policy. Rather, the model provide multiple solutions quickly, from which the Academy’s decision makers can determine the trade-offs and classification policy they prefer.

4.2. Results for Class of 2002

We developed the prototype of this model while the Class of 2001 classification was in-process. Because of the apparent success of the model as applied to the Class of 2001, we conducted a full proof-of-concept during the Class of 2002’s classification process. This class had similar characteristics to the previous class: 348 graduates were classified into 37 non-rated AFSCs.

4.2.1. Comparison with greedy classification

During the 2002 classification process, senior decision-makers specified an additional condition: the top 25% cadets in the class were guaranteed their top AFSC preference (assuming they were qualified). We added the following constraints to the model to enforce this requirement:

$$y_{ij} = 1, \quad \forall (i, j) : \text{rank}(i) \leq \frac{N}{4}$$

and $\text{pref}(i, j) = 1,$

where N is the number of cadets. In addition, we penalized heavily deviations from AFSC quotas until the model filled the targets exactly. Finally,

Table 2
Comparison of Greedy Heuristic and the actual classification with the optimization model for the Class of 2002

Job preference	Greedy Heuristic		Actual classification		Model results	
	# Assigned	Cum%	# Assigned	Cum%	Board score objective	
					# Assigned	Cum%
1	202	58.2	218	62.8	240	69.2
2	59	75.2	57	79.3	54	84.7
3	29	83.6	25	86.5	15	89.0
4	12	87.0	16	91.1	9	91.6
5	14	91.1	9	93.7	11	94.8
6	9	93.7	10	96.5	15	99.1
No. pref	22	100.0	12	100.0	3	100.0

we used the *board score* (Section 3.2.3) approach within the objective function.

The results are displayed in Table 2. The model outperforms the greedy heuristic, moving the percent of cadets receiving their first choice from 58% to just over 69%. The number of cadets receiving none of their job choices dropped from 22 to 3. Thus, applying the optimization-based classification model to the Class of 2002 data yield similar results to those found with the Class of 2001 data. In addition, Table 2 shows a summary of the actual classification created by the AFPC who also used an optimization-based approach. We see the clear advantage of our model both in terms of the increase in the number of graduates assigned to their top choice and in terms of the decrease in the number of graduates assigned to none of their choices.

4.2.2. Preference experiment

Prior to the actual classification process, during the spring of their junior year, we collected information regarding AFSC preference beyond the 1–6 ranking. We surveyed 350 participants (both rated and non-rated cadets) and compared the AFSC classification using this additional information {see (9)} with the existing 1–6 ranking procedure (8).

The method we used to capture relative preferences follows. We gave each cadet 1000 points to allocate among six AFSCs. Assigning an equal number of points to all six AFSCs indicates indifference among the six. Allocating all points to a single AFSC indicates a strong preference for that

single career field, but indifference among *all* remaining AFSCs. All non-selected AFSCs received zero preference points.

The relative satisfaction of a cadet-AFSC classification was found by dividing the points the cadet assigned to that AFSC by the maximum points the cadet assigned to any single AFSC. That is, for each cadet the value of $r(i, j)$ used in Eq. (9) is found by dividing the number of points assigned to AFSC j , $p(i, j)$, by the maximum awarded to any AFSC by that cadet:

$$r(i, j) = \frac{p(i, j)}{\max_{j'} p(i, j')}. \quad (11)$$

While there are a number of ways to collect and measure these relative preferences, this method renders a [0–1] scale.

Using *board scores* (Section 3.2.3) as the basis for the objective function coefficients, we first compare the models with the different preference measures in the objective function. In Table 3, we show how each model performs by counting the number of 1–6 preferences actually assigned. On the surface, using the relative preferences appears to hurt the quality of the solution. Fewer people receive their top choice. A more precise performance measure is how the models compare with respect to the *relative preferences*.

In Table 4, we see the positive effect of using the relative preferences in the objective function. The scale used in the left column shows the relative preferences. If a cadet receives a job for which they had assigned maximum points, their satisfaction level is 100%. Furthermore, if a cadet was indif-

Table 3

Preference received (on the 1–6 scale) using either 1–6 preferences or relative preferences in the objective function

Job preference	1–6 preference used in objective function		Relative preference used in objective function	
	# Assigned	Cum%	# Assigned	Cum%
1	227	64.9	209	59.7
2	43	77.1	56	75.7
3	17	82.0	25	82.9
4	20	87.7	19	88.3
5	19	93.1	17	93.1
6	19	98.6	19	98.6
No. pref	5	100.0	5	100.0

Table 4

Preferences received (measured with the relative scale) using either 1–6 preferences or relative preferences in the objective function

