Proceedings of the Annual General Donald R. Keith Memorial Conference
West Point, New York, USA
May 2, 2019
A Regional Conference of the Society for Industrial and Systems Engineering

Optimizing the Special Forces Qualification Course: A Monte Carlo Simulation

Kayleigh Stallings[1] and Christopher Fisher[2]

[1]Department of Systems Engineering
United State Military Academy, West Point, NY

[2]Department of Mathematics
United State Military Academy, West Point, NY

Corresponding author: kayleigh.stallings@westpoint.edu

Author Note: Cadet Kayleigh Stallings is a System Engineering major at the United States Military Academy (USMA). MAJ Christopher Fisher is an Assistant Professor in USMA's Department of Mathematics currently serving as an analyst in the Operations Research Center.

Abstract: This study explores the success and attrition rates of the Special Forces Qualification Pipeline. It answers the question of what percentage of SFQC starts will become graduates. The model calculates a beta distribution of the success rates. The tool created is a comprehensive, modular, adaptable and interactive system that calculates and displays the requested success rate and graduating class size along with supporting statistics such as confidence interval and standard deviation. The success of this study relies on the accuracy of the calculated success rate.

Keywords: Special Forces Qualification Course, Monte Carlo Simulation, Bayesian Updating

1. Background

The current Special Forces Qualification Course (SFQC) pipeline is a system of six sequential courses that all must be completed to graduate. The six courses are Orientation (OR), Small Unit Tactics (SU), Survival Evasion Resistance and Escape (SERE), Military Occupational Specialty (MOS), Robin Sage (RS), and Language (LANG). Additionally, there is an administrative Graduation phase (GRAD). There are many different components of Soldiers which go through the SFQC pipeline, but this study focuses on Active Duty Officers (ADO). Every year the Special Forces (SF) branch receives a mission of how many ADO SF graduates they need (M. Gorevin, personal communications, February 11, 2019). Knowing how many trainees are required to achieve that quota is an important capability. Potential stakeholders include leadership and program directors for the SFQC. This analysis will allow the Army to have more precise control over the number of SF qualified personnel. There are other factors such as the assessment and selection course and timing of the Special Operations Captain's Career Course that could influence the system; these impacts should be examined in follow-on research.

2. Introduction

This study attempts to analyze the throughput of the Special Forces Qualification Course (SFQC) pipeline. Currently, when the SF branch receives the graduate quota, they use a deterministic method to calculate the necessary initial starts to achieve the quota, based on the average historical success rate for each course. This study attempts to account for uncertainty by creating a stochastic model to represent the SFQC pipeline. By incorporating uncertainty, this study provides more accurate assessment and allows decision makers to understand the risk of not reaching the quota. Having confidence intervals for the final number of graduates, based on the number of initial trainees, allows leadership to make educated decisions on how many ADO trainees to enroll in each class. Bayesian updating was used to estimate the distribution of success rates for each course in the SFQC pipeline. These distributions were then used in a Monte Carlo Simulation to calculate the success rates. The model simulated 10,000 SFQC iterations to create a distribution of SFQC final graduation rates. This methodology was built into a tool which SFQC analysts can use to quickly update their estimates based on the number of trainees and evolving success rates for each course. It also provides success and attrition rates for each individual course in the pipeline. This paper explains the method and tools used, and final product created.

3. Methodology

This section describes the method of calculating the distribution of SFQC success rates. The data used for this project was acquired from AR 350-10, Army Training Requirements &Resources System (ATRRS), and the United States Army John Fitzgerald Kennedy Special Warfare Center and School (USAJFKSWCS). The data provides entries of individual classes for each specified course. In this paper, each instance of a course is referred to as a “class,” and its graduation rate is referred to as a “success rate”. The six courses of the pipeline are referred to as “courses” and their graduation rate is referred to as a “pass rate.” Finally, the sequence of the six courses makes up the whole SFQC Pipeline and is referred to as the “pipeline,” and its graduation rate is referred to as the “graduation rate.” The data available consist of 930 classes which took place between 2010 and 2018.

