Proceedings of the 2023 Winter Simulation Conference
C. G. Corlu, S. R. Hunter, H. Lam, B. S. Onggo, J. Shortle, and B. Biller, eds.

SIMULATING TECHNICIAN POPULATIONS WITH TANDEM ANALYTIC AND DISCRETE EVENT MODELS

George Ryan Ambrose
François-Alex Bourque

Centre for Operational Research and Analysis
Development Research and Defence Canada
60 Moodie Drive
Ottawa, ON K1A 0K2, CANADA

ABSTRACT

Military workforce modelling is typically limited to either a series of analytic equations, or a simulation model. However, developing two such models in tandem has the benefit of cross-validation as well as the opportunity to explore problem space not easily accessed by a single approach. In particular, business rules for force employment are not easily described by closed-form equations while simulation models require exceedingly large computational resources to reach the asymptotic behaviour provided by analytic equations. This work leverages the benefits of both approaches to describe the population and career trends of technician individuals. As this career tends to have well defined training requirements, hence clear delineation between semi-functional apprentices and fully-functional journeymen, it is well suited to population modelling. Notional distributions for career parameters are assumed and the results for career progression and fleet readiness are compared.

1 INTRODUCTION

For military occupations in the Canadian context, an important element of professional development is on-the-job training. Mentoring for technical trades is even formalized through an apprentice-journeymen progression. This takes the form of personnel joining units as apprentices and training under the supervision of journeymen, before becoming journeymen and, eventually, managers. Importantly, only journeyman are fully functional and productive. As such, the population of journeymen is an essential factor in the unit productivity. For example, to maintain aircraft operational readiness requires a number of experienced technicians.

Most models of personnel based on differential equations (Boileau 2012; Vincent and Okazawa 2019), Markov decision processes (Zais and Zhang 2015; Diener 2018; Suvorova et al. 2019), and discrete-event simulations (Novak et al. 2015; Henderson and Bryce 2019) come short of accounting for career progression other than rank. Nor do the militarily workforce models reviewed by (Bastian and Hall 2020) include an apprentice-journeyman progression. Notable exceptions are a study on pilots based on system dynamics (Séguin 2015) and a series of theoretical works on technicians that model the dynamical aspect of the apprentice-journeyman population through a modified predator-prey model (Schaffel et al. 2021; Lahteenmaa-Swerdlyk. et al. 2023).

In this paper, we introduce two models in tandem that account for the population dynamic of the apprentice-journeyman-manager progression in technician occupations: one analytically solvable and one implemented as a simulation. The former follows (Schaffel et al. 2021), but includes managers and excludes training limitations due to the lack of journeymen (mentors). The latter is based on using discrete-event simulations to track the progress of personnel through a military training system as in (Novak et al. 2015;

Henderson and Bryce 2019). An immediate benefit of introducing two very different models is the ability to validate not only the implementation of the complex simulation model, but also the usefulness of the faster analytical model.

The goal for these models is to develop strategies that mitigate any drop in operational readiness during periods of transition or change for the technician career. For example, the emergence of a new occupation may cause concerns about personnel drain in other fields to compensate, or even failing to reach a target readiness by a particular time. By modelling the population of apprentices, journeymen, and managers, as well as the process by which individuals traverse this career, strategies are developed to ensure a healthy population can be established and maintained. This is done by exploring the effects of changing the various model parameters on the fully-functional, journeyman population.

The remainder of this report is organized as follows. Section 2 describes the analytic model of population dynamics, with Section 2.1 deriving the equations themselves and Section 2.2 highlighting how the equation parameters are obtained. Conversely, Section 3 delves into the discrete event model. Section 3.1 summarizes the simulation framework Origame while Section 3.2, and all the subsections within, explains how the model is designed. Section 4 provides a comparison between the two models with interpretation of these results. Lastly, this work and the impact it has is summarized in Section 5.

2 ANALYTICAL MODEL

An analytic model is well suited to capture the average behaviour of population dynamics. The results from this analytic method can serve as an important comparison to both the simulated results as well as real world data.

This section will outline how the analytic results are obtained for comparison in Section 4. Section 2.1 will motivate, and give the form of, the population dynamics. Section 2.2 delves into how the dynamical variables are estimated from underlying distributions.

2.1 Population Dynamics

Technician individuals in this model arrive from various sources (arrivals, with time-to-arrival), to dwell within the scope, and depart by falling into a sink (departures, with time-to-departure) after some time. Unsurprisingly, this lends itself well to a differential equation formulation relating these changes in population to various parameters obtained either empirically, or externally imposed. As a basic description, the differential equations built must relate the changing populations to those mechanisms that either increase or decrease the number of individuals. Each career milestone has unique influences on the corresponding population, leading to a set of three differential equations

$$\frac{\mathrm{d}P_{\mathrm{A}}(t)}{\mathrm{d}t} = I_{\mathrm{A}} - TP_{\mathrm{A}} - L_{\mathrm{A}}P_{\mathrm{A}}, \tag{1}$$

$$\frac{\mathrm{d}P_{\mathrm{J}}(t)}{\mathrm{d}t} = I_{\mathrm{J}} + TP_{\mathrm{A}} - L_{\mathrm{J}}P_{\mathrm{J}}, \tag{2}$$

$$\frac{\mathrm{d}P_{\mathrm{M}}(t)}{\mathrm{d}t} = I_{\mathrm{M}} - L_{\mathrm{M}}P_{\mathrm{M}}. \tag{3}$$

Where Equation 1 describes how the apprentice population is expected to change, Equation 2 similarly applies to the journeyman population, and Equation 3 finishes with the manager population. For each of these equations: $P_x(t)$ represents the current population at time t, I_x gives the influx of individuals per unit time, and L_x is the proportion of the population that leaves per unit time, where x is the career milestone (A is apprentice, J is journeyman, and M is manager). Lastly, T is the proportion of the population that completes training per unit time. As expected, the influx of technicians is a positive contribution to the changing population while the efflux has a negative contribution. The effect of training is more complicated; while it causes a decrease in the apprentice population, it is equally compensated by an increase in the journeyman population coupling the two equations.

Equations (1)-(3) can be solved, using the methods of ordinary differential equations, to get relations between technician population and time. The resultant equations are

$$P_{\mathrm{A}} = \frac{I_A}{T+L_A}\left(1-e^{-(T+L_A)t}\right), \tag{4}$$

$$P_{\mathrm{J}} = \frac{e^{-L_J t}}{L_J}\left(\frac{I_A(L_A-T)(L_J e^{(L_J-L_A-T)t}+e^{L_J t}(-L_J+L_A+T)-L_A-T)}{(L_A+T)(L_J-L_A-T)}+I_J(e^{L_J t}-1)\right), \tag{5}$$

$$P_{\mathrm{M}} = \frac{I_M}{L_M}\left(1-e^{-L_M t}\right), \tag{6}$$

where the parameters are the same as described before. In particular, Equation 4 shows the apprentice population as a function of time, Equation 5 similarly represents journeyman levels, and Equation 6 applies to management positions.

2.2 Coefficient Estimation

While the numbers of input technicians (I_x) will be adjusted for the dynamic to reach the desired population levels, the coefficients for technician loss per unit of time (L_x and T) are assumed to be obtained from the means of their underlying probability distributions. For T, this procedure is straightforward to follow as its distribution is assumed given. For a L_x, the assumption is that the distribution of dwell time in the milestone x must be derived from two other distributions, namely that for the years-to-arrival and years-to-departure. A further complication is the non-independence of these two distributions, the numbers of years-to-departure must necessarily depend on the number of years-to-arrival.

To resolve the dwell time for technicians, the joint probability distribution is built between the years-to-arrival and the years-to-departure distributions for a given milestone x. This approach is taken as it yields a probability distribution on all pairs of inputs, while encoding the conditional probability distribution between them.

The joint probability distribution is built by

$$p_{\text{Arr, Dep}}(t_1,t_2) = P(\text{Dep}(t=t_1)|\text{Arr}(t=t_2))\cdot P(\text{Arr}(t=t_2));\ t_1 \geq t_2. \tag{7}$$

Here, $\text{Dep}(t)$ refers to the time-to-departure distribution and $\text{Arr}(t)$ refers to the time-to-arrival distribution, where t_1 and t_2 are two independent times from each distribution respectively. Also notice that the above is specifically conditioned on the arrival distribution, as this determines which values of the departure distribution are feasible (the technician cannot leave before they enter, $t_1 \geq t_2$).

With Equation (7) the probability distribution function for dwell time can be obtained from

$$p_{\text{Dwell}}(t_3) = \int_\tau p_{\text{Arr, Dep}}(\tau,\tau-t_3)\,\mathrm{d}\tau, \tag{8}$$

where τ is the integration variable for time-to-departure and t_3 is the dwell time, yielding $\tau - t_3$ as the time-to-arrival as needed. The final step is obtaining the average, or more formally calculating the expectation value, for probability distributions like Equation (8). This is calculated according to

$$\mathrm{E}[X] = \int_{\mathbb{R}} x f(x)\,\mathrm{d}x, \tag{9}$$

for continuous random variable x on probability distribution function $f(x)$. With this, all of the parameters for Equations (4)-(6) are either externally set or an input, the population levels can be determined.

As a supplementary note, and perhaps motivation for Section 3, one can see potential issues with applying Equation (9) to real-world data. In particular non-monotonic, complex probability distribution functions are quite likely. The restriction of reducing a probability distribution to a single coefficient causes a significant amount of information to be lost, where this lost nuance may be critical for making informed real-world decisions. Thus, a procedure to randomly sample these series of distribution and build population trends accordingly would be beneficial leading to the formulation of a discrete event model.

3 DISCRETE EVENT MODEL

In a perfect world, the analytic model and equations presented in Section 2 would fully capture the technician population and all relevant career milestones. All that would be required is sufficiently detailed information input into the analytic model to recreate the technician population, and changing these inputs to learn how these groups would change. However, the real world is more complex than this analytic model. Various aspects of the business rules dictating the technician career are difficult, if not impossible, to translate into a mathematically closed form as per Equations (4)-(6). This is to say nothing about the unofficial rules and exceptions that only serve to complicate an already intractable problem.

A simulated model allows for the flexibility to capture some of these complicating aspects, however the choice of model is important. In particular, the simulation must accommodate the following requirements: at any point in time technicians are in a particular state defined by their individual history, the flow between states is calculated using empirical probability distributions conditioned on individual technician attributes, and the final output from the simulation are career trajectories of individual technicians so that various population metrics can be evaluated. Discrete-event simulated approach satisfies these requirements and is thus the framework employed.

This section will outline how the discrete-event simulation was designed, implemented, and interpreted. Section 3.1 introduces the software chosen to build the discrete-event model, Origame. Section 3.2 provides an overview of the simulation scope, and what aspects are captured. Section 3.2.1 describes the process of building the probability distributions and other information required to run the simulation. Section 3.2.2 details the process by which technicians enter the simulation. Section 3.2.3 describes the career trajectories of technicians while in the simulation. Lastly, Section 3.2.4 explains how technicians leave the simulation.

3.1 Origame

Discrete event simulation is a mature tool in the repertoire of computational analysis. Unsurprisingly then, there is a wide variety of potential software programs to satisfy the role of discrete event simulator. For this particular project, the Defence Research and Development Canada, Centre for Operational Research and Analysis (DRDC-CORA) developed, python-based Origame (Okazawa 2013) simulation environment was selected. There are two reasons for this selection: flexibility and security.

In general, discrete-event simulation packages can be placed along a structured vs flexibility axis. That is to say, these environments typically have either specific, robust functionality (e.g. predefined objects and relations) or allow greater flexibility (e.g. allow user to write custom classes). Origame falls into the latter class; it provides the user with a convenient GUI to organize the overall structure of the simulation but the actual simulation procedure is user implemented. Specifically, users in Origame must write their own python code to dictate what the subject of the simulation is, what properties they have, how they are stored, how they traverse trajectories, etc. Given the previously mentioned complexity of the technician career trajectory, this flexibility allows for the simulation to better replicate reality.

A more subtle reason to proceed with Origame is the advantage provided by the fact that this is a DRDC-CORA developed tool. This greatly simplifies the process to port the simulation to machines that are capable of handling sensitive information. As this project is intended to handle real-world data in the future, this is a serious aspect to consider.

3.2 Implementation

The subjects of this simulation are individual technicians where their career progression is the topic of interest. Thus the design of the technician subject and the various flows allowed for their career need careful consideration. As alluded to in Section 3.1, technicians are implemented in the simulation as a custom python class. This allows for individual technicians to contain all relevant human resource information as attributes (ranks, occupations, units, training, milestones, and all associated dates) as well as methods to set and get these properties.

Flow of technician objects along a career trajectory is performed by signalling functions to execute at particular times. For example, once a technician object is created a function to set some training qualification is called after some elapsed simulation time. This flow between functions is divided into three categories: source, state change, and sink. Source are the collection of functions that dictate how technician objects are created. State change determines how these technician objects change while within the scope of the problem. Sink is the final destination for technician objects, when these functions are called the object is taken out of the scope of the study. A summary graphic of this general flow is provided in Figure 1

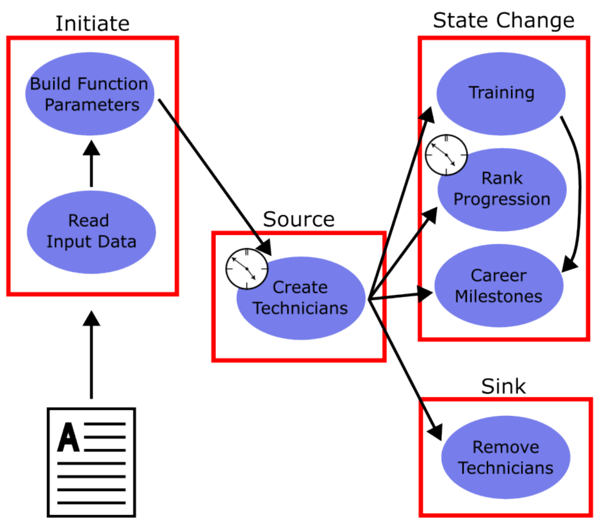

Figure 1: Summary of the simulation design to capture technician careers. Starting from the bottom left, clockwise: data is read into the **Initiate** group of functions to build various probability distributions, the **Source** set of functions are triggered to create custom technician objects, **State Change** functions controls career development, and **Sink** functions remove technicians from the task queue. The clock icon indicates functions that are self-signalling after a fixed interval.