Satisfaction level (%)	Job preference measure			
	1–6 ranking		Relative preferences	
	#	Cum%	#	Cum%
100	234	66.9	238	68.0
90+	0	66.9	5	69.4
75+	8	69.1	14	73.4
50+	25	76.3	37	84.0
25+	20	82.0	17	88.9
10+	22	88.3	15	93.1
<10	29	96.6	19	98.6
0	12	100.0	5	100.0

ferent between his two top AFSC choices (which is impossible to determine in existing 1–6 ranking system), then being classified into *either* of those AFSCs would yield 100% satisfaction. Using the relative preferences, we see a large increase (from 76.3% to 84%) in the number of cadets at least 50% satisfied and a corresponding reduction in the number of *unsatisfied* cadets. Thus, adding relative preferences provides additional flexibility to better meet the desires of the graduation class.

5. Concluding remarks and future work

This optimization-based approach for the AFSC classification process provides a flexible tool to support the decision authorities at the Academy and the AFPC. We emphasize that from year-to-year, variation in the data exists such that selecting

a single method for creating the objective function is not sensible. The ease of incorporating different objective functions and the fast run time of the model allow us to quickly create a set of solutions that best reflect the philosophies of decision-makers.

In addition, our explorations of using actual AFSC preferences indicate value-added in changing from the current 1–6 ranking system. Clearly actual preferences provide more relevant information than arbitrarily assigning preferences to each classification. However, obtaining true preferences ignores the presence of gamesmanship, which is likely. We have not explored this effect.

Future extensions to the system include a web-base mechanism for displaying information about each AFSC, its projected availability, and current cadet interest in each AFSC. This website will collect preference data supporting the optimization. It will alleviate much of the time-consuming process of advising and counseling cadets about the likelihood of receiving certain AFSCs. The manner in which these relative preferences are collected is an area of further study. Finally, the fact that this model is a network flow problem (thus, the optimal linear solution is always integer) means we can exploit dual information to provide feedback on the relative interest of certain AFSCs. The duals could be the basis to coordinate target quotas and the AFSC classification across all three commissioning sources: the Academy, the Reserve Officer Training Corps (ROTC), and the Officer Training School (OTS).

Because of the demonstrated success of our model, AFPC has revised their process. They have implemented and will utilize this optimization

model for the graduating class of 2003. In addition, the Personnel Center is modifying existing software to incorporate this classification model for both ROTC and OTS career field classifications.

Acknowledgements

This work has been made possible through effort and cooperation of the AF Academy's Office of Cadet Personnel and the Air Force Personnel Center in San Antonio, TX. We wish to recognize the invaluable contributions of Randy Nicholson, Ed Duffett, Julia Phillips, Rich Linton, and Matt Speer.

References

- [1] R.K. Ahuja, T.L. Magnanti, J.B. Orlin, *Network Flows: Theory, Algorithms and Applications*, Prentice Hall, Englewood Cliffs, NJ, 1993.
- [2] T.A. Blanco, R.C. Hillery, A sea story: Implementing the Navy's personnel assignment system, *Operations Research* 42 (5) (1994) 814–822.
- [3] R.B. Chase, N.J. Aquilano, F.R. Jacobs, *Operations Management for Competitive Advantage*, McGraw-Hill, 2001.
- [4] S.C.K. Chu, C.K.Y. Lin, Manpower allocation model of job specialization, *Journal of the Operational Research Society* 44 (10) (1993) 983–989.
- [5] Dash Optimization, *Xpress essentials: An introduction to modeling and optimization*, Dash Optimization, Inc., Englewood Cliffs, NJ, 2001.
- [6] D.C. Dietz, M. Rosenshine, Optimal specialization of a maintenance workforce, *IIE Transactions* 29 (5) (1997) 423–433.
- [7] Frontline Systems, *Premium Solver Platform user guide*, Frontline Systems, Inc., Incline Village, NV, 2000.
- [8] I.A. Krass, M.C. Pinar, T.J. Thompson, S.A. Zenios, Network model to maximize Navy personnel readiness and its solution, *Management Science* 40 (5) (1994) 647–661.
- [9] T.T. Liang, T.J. Thompson, A large-scale personnel assignment model for the Navy, *Decision Sciences* 18 (1987) 234–248.
- [10] G.R. Reeves, R.C. Reid, Military reserve manpower planning model, *Computers and Operations Research* 26 (12) (1999) 1231–1242.
- [11] H.S. Weigel, S.P. Wilcox, Army's personnel decision support system, *Decision Support Systems* 9 (3) (1993) 281–306.
- [12] A.W. Whisman, J. Silverman, R.E. Steuer, Embedding a manpower optimization model in a managerial environment, *Mathematical and Computer Modeling* 12 (10–11) (1989) 1383–1391.