3.1 Current State

Currently, the SFQC pipeline consists of six courses that must be completed in sequence. ATRRS specifies two input types and three output types a candidate can be categorized as when entering and exiting a class. They can start as either an initial start (I), or a recycle start (Q). A soldier exiting the class can exit with one of three statuses, graduate (G), recycle (L), or failure (Z). A candidate starts the pipeline as an initial start and is allowed two additional restart attempts if recycled.

3.2 Variables

The following variables are used and referenced throughout the paper.

$$v \in \{1,2,3\}, \textit{attempt number for a given course}$$
$$i \in \{1,2,3,4,5,6\}, \textit{indicator for course number}$$
$$j \in \{I, Q\}, \textit{candidate input status}$$
$$k \in \{G, L, Z\}, \textit{candidate output status}$$

$$S_{iv} = \textit{class success rate rate for course } i, \textit{attempt number } v$$
$$P_i = \textit{pass rate for course } i$$
$$G = \textit{graduation rate for SFQC Pipeline}$$

$$\alpha_i = \textit{alpha parameter for course } i$$
$$\beta_i = \textit{beta parameter for course } i$$

$$N_{ijk} = \textit{Number of candidate starts for course } i \textit{ with input status } j \textit{ and output status } k$$
$$d = 1 - \textit{dropout rate} = .95 = 95\%$$

Figure 1 depicts a model of the SFQC pipeline along with the respected course pass rates. This model is consistent with the deterministic model and only serves to demonstrate the construct of the pipeline.

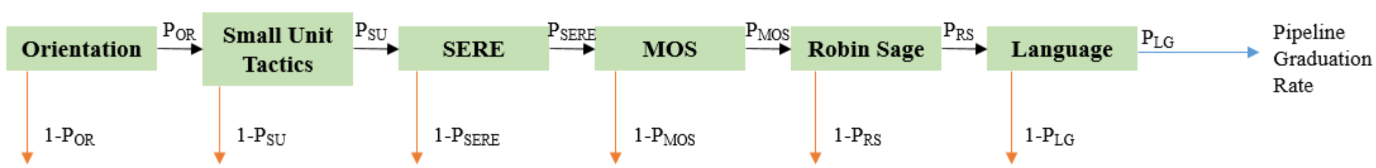

Figure 1. SFQC Pipeline Model

3.3 Model Overview

The problem at hand is transforming a deterministic pass rate into a distribution of pass rates that in aggregate can calculate a distribution of pipeline graduation rates. Due to the binomial nature of a pass rate (pass/fail), along with the iterative data used to calculate such pass rates, the method of Bayesian updating was used to calculate the distribution

(Holmes and Held, 2006). A Beta distribution has two parameters, α and β, that represent the number of successes and failures, respectively. Beta(1,1) was chosen as the prior distribution, with the intention of being non-informative prior (Kim, 2009). Once the observations from the data were used to find the posterior distributions of individual class success rates, Monte Carlo Simulation was used to determine the distribution of overall course success rates.

The method can be broken down into three steps. The first step is calculating the individual class success rates, S_{iv}. The second step is estimating the pass rate of each course referred to as pass rates, P_i. The third step is calculating the final graduation rate distribution, referred to as the final graduation rate, G. At the end of the first step, every class in every course will have its own pass rate and values for success and failures. In the second step, these individual class success and failure values are summed up for the entire course, and are added to the prior values for α and β, respectively. The third step combines all course pass rates by multiplying the beta distributions together to return a final graduation rate.

3.3.1 Determining Individual Class Success Rates

The first step in determining the final graduation rate is finding the intermediate pass rate for each class, in each course S_i. The success rate for each class is defined as the probability of success (graduation) given any start type. Because of the possibility of a Soldier recycling, the success rate was calculated as the sum of three variables. The first variable, S_{i1}, represents the success rate of soldiers with an initial start status. To calculate this variable, the sum of all initial starts with a graduation output was divided by the total number of initial starts, per class, shown in Equation 1. The second variable, S_{i2}, represents the probability for an initial start to succeed after being recycled exactly once. In other words, it is probability of recycle output given initial start, and the probability the recycled trainee decides to recycle back in with the probability of success given recycle start, shown in Equation 2. The third variable, S_{i3}, represents the probability for an initial start to succeed after being recycled exactly twice. This is equal to the probability of two recycles followed by a success, as shown in Equation 3.

$$S_{i1} = \frac{N_{iIG}}{\sum_k N_{iIk}} \tag{1}$$

$$S_{i2} = \frac{N_{iIL}}{\sum_k N_{iIk}} * \left[d * \frac{N_{iQG}}{\sum_k N_{iQk}}\right] \tag{2}$$

$$S_{i3} = \frac{N_{iIL}}{\sum_k N_{iIk}} * \left[d * \frac{N_{iQL}}{\sum_k N_{iQk}}\right] * \left[d * \frac{N_{iQG}}{\sum_k N_{iQk}}\right] \tag{3}$$

$$S_i = S_{i1} + S_{i2} + S_{i3} \tag{4}$$

The overall success rate for a class, then, is the sum of S_{i1}, S_{i2}, and S_{i3} as shown in Equation 4. The addition operator was used because the probability of success for an initial start is the probability they succeed on their first start attempt, *or* they succeed on their second attempt *or* they succeed on their third attempt. This method was used for each class in the data set. This resulted in a distribution of success rates, which was used to inform the course pass rates as explained in the following section.

3.3.2 Estimating Distribution of Course Success Rates using Bayesian Updating

Bayesian updating is an iterative way of creating a distribution to fit a data set. In this case, Bayesian updating was used to determine the distribution of pass rates for each course, P_i. Equation 5 simplifies the Bayesian method of using a likelihood function to update a prior distribution, resulting in the refined posterior distribution.

ISBN: 97819384961-6-5

$$Posterior \propto Likelihood\ X\ Prior \tag{5}$$

The likelihood function is every observation, or in this case, every individual class success and failure, previously calculated from S_i . The initial prior distribution is what subject matter experts believe the distribution to look like. Due to the sufficiently large number of available observations, this model uses a non-informative prior of Beta(1,1). This uninformed prior distribution is valid because the number of observations (classes) is sufficient to make the posterior distribution robust and credible (Chen, 1984). After the initial prior distribution is updated with the first class's data, the posterior distribution then becomes the prior distribution and is iteratively updated with every new observation (likelihood function). With both the prior and likelihood function being beta distributions, the final pass and graduation rate distributions can be determined by adding all successes to α, and all failures to β. Both the likelihood and prior distribution are Beta distributions whose parameters α and β are added together to calculate the posterior distribution (Jaffray, 1992).

As mentioned above, the number of successes and failures from each class can be calculated and summed to determine the α and β parameters for each course. To calculate the number of successes from the success rate, the respective class success rate is multiplied by number of initial starts in each class. All class successes for a given course are then summed to calculate the α parameter, as shown in Equation 6. The β parameter is the number of failures per class in each course, which is found by subtracting the α parameter from the total number of initial starts across all classes. This calculation is found in Equation 7. α and β are now the parameters for the beta binomial distribution of the course success rate, as shown in Equation 8. Supporting statistics can be pulled from this distribution to further understand the data. P_i, then, is a distribution of success rates. When used in a Monte Carlo simulation, the output is 10,000 iterations of the specified course success rates.

$$\alpha_i = \sum (S_i * (N_{ilk})) \tag{6}$$

$$\beta_i = \left(\sum N_{ilk}\right) - \alpha_i \tag{7}$$

$$P_i = Beta\,(\alpha_i, \beta_i) \tag{8}$$

3.3.3 Calculate the Final Graduation Rate

The final graduation rate, $\boldsymbol{G}$, is found by taking the product of all course success rate distributions, $\boldsymbol{P_i}$, shown in Equation 9. This assumes that all pass rates are independent allows G to be calculated as the multiplicative product of $\boldsymbol{P_i}$. This seems to be a valid assumption, since the courses are strictly sequential, and the observations come from different classes spread out over nearly a decade.

$$G = \prod_i P_i \tag{9}$$

3.3.4 Monte Carlo Simulation

The Excel add-in SipMath was used to conduct a Monte Carlo simulation with 10,000 iterations, which provide a distribution of pass rates for each course (Mittal, 2018). Equation 9 was then applied to find the corresponding distribution of SFQC graduation rates. Section 4 details the analysis of the simulation results.

4. Results and Analysis

4.1 Results

The model outputs consist of averages, standard deviations, and confidence intervals for pass rates of every course in the SFQC pipeline, as well as the cumulative graduation rate. Table 1 provides the individual course pass rates for the ADO component. The mean pass rates align with the deterministic pass rates.

Table 1. Course Breakdown

Course	Mean Pass Rate	Standard Deviation	5th Percentile	95th Percentile
OR	96%	0.49%	95%	97%
SU	90%	0.70%	90%	93%
SERE	99%	0.24%	98%	99%
MOS	92%	0.70%	90%	93%
RS	96%	0.54%	95%	97%
LANG	98%	0.38%	97%	99%
GRAD	99%	0.01%	99%	100%
TOTAL	75%	1.12%	72%	77%

Figure 2 displays the histogram of the ADO graduation rate, created from the Monte Carlo Simulation. The horizontal axis shows the graduation rate, and the vertical axis shows the frequency of that success rate across 10,000 iterations of the simulation.

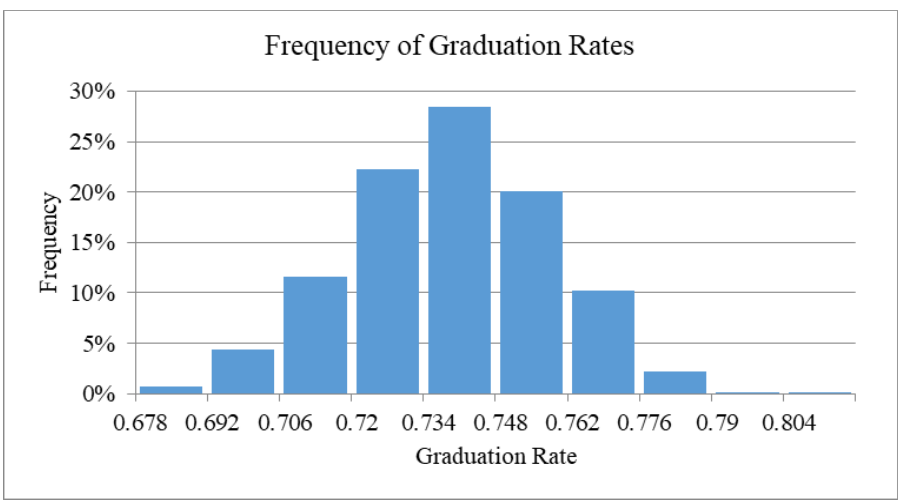

Figure 2. Histogram of SFQC Graduation Rates

Figure 3 shows a graphical depiction of the attrition of the ADO pipeline given an arbitrary start value of 100 recruits. This graph includes the expected value, upper 95th percentile, and lower 5th percentile for how many of the 100 initial trainees will remain in the SFQC after each phase. This is a key component of the tool that cannot be provided by a deterministic model. Note that the lower percentile line represents the fifth percentile graduation rate occurring for every course, the upper percentile represents the 95th percentile graduation rate occurring for every course, and the mean line follows suit. These are meant to provide a plausible range of graduates at the completion of each course as opposed to just the average, which is what their current deterministic models provide. By displaying the “best” and “worst” case scenario, it

ISBN: 97819384961-6-5

allows leadership to identify outlier situations and adjust accordingly. Additionally, if a course has a success rate drastically lower than the 5th percentile value it should alert leadership of potential issues within the course.

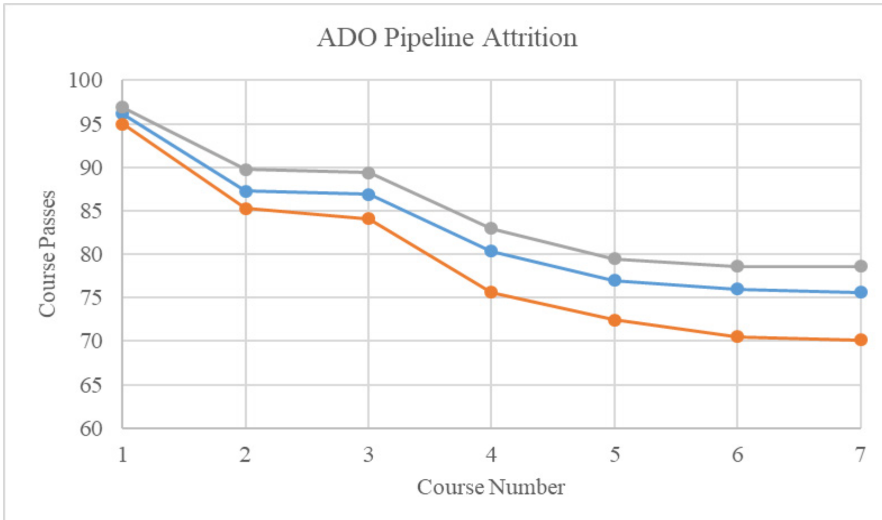

Figure 3. ADO Pipeline Attrition

4.2 The Tool

All of the analysis conducted, including reading the data, determining the course pass rates, executing the simulation, and generating the visualizations of the simulation outputs, was built into an easy-to-use tool within Excel. The purpose of the tool is to automate the creation of a report for SF leadership, enabling quick analysis going forward. The tool provides a confidence interval for the number of graduates, given a user input number of recruits. It is user- friendly with clear and concise input fields required. The tool is modular, it is easy to combine with other technologies such as discrete event simulations, and easy to add graphs or other required statistics. If SF leadership chose to rearrange the pipeline, the tool can adapt to accurately represent the new configuration. Finally, it is easily updated which will allow it to be maintained by SF analysts as they simply add more observations of new classes to the model. The tool will provide SF leadership with dynamic and accurate analysis to inform recruiting goals and the current risk of not meeting SF quotas based on the number of recruits.

The tool has two main reports. The first is the Course Breakdown and the second is the SFQC Overview. The Course Breakdown is an interactive page which displays detailed analysis for each course in the pipeline one at a time. It allows users to input the number of initial starts for the course and the confidence interval they wish to explore. It also includes a date filter to include or exclude course data, in the event that course structure changes. The outputs of this report are the average value of distribution for the desired course success rate, G_i, and the number of anticipated graduates calculated from the given input of initial starts.

The SFQC Overview displays the whole pipeline success rate with supporting statistics, including the breakdown of successes and failures of each course. The purpose of this page is to provide a general overview of the throughput of the SFQC pipeline. Future iterations of both reports will include other input components besides Active Duty Officers, such as Enlisted and National Guard trainees.

4.3 Comparison with Deterministic Model

The tool described above provides insight that a deterministic model cannot offer. Current deterministic models can simply provide a mean graduation rate that offers little insight to the full distribution of possible graduation rates. For example, by only providing a mean rate, currently recruiting efforts can either achieve mission success (recruit the number required based on the mean), or not (recruit fewer than this number). The stochastic model eliminates this binary evaluation system by allowing leaders to understand, based on a given number of recruits, the probability of achieving their quota for

ISBN: 97819384961-6-5

graduates. By assessing the current situation and evaluating risk, leaders can adjust recruiting goals for future classes in order to achieve a desired balance.

Another advantage of this stochastic approach is that it allows leadership to understand how variance in graduation rates can affect the number of graduates. For example, Figure 4 shows the distributions for the Language course compared to the SERE course distribution. A leader can accept risk differently after analyzing the two graphs below. Even though the two courses have similar mean pass rates, the Language course is normally distributed with a greater variance. The Language course has an equal probability of the actual graduation rate being lower or higher than the mean, as opposed to the SERE graduation rate that has a greater probably of being on or above the mean.

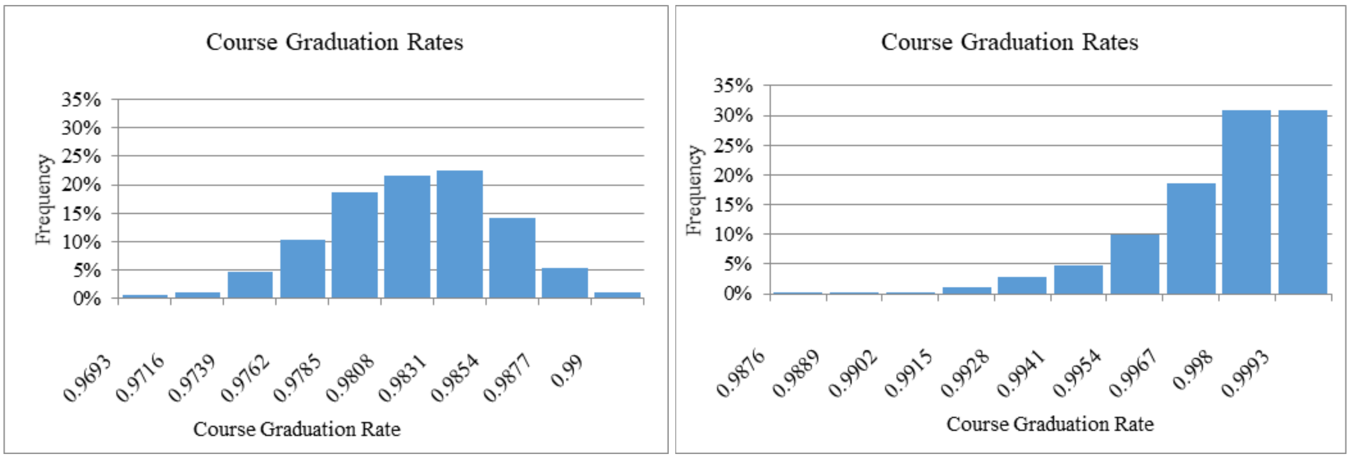

Language: Mean= 98% SERE: Mean=99.8%

Figure 4. Language and SERE Course Distributions (2010-2020)

6. Conclusion and Future Work

6.1 Conclusion

Historically, SF analysts have used a deterministic model to analyze average graduation rates and projected number of graduates from the SFQC. By incorporating stochastic methods using Bayesian updating, leaders can now make better-informed decisions about recruiting goals and can better understand how trainees attrite throughout the SFQC pipeline. With this stochastic model, a leader can now intelligently assume risk when necessary.

Importantly, the analysis was conducted using a custom tool which was built in Excel. While other, more advanced software are available for this type of analysis, conducting the work in a program which the unit has access to was critical to ensuring that the work done can be easily repeated, and is not one time only.

6.2 Future Work

The data filtering feature of the current tool allows analysts to find the most current statistics or view the entirety of the data to view long term change. It should be noted, however, that given the nature of Bayesian updating and using a non-informative prior, the more classes (wider date range) examined, the tighter the distribution of graduation rates will be. The opposite is true when evaluating fewer classes: the tighter the date range, the wider the resulting distribution will be. If a leader was to look at a tight date range with few classes per course, analysts may consider changing the prior from Beta(1,1) to a more informative prior based on subject matter expert opinion. Follow-on research can be conducted on how to appropriately assess small and large date ranges.

This project's scope was limited to the SFQC courses and the pipeline as a whole, however further research can be done examining the attrition rate of each instructional block inside each course. Generally, it is most beneficial from a cost and time perspective, for most of the attrition to take place at the beginning of a system. By analyzing the instructional blocks, leadership can choose to reorganize them according to pass rates where possible.

ISBN: 97819384961-6-5

Additional follow-on research topics include incorporating the Enlisted components in the model, the creation of a tool that calculates the necessary recruits for a given graduation quota and confidence and replicating the model calculation technique with classical statistics opposed to Bayesian updating.

7. References

Chen, J. J., & Novick, M. R. (1984). Bayesian analysis for binomial models with generalized beta prior distributions. *Journal of Educational Statistics*, *9*(2), 163-175.

Holmes, C. C., & Held, L. (2006). Bayesian auxiliary variable models for binary and multinomial regression. *Bayesian analysis*, *1*(1), 145-168.

Mittal, V. (2018, June). The Use of Stochastic Value Models to Create Technology Roadmaps. In *2018 IEEE Technology and Engineering Management Conference (TEMSCON)* (pp. 1-6). IEEE.

Jaffray, J. Y. (1992). Bayesian updating and belief functions. *IEEE transactions on systems, man, and cybernetics*, *22*(5), 1144-1152.

Kim, B. C., & Reinschmidt, K. F. (2009). Probabilistic forecasting of project duration using Bayesian inference and the beta distribution. *Journal of Construction Engineering and Management*, *135*(3), 178-186.