3.2.1 Initialization

The very first step of the simulation is to build the various distributions that dictate the career flows, create structures to store the data, and begin the function calling cascade. This is all performed in the initiation section of Figure 1. Put simply, it is assumed that a spreadsheet exists that contains all the desired empirical information which is then read into the simulation. Specifically, human resource fields (occupations, units, etc) are used to build dictionaries to store technician objects, organizational charts, probabilities, and training definitions. Any supplied empirical probability distributions undergo a Gaussian kernel smoothing procedure (Tanner and Wong 1983) to build probability distribution functions for the simulation. All other details are read in where appropriate, for example if an initial population is specified it will fill the technician object dictionary. The last step of the initialization procedure is to call each source function, propagating the function cascade.

3.2.2 Source

Once initialized, the simulation can begin in earnest by creating new technician objects in the source functions. From this origin point, they will undergo some career progression to potentially contribute to the readiness of the fleet before falling outside of this scope by some means.

At this point, all potential real-world sources for technician individuals must be defined in concrete terms. As elaborated in Section 1, it is assumed that there are three possible technician career milestones: apprentices, journeymen, and managers. Accordingly, there are a set of three functions to create technicians in these specific milestone stages and begin their unique career trajectories. Each of these functions are informed by real-world data to faithfully re-create the existing technician population. For example, suppose that an occupation of technicians is being considered where there is a single source of apprentices, a trades school, and two sources of journeymen, inter-occupational transfer and advanced training. In the model for this occupation, the apprentice source function would create apprentice technician objects at the input graduation rate for this trades school. In the journeyman source function, the rate of technician object creation is instead the sum of the individual sources.

When a new technician object is created, all of their attributes are updated to reflect the human resource information for this new individual. Among these attributes are some that have some fixed value: hire date, career milestone, component, and unit. Others draw from supplied probability distribution functions: birth date, rank, occupation, and years served upon entry. Again, consider the example from the previous paragraph. Here, apprentice technicians only have a single source and so all human resource information is pulled from this single record. Journeymen technicians on the other hand, while sharing fixed values of hire date, milestone, etc., will have probability distributions informed by multiple sources.

After the technician object is fully initialized the career flows for the individual need to be triggered. This includes training certifications, rank progression, and career milestone development. As expected, not all state change functions are appropriate for all incoming technicians (like managers requiring training certifications) and so typically only a subset of these functions are called where appropriate.

The last aspect of the technician career that must be specified is the duration. If the fate of the technician is determined at arrival, then the dwell time (see Section 2.2) is used to trigger the sink function; otherwise, the determination is done during the simulation. In both cases, the actual fate is randomly choosen based on the relative probabilities of possibilities (i.e., retirement, transfers, etc) at the time of departure following the usual rule used in the Gillepsie algorithm (Gillespie 1976); the only difference between the approaches being when the departure time is actually resolved.

3.2.3 State Changes

While stationed within the fleet, technicians contribute to the readiness only after they reach full functionality. This is assumed to be achieved once they complete required training, or enter the journeyman career milestone. Some technicians enter the fleet fully functional, but most enter as an apprentice and require on-the-job training to reach this state. This is the primary purpose of the state change section of Figure 1, to properly model how technicians can become fully functional and contribute to fleet readiness. To this end, there are three aspects that are considered: obtaining training, rank promotion, and career milestone progression.

Required training for technicians is among the input of Section 3.2.1 and is typically dependent on the individual occupation. In real-world data the list of required training can be obtained from documents like fleet employment plans, where there are often dozens of subjects with different tiers of proficiency. However, the true challenge is obtaining timing data for this level of delineation. Rarely is it feasible to get a sufficient volume of training records to confidently determine the length of time until each training certification is obtained. Thus, a simplified version of simply becoming "trained" is assumed (i.e. obtaining some set of training simultaneously) which is closer to a on-the-job training assumption.

Rank progression is another important aspect of the technician career path. It not only is a barometer for the efficacy of a technician, but is crucial for the organizational structure of the entire technician career to ensure a healthy, sustainable population is present. Currently, rank progression is handled in a simple manner: if a technician has the required training and years of service, they can potentially raise in rank assuming space exists for them in the organizational chart. This rank progression is treated as another fixed interval event; there are periodic checks for technicians to move up the rank chain. However, as mentioned earlier, career milestone is the more apt criteria for if a technician is considered fully functional. Thus, rank progression serves as more of an organizational chart restriction rather than a procedure to indicate functional technicians.

The last aspect of technician state change to consider is how career milestones are reached. This is understood to correlate strongly with technician functionality, as well as rank, and is thus used as the primary indicator of readiness level for the technician population. The specific criteria for journeyman is dependent on several human resource attributes, such as their occupation and unit, but achieving training is the primary bottleneck to reaching the journeyman stage. Once the required training is obtained, a progression to the next milestone is triggered with some optional delay.

3.2.4 Sink

All good things must come to an end, including tenure within the technician population. As mentioned in Section 3.2.2, the technician fate can be determined as they are created or during the simulation. In this section, both avenues will be considered.

First considered is the sink triggered upon entry, which unsurprisingly provides a computational advantage. As the technician population is quite large with thousands of individuals, having to test each individual for the sink procedure can be computationally expensive. This is particularly true if fine time resolution is needed, for example, to investigate seasonal effects. Once the cause of the technician's departure is determined: the technician is placed into the appropriate unit, the time for sink is recorded, and all subsequent functions involving the technician are removed.

While involving more computational overhead, checking if technicians fall out of the scope on-the-fly allows for a greater degree of flexibility when compared to the upfront method. By checking for attrition at fixed intervals, changes in the technician state can be incorporated into the sink probability calculation. This becomes centrally important if significant changes to the technician career need to be considered, such as the creation of new units, occupations, etc. While these can also be included in the upfront method, it requires knowing exactly when these changes will occur and so restricts some of the simulation flexibility. The actual procedure for this on-the-fly method is identical to the upfront procedure, except that it occurs for each technician at a fixed rate instead of just once. Just as before, if it is determined that an individual has left the fleet all human resource information is updated accordingly.

4 MODEL COMPARISON

Before embarking on more complex data sets and assumptions, the critical first test is ensuring the simulation agrees with the analytical expectation. To this end, artificial datasets are constructed for common input to generate results for comparison. Section 4.1 will highlight that the assumptions and procedures are identical in both approaches. Section 4.2 will turn to the population trends and compare them between the two models. Section 4.3 will discuss how the discrete-event simulation compares to real-world expectations.

The input data used for this comparison are taken from nominal estimates for a typical military technical professions. In particular, it is assumed that sources for all three career milestones exist (apprentice, journeymen, and managers) with four potential departments denoted by a numerical identifier. Four occupations are also possible, also with unique numerical identifiers, where the proportion for each is assumed to be identical. To simplify the results, each milestone is assumed to correspond to a single rank. Another simplifying measure is to assume that the number of incoming technicians is Poisson distributed, that is, the probability of some number of incident technicians is constant over time. Input years-to-arrival are stratified according to career milestone and time until on-the-job trained is applied for only the apprentice technicians.

For the sink procedure, four potential causes for technicians to leave the scope are considered. These correspond to broad categories of typical military technician career progress: retirement (or, entering non-effective strength), inter-fleet transfer, inter-occupational transfer, and promotion into administrative roles. To test the generality of the simulation approaches, the input distribution contains a bimodal feature. Such a distribution is not unusual for real-world human resource data, as these gates could correspond to deployment duration, retirement requirements, etc.

To obtain simulated population results, 1000 repetitions of the discrete-event simulation are run. Results are taken from an average across these repetitions, where the standard error of the mean for the various populations is less than 0.1%.

4.1 Basic Results

In order to rely on any more complex quantities yielded from the simulation, a consistency check on the basic results is needed. To this end, Figure 2 compares the time-to-arrival cumulative distribution functions for the apprentice and manager career milestones. To make these probability functions, the empirical data undergoes a smoothing procedure to allow continuous inferences to be made from the finite set of available data. As both models follow the same mathematical procedure to generate these distributions, it is expected that they should be identical. In this figure, the analytic result of Section 2 is given by the solid orange line and the discrete event simulation of Section 3 is given by the blue scatter points. Identical comparisons could be made for the time-to-departure and training times as well, since they all follow this same Gaussian-kernel smoothing procedure. As can be clearly seen, the distributions are consistent indicating that the discrete event simulation is working as expected.

The next verification stage is to compare derived distributions from the established years-to-arrival, years-to-departure, and training time distributions. This comparison will establish that the sampling procedure employed in the simulation yields the distributions expected analytically, given a sufficient number of samples. A prime candidate for this comparison is the dwell time distribution, described in Section 2.2. In essence, this quantity depends on sampling both the time-to-arrival and time-to-departure distributions, with the imposed restriction that the arrival time be less than the departure time, and then summing over all departure times. Analytically, this is equivalent to building the joint probability distribution from these two distributions and then integrating along the departure time.

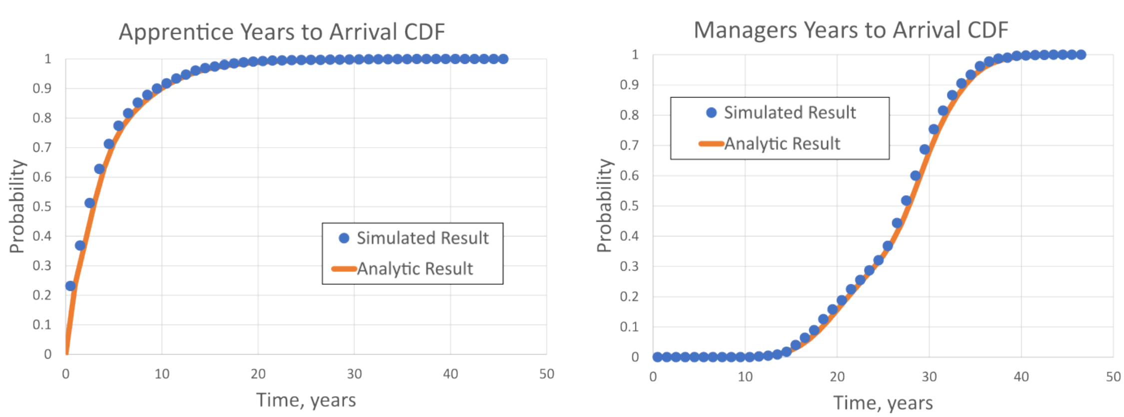

Figure 2: Comparison of the time-to-arrival cumulative distribution function (CDF) input into both analytic and the discrete event models and reconstructed from the simulation results. Apprentices are shown left and managers right. The input distribution is shown by a solid orange line and output results by the blue circular points.

The dwell time probability distribution functions obtained from the analytic and discrete-event simulation models are shown in Figure 3. Just like the previous comparison, two different professional milestones are compared, apprentice on the left and journeyman on the right, where the analytic result is the solid orange line and the discrete event result is the blue scatter points. While there is good agreement between the

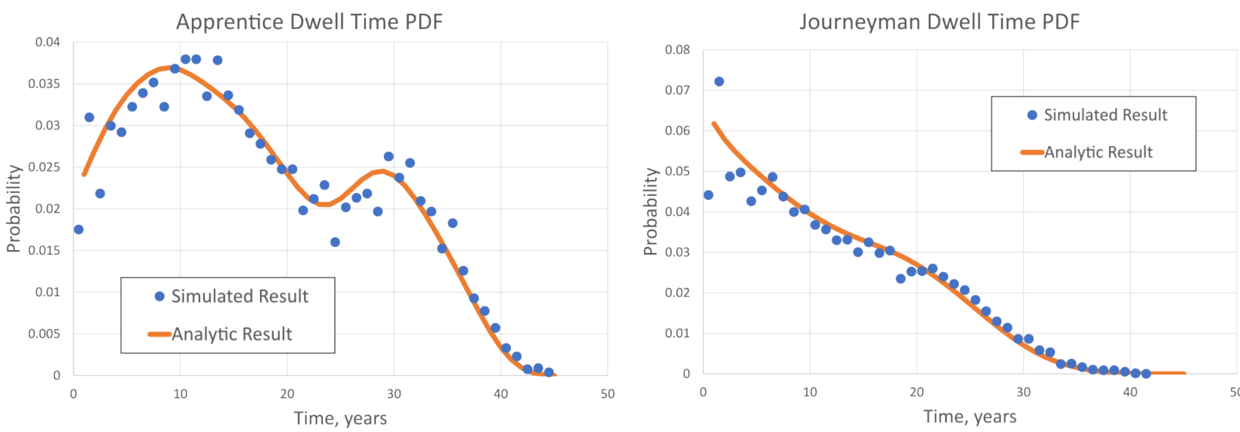

Figure 3: Comparison of the probability distribution function generated from analytic equations and the discrete event model for the dwell time. Apprentices are shown left and journeymen right. Analytic results are shown by a solid orange line and discrete event results by the blue circular points.

two approaches, for example the bimodal nature of the apprentice dwell time is well captured by both, the differences are significantly larger than the time-to-arrival distributions in Figure 2. This is due to the noise inherent in the discrete event simulation case: each point is obtained by the average count of instances in a certain interval and thus one must consider the standard error of the mean. Since the discrete event simulation was run for one hundred years, with an average of twenty technicians input each year, the total number of technicians in the simulation is approximately two thousand. Taking a best-case estimate of probability as 5%, which corresponds to approximately one hundred technicians, this yields a standard error of 10%. Keeping this value in mind, the discrepancy between the two approach falls within expectations.

4.2 Population Results

With the verification that the implementation of the simulation is consistent with expectations, it is now appropriate to transition to the central result of this work, the technician population trends. The result from the discrete event simulation is shown in Figure 4, where the population of a single unit is selected. Furthermore, each career milestone is denoted by a different colour and shape: small cyan circles for apprentice, large blue circles for journeyman, and pink triangles for manager. The overlaying solid lines correspond to Equations (4)-(6) after a curve fitting procedure is applied to the simulated population data.

The reason to apply a curve-fitting procedure at all is to obtain values for the coefficients of Equations (4)-(6) from the simulated results. Since the analytic model is able to obtain these coefficients directly from the input distributions (see Section 2), by obtaining estimates for these coefficients from simulated data, a direct comparison between the population results is possible. The curve-fitting procedure itself is a least-squares fitting algorithm (specifically, trust-region-reflective) which minimizes the difference between the input population data and the expected trends given from the equations. Initial estimates are given for each coefficient, obtained from the analytic model, and bounds are placed to obtain only physical results (e.g. ensuring non-negative parameters).

This fitting procedure generates the coefficients of technician loss per unit of time in Equations (4)-(6) allowing for direct comparison to the analytically determined values (i.e., through taking the mean of the corresponding distribution as explained in Section 2.2). These results are presented in Table 1. Uncertainties

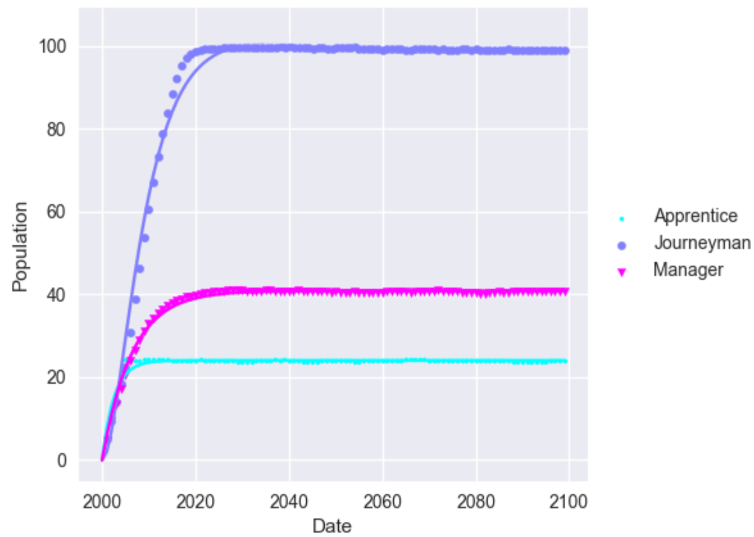

Figure 4: Technician population trends generated from the discrete event model for a single unit. Simulation results are shown by discrete points, solid lines indicate the equation fit results. Apprentices are cyan small circles, journeymen are blue large circles, and managers are pink triangles.

on the discrete-event simulation result are based on the standard deviation of the parameter between units. Two interesting features of this table merit some explanation. Firstly, since so few technicians leave the simulation as apprentices, or indeed stay an apprentice very long, there is a larger variation on their dwell time due to the lower number of available events. Secondly, the apprentice flow into the journeyman career milestone causes a larger spread in journeyman dwell times due to the distribution of training times.

Comparing these parameters of interest, one can see that there is strong agreement as expected from the results of Section 4.1. In fact, the only discrepancy is between the journeyman dwell time. As discussed, the uncertainty of this parameter is expected to be large for the discrete event modelling case. However, even by examining Figure 4, Equation 5 does not capture the journeyman population as well as the other two career milestones. This may be due to the attempt to capture the population behaviour with a single parameter (reducing the joint probability distribution to a single expectation value, see Section 2.2) which impacts the population dynamic before reaching the steady state.

4.3 Validation

Now that strong agreement between the two model approaches has been established, the discussion can turn towards comparing the model results to the real world. Unfortunately, this comparison cannot be completed as easily as it was for Section 4.2. First and foremost, security concerns restrict the access and distribution of population information. But beyond this, obtaining historical population trends for technicians is challenging due to: career restructuring, policy changes on hiring and retention, and changing operational

Table 1: Comparison of the population dynamic parameters between the discrete event model and the analytic model. Uncertainties correspond to the 95% confidence interval.

Parameter	Discrete-event Simulation Result	Analytic Result
Training Time (years)	1.7 ± 0.4	1.5
Apprentice Dwell Time (years)	19 ± 2.0	18.9
Journeyman Dwell Time (years)	14.2 ± 1.0	12.3
Manager Dwell Time (years)	6.0 ± 0.4	6.4

demands. This yields historical records without a steady state and with population levels disconnected from current parameters. However, the proportion of trained technicians (for a particular subset of career details) is a useful metric that has well-defined available data. Specifically, training is broken down into two tiers (Level A and Level B) where real-world records indicate that approximately 95% of technicians hold Level A and 80% hold Level B. Furthermore, to reach Level B, one must have obtained Level A and some predefined number of years of experience. Thus, by inputting the real-world time-to-training and time-to-departure distributions, these steady-state numbers should be achieved. These population levels are displayed in Figure 5, where these expected proportions are reached.

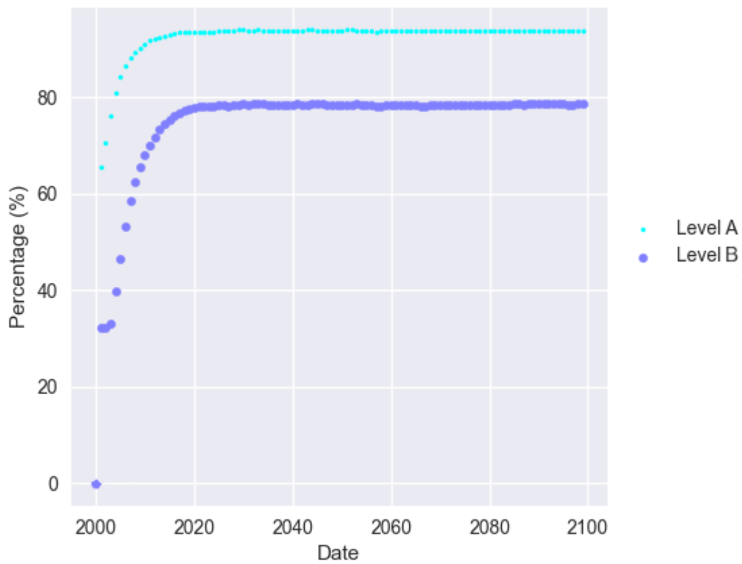

Figure 5: Proportion of technicians that obtained some degree of training. Level A training is cyan small circles, Level B training is blue large circles.

5 CONCLUSION

At this stage, it has been demonstrated that the tandem development of analytic and discrete event modelling tools allows for a unique method to cross-validate results. Additionally, each procedure is capable of taking the complexity of real-world data as input and generate meaningful results. The importance of this feature cannot be overstated, if this tool is to be used to inform strategy concerning technician population levels, it must be informed by real-world data specifically.

Turning towards the future, the next test for this simulation tool is a scenario associated with implications for technician readiness. For example, assuming a new operational requirement occurs *in media res* that creates a set of new occupations that need to filled as soon as possible, while ensuring existing occupations do not fall below some minimal level. Impose on this organizational chart restrictions, mentor-mentee dynamics, and other business rule complications and now the tool has a robustness to accurately capture the military technician profession. Currently, the discrete event simulation tool has the capacity to impose these business rule restrictions and is ready to work with real-world data.

ACKNOWLEDGMENTS

The authors would like to thank Dr. René Séguin and Dr. Sławomir Wesołkowski for their input and feedback on this work.

REFERENCES

Bastian, N. D., and A. O. Hall. 2020. "Military Workforce Planning and Manpower Modeling". In *Handbook of Military and Defense Operations Research* (1st ed.)., edited by N. M. Scala and J. P. Howard, 30. Boca Raton, FL, USA.

Boileau, M. L. A. 2012. "Workforce Modelling Tools used by the Canadian Forces". In *International Workshop on Applied Modelling and Simulation*, edited by A. G. Bruzzone, W. Buck, E. Cayirci, and F. Longo, 18–23. Rome, Italy.

Diener, R. 2018, 01. "A Solvable Model of Hierarchical Workforces Employed by the Canadian Armed Forces". *Military Operations Research* 23:47–57.

Gillespie, D. T. 1976. "A General Method for Numerically Simulating the Stochastic Time Evolution of Coupled Chemical Reactions". *Journal of Computational Physics* 22(4):403–434.

Henderson, J., and R. Bryce. 2019, 12. "Verification Methodology for Discrete Event Simulation Models of Personnel in the Canadian Armed Forces". In *Proceedings of the 2019 Winter Simulation Conference*, edited by N. Mustafee, K.-H. G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, 2479–2490. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Lahteenmaa-Swerdlyk., T., F. Bourque., and S. Wesolkowski.. 2023. "Workforce Modelling with Experience Accumulation". In *Proceedings of the 12th International Conference on Operations Research and Enterprise Systems - (ICORES)*, edited by F. Liberatore, S. Wesolkowski, and G. H. Parlier, 27–37. Lisbon, Portugal.

Novak, A., L. Tracey, V. Nguyen, M. Johnstone, V. Le, and D. Creighton. 2015. "Evaluation of Tender Solutions for Aviation Training using Discrete Event Simulation and Best Performance Criteria". In *Proceedings of the 2015 Winter Simulation Conference*, edited by L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, 2680–2691. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Okazawa, S. 2013. "A Discrete Event Simulation Environment Tailored to the Needs of Military Human Resources Management". In *Proceedings of the 2013 Winter Simulations Conference*, edited by R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, 2784–2795. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Schaffel, S., F.-A. Bourque, and S. Wesolkowski. 2021. "Modelling the Mentee-Mentor Population Dynamics: Continuous and Discrete Approaches". In *Proceedings of the 2021 Winter Simulation Conference*, edited by S. Kim, B. Feng, K. Smith, S. Masoud, Z. Zheng, C. Szabo, and M. Loper, 1–10. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Séguin, R. 2015. "Parsim, