Proceedings of the Annual General Donald R. Keith Memorial Conference
West Point, New York, USA
April 23, 2026
A Regional Conference of the Society for Industrial and Systems Engineering

Optimizing Air Force Officer Accessions: Integrating Cadet Preferences Across AFSC, Base, and Training Assignments

Gershon Ellis, Katelin Marin, Luke Schlimm, Julian Whang, Benjamin Kallemyn, Taylor Leonard, and Gerry Gonzalez

Operations Research Program, United States Air Force Academy, Colorado Springs, Colorado 80841

Corresponding author's Email: gerry.gonzalez@afacademy.af.edu

Author Note: Gershon Ellis, Katelin Marin, Luke Schlimm and Julian Whang are first-class cadets at the U.S. Air Force Academy (USAFA) majoring in Operations Research and Gerry Gonzalez, Taylor Leonard, and Benjamin Kallemyn are faculty advisors. The views expressed herein are those of the authors and do not reflect the position of the United States Air Force Academy, the Department of the Air Force, or the Department of War.

Abstract: The Office of Labor and Economic Analysis (OLEA) and the Air Force Personnel Center (AFPC) currently use the One Market model to assign graduating officers to Air Force Specialty Codes (AFSCs). While the model fills AFSC quotas, it does not incorporate two factors that influence officer utility and retention: initial base assignments and initial skills training (IST) start dates. This capstone project develops a multi-attribute optimization framework that integrates cadet preferences across AFSCs, base locations, and IST start dates into a unified assignment model. Using synthetic datasets calibrated to AFPC and USAFA distributions, we evaluate how incorporating these preferences affects assignment feasibility and cadet utility. Results indicate that integrating these dimensions improves assignment feasibility and overall cadet utility while highlighting computational challenges at operational scale.

1. Introduction

The Office of Labor and Economic Analysis (OLEA) at the United States Air Force Academy (USAFA) provides quantitative analysis to support personnel policy across the Department of the Air Force. Our capstone team partnered with OLEA to evaluate improvements to the current officer accessions model used by the Air Force Personnel Center (AFPC). This model assigns graduating cadets and ROTC students to Air Force Specialty Codes (AFSCs), while base locations and initial skills training (IST) start dates are determined later in the process.

1.1 Background

AFPC currently uses the One Market model to assign graduating cadets to AFSCs. The model accounts for merit rankings and limited preferences while ensuring AFSC quotas are satisfied. However, base assignments and IST start dates are determined later in the accession process.

Because base location and training schedules strongly influence early career satisfaction, separating these decisions from the initial AFSC assignment may reduce overall cadet utility. Many new officers therefore receive assignments that only partially align with their preferences.

This project investigates whether a multi-attribute optimization framework can improve assignment outcomes by jointly allocating cadets to AFSCs, bases, and training pipelines. We treat the accession process as a constrained assignment problem and evaluate whether integrating these preference dimensions improves assignment feasibility and cadet utility.

1.2 Related Work

Previous research has applied optimization methods to officer career field assignments. Armacost and Lowe (2005) developed an optimization model for AFSC assignment at USAFA that balanced cadet preferences, merit rankings, and AFSC quotas. Their approach improved the percentage of cadets receiving preferred career fields compared to earlier manual assignment procedures.

Related work also examines preference-based matching systems in military assignments. Sönmez and Switzer (2013) propose a market design approach for branch assignment at the United States Military Academy that improves preference satisfaction while maintaining fairness and stability. These mechanisms draw on the broader literature on stable matching—an assignment framework in which no pair of participants would both prefer to be matched with each other over their current assignments—such as the Gale–Shapley algorithm and its extensions (Iwama & Miyazaki, 2008).

Finally, Patki, Wedge, and Veeramachaneni (2016) introduce the Synthetic Data Vault (SDV), a framework for generating realistic synthetic datasets while preserving privacy. We apply similar ideas to construct synthetic cadet populations that replicate observed assignment patterns without exposing sensitive personnel data.

2. Data and Methodology

To evaluate the impact of incorporating base and training preferences into the accession process, we develop a data-driven modeling framework grounded in both observed cadet data and synthetically generated populations.

2.1 Data

We integrate several datasets provided by OLEA and AFPC to model cadet preferences, historical assignment outcomes, and training pipeline constraints.

The 2025 Base of Preference (BOP) dataset contains cadet-submitted base rankings and selection outcomes for the Class of 2025. Figure 1 illustrates the distribution of first-choice base selections. Preferences are concentrated among a small number of installations such as Hill AFB, Hurlburt Field, and Eglin AFB. We use these distributions to generate realistic base preference rankings when constructing synthetic cadet populations.

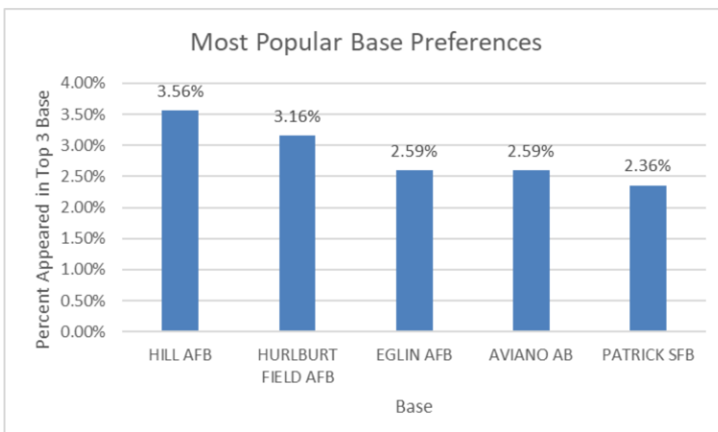

Figure 1: Most Popular Base Preferences in BOP dataset

The 2025 Aggregated Dataset contains survey results collected by OLEA describing AFSC preferences, merit scores, commissioning sources, and other cadet characteristics. These distributions parameterize the synthetic data generator used in the experiments.

The Officer Accessions Tracking System (OATS) dataset provides historical cadet-to-AFSC assignment outcomes and eligibility parameters used by AFPC. Although personally identifiable information has been removed, these records provide a baseline for validating that simulated assignment scenarios approximate operational conditions.

Additional inputs include IST training schedules, which specify course start dates and seat capacities for each AFSC training pipeline. These schedules define the training capacity constraints used in the assignment model.

2.2 Methodology

This project extends the existing One Market model into a multi-attribute assignment framework that incorporates cadet preferences across three dimensions: AFSC assignment, base location, and IST start date. The objective is to evaluate whether integrating these preference dimensions improves assignment feasibility and cadet utility while maintaining the

ISBN: 97819384962-5-7

operational constraints used in the current accession process. We begin by replicating the baseline One Market model using the existing framework. This step ensures consistency with the current AFPC assignment process, which allocates cadets to AFSCs while respecting eligibility constraints, commissioning-source quotas, and merit rankings.

Our analysis uses both empirical 2025 datasets and synthetically generated instances. The 2025 datasets—including the BOP dataset, aggregated preference survey data, and the OATS dataset—characterize observed cadet preferences, assignment outcomes, and training pipeline constraints. We use these empirical distributions to parameterize the synthetic data generator so that simulated cadet populations reflect realistic preference patterns and demographic characteristics.

Synthetic cadet populations are generated using an existing data generation module. This generator produces cadet instances that include ranked preferences over AFSCs, base locations, and IST start dates while preserving the statistical structure of the original datasets.

Using these datasets, we evaluate two assignment approaches. The first approach applies a greedy baseline algorithm that sequentially assigns bases and IST start dates after AFSC assignments are determined. The second approach uses a multi-attribute optimization model that simultaneously assigns AFSCs, bases, and training pipelines while maximizing total cadet utility subject to eligibility and capacity constraints.

Model performance is evaluated using simulated assignment scenarios that vary cadet population size and assignment conditions. Metrics include assignment feasibility, and cadet utility. We also conduct sensitivity analysis to examine how changes in training pipeline capacity affect assignment feasibility and preference satisfaction.

2.3 Greedy Baseline Algorithm

To benchmark the proposed optimization model, we implement a greedy baseline algorithm that assigns bases and IST start dates after AFSC assignments are determined. A greedy algorithm is a heuristic that makes locally optimal decisions at each step without reconsidering earlier assignments. Accordingly, the algorithm processes cadets sequentially in descending order of merit, assigning each cadet the highest-ranked feasible option from their preference list while respecting AFSC compatibility and capacity constraints.

Base assignments are determined by scanning each cadet's ranked base preferences and assigning the first base that is valid for the cadet's AFSC and has remaining capacity. IST assignments follow a similar procedure in which cadets are assigned to the first available training class associated with their AFSC.

Because the algorithm makes assignments sequentially without reconsidering earlier decisions, it may exhaust limited base or training capacity before lower-merit cadets are processed. Although this approach does not guarantee globally optimal outcomes, it provides a computationally efficient benchmark for evaluating the optimization model.

2.4 Modeling

Our model extends the One Market formulation by incorporating base assignment and IST start date decisions into the assignment framework. We formulate the problem as an unbalanced assignment model that maximizes total cadet utility subject to manpower, eligibility, capacity, and feasibility constraints. We implement the optimization model in Python using the Pyomo optimization framework and solve it using the Gurobi mixed-integer linear programming solver.

The extended formulation builds upon the One Market model developed by Laird, reproduced in Attachment 1. To represent the additional assignment stages, we introduce decision variables for base assignment and IST course assignment. These variables link AFSC assignments to valid bases and training pipelines while enforcing base capacity limits and training course quota constraints.

The additional constraints ensure that base and IST assignments are feasible and consistent with each cadet's AFSC. Base constraints restrict cadets to eligible bases and enforce installation capacity limits, while training constraints assign cadets to valid IST start dates subject to course seat availability. Linking constraints connect AFSC, base, and training decisions to prevent infeasible combinations.

Together, these constraints ensure operational feasibility while allowing the model to maximize cadet utility across all three assignment dimensions. The additional variables and constraints used to extend the baseline formulation are presented below.

Additional Decision Variables:

$$v_{ib},\ q_{ijc} \in \{0,1\}, \qquad CV_i \tag{1}$$

ISBN: 97819384962-5-7

where $v_{ib} = 1$ if cadet i is assigned to base b, and $q_{ijc} = 1$ if cadet i is assigned to course c for AFSC j. CV_i represents the value cadet i receives from their assignment outcome.

Additional Constraints:

$$\sum_{b \in B_i^E} v_{ib} = \sum_{j \in J^B \cap J_i^E} x_{ij} \qquad \forall i \in I \tag{2}$$

$$x_{ij} \le \sum_{b \in B_j^A} v_{ib} \qquad \forall i \in I,\ j \in J^B \cap J_i^E \tag{3}$$

$$x_{ij} = \sum_{c \in C_{ij}^E} q_{ijc} \qquad \forall i \in I,\ j \in J_i^E \tag{4}$$

$$lo_{jb}^B \le \sum_{i \in I_j^E} v_{ib} \le hi_{jb}^B \qquad \forall j \in J^B,\ b \in B_j^A \tag{5}$$

$$lo_{jc}^C \le \sum_{i \in I_{jc}^A} q_{ijc} \le hi_{jc}^C \qquad \forall j \in J,\ c \in C_j \tag{6}$$

3. Results and Analysis

We evaluate the greedy baseline assignment procedure and the proposed optimization model using an augmented 2025 accession dataset and synthetic cadet populations. The augmented dataset reflects operational assignment conditions, while synthetic datasets allow controlled experiments across different population sizes.

Across all experiments, we compare assignment feasibility and cadet utility between the two assignment approaches. These results illustrate how incorporating base and IST training preferences into the assignment framework affects overall assignment quality and computational performance.

In the first set of experiments, we evaluate both assignment approaches using the augmented 2025 dataset representing approximately 1,000 cadets that commission through USAFA. Table 1 summarizes the assignment outcomes that the greedy baseline assignment procedure and the proposed optimization model produce.

Table 1: Assignment Outcomes for 2025 Data, 1000 Cadets

Metric	**Greedy Baseline**	**Optimization Model**	**One Market Model (2850 cadets)**
Cadets Assigned to AFSC	1000/1000	1000/1000	-
Cadets Assigned to Base	924/1000	943/1000	-
Cadets Assigned to IST	885/1000	985/1000	-
Cadets Fully Assigned (AFSC + Base + Training)	878/1000	938/1000	-
Average Cadet Utility	0.7128	0.8604	0.9204

The optimization model produces higher-quality assignment outcomes than the sequential greedy baseline. While both methods assign all cadets to AFSCs, the optimization model increases the number of cadets receiving feasible base and training assignments and substantially improves average cadet utility. Because the optimization model considers AFSC assignments, base locations, and training pipelines simultaneously, it distributes cadets across available base and training capacity more effectively than the sequential greedy procedure.

To benchmark these results against the current AFPC accession process, we compare the proposed framework to the existing One Market model, which optimizes AFSC assignments only. When applied to the full 2025 accession population of 2,850 cadets, the One Market model achieves an average AFSC-assignment utility score of 0.9204. This benchmark reflects the high-quality AFSC-only matching currently produced by AFPC's assignment process. However, because the One Market model does not incorporate base or IST preferences, this utility score does not capture downstream assignment satisfaction associated with installation location or training timeline. The proposed multi-attribute optimization framework seeks to preserve similarly strong AFSC assignment performance while extending optimization across all three assignment dimensions.

Table 2: Greedy Baseline Performance Across Synthetic Population Sizes

Cadet Population	Fully Assigned	Average Cadet Utility
100	100/100	0.7502
250	250/250	0.7420
500	487 / 500	0.7231
1000	878 / 1000	0.7128

Table 3: Optimization Model Performance Across Synthetic Population Sizes

Cadet Population	Fully Assigned	Average Cadet Utility
100	100/100	0.8126
250	250/250	0.8348
500	490/500	0.8491
1000	938/1000	0.8604

To examine the robustness of the assignment framework, we conduct a sensitivity analysis that evaluates how changes in training pipeline capacity affect assignment feasibility and cadet utility.

Training pipeline capacity emerges as the primary limiting factor in both the augmented 2025 dataset and the synthetic experiments. To examine this effect, we vary total training seat availability by ±10% and ±20% relative to baseline capacity while holding all other parameters constant. Table 4 summarizes the resulting assignment outcomes for the optimization model using the 1,000-cadet synthetic population.

Table 4: Sensitivity of Assignment Outcomes to Training Capacity

Training Capacity	Fully Assigned	Average Cadet Utility
-20% capacity	842/1000	0.8121
-10% capacity	901/1000	0.8360
Baseline	938/1000	0.8604
+10% capacity	962/1000	0.8728
+20%	979/1000	0.8832

The sensitivity results confirm that training pipeline capacity is the primary limiting factor in the accession assignment process. Reductions in training seat availability significantly decrease the number of cadets receiving complete AFSC–base–training assignments, while increases in capacity produce measurable improvements in both feasibility and cadet utility.

Taken together, these experiments demonstrate that incorporating base and training preferences directly into the assignment framework improves assignment outcomes relative to the sequential greedy approach. At the same time, the computational demands of solving the optimization model at larger cadet populations highlight the need for scalable solution approaches that balance assignment quality with feasibility constraints.

4. Conclusion and Future Work

This study evaluates the impact of incorporating cadet preferences across AFSC assignments, base locations, and IST start dates into a unified accession assignment model. We compare a sequential greedy assignment procedure with a multi-attribute optimization model using both augmented operational data and synthetic cadet populations. Results indicate that integrating base and training preferences directly into the assignment framework improves overall assignment quality. Across both datasets, the optimization model produces higher cadet utility and a greater number of fully feasible AFSC–base–training assignment pipelines than the greedy baseline. In the augmented 2025 dataset, the optimization model increases the number of fully assigned cadets from 878 to 938 while substantially improving average cadet utility.

Comparison to the existing One Market model further contextualizes these findings. The One Market model achieves an AFSC-assignment utility of 0.9204 on the full 2025 accession population of 2,850 cadets, demonstrating strong performance when optimizing AFSC placement alone. However, because the One Market model framework does not incorporate base or IST preferences, it cannot optimize total assignment satisfaction across the full accession pipeline. The proposed model extends this framework by jointly optimizing AFSC, base, and training decisions within a single assignment process.

Synthetic experiments further demonstrate that the optimization approach consistently outperforms the greedy baseline, though solver runtimes increase as cadet population size grows. Sensitivity analysis identifies training pipeline capacity as the primary limiting factor in assignment feasibility, with training seat availability significantly affecting both complete assignment rates and cadet utility.

These findings have important implications for Air Force readiness and talent management. Better aligning cadet preferences with assignment outcomes may improve early-career satisfaction and retention, while more effective allocation across AFSCs, bases, and training pipelines supports the placement of talent where it is most operationally effective.

Overall, the results suggest that extending the One Market model framework to incorporate base and training preferences can improve assignment feasibility and preference satisfaction while preserving the strengths of AFSC assignment optimization. However, the computational complexity of solving the full model at operational scale indicates that additional research is needed to develop scalable solution methods for large accession populations.

Future research may extend this work by developing hybrid solution approaches that combine heuristic assignment methods with optimization techniques to balance computational efficiency and assignment quality. Additional model extensions could incorporate specialized training pipelines such as helicopter pilot training, Combat Rescue Officer (CRO) and Special Tactics Officer (STO) programs, and the Euro-NATO Joint Jet Pilot Training (ENJJPT) program. These pipelines introduce additional eligibility constraints and highly limited training capacity.

Finally, future work may also incorporate operational personnel constraints such as joint spouse assignments, coordinated training timelines, and assignments involving multiple commissioning sources, including ROTC and Officer Training School. Integrating these factors would allow the model to more accurately represent the full Air Force accession process.

5. References

Armacost, A. P., & Lowe, J. K. (2005). *Decision support for the career field selection process at the US Air Force Academy*. European Journal of Operational Research, 160(3), 839–850.

Iwama, K., & Miyazaki, S. (2008). *A survey of the stable marriage problem and its variants*. International Conference on Informatics Education and Research for Knowledge-Circulating Society, 131–136.

Laird, D. G., "A New Approach to Career Field Matching for Commissioning Air Force Cadets" (2022). Theses and Dissertations. 5346. https://scholar.afit.edu/etd/5346

Patki, N., Wedge, R., & Veeramachaneni, K. (2016). The synthetic data vault. *IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, 399–410.

Sönmez, T., & Switzer, T. B. (2013). *Matching with (branch-of-choice) contracts at the United States Military Academy*. Econometrica, 81(2), 451–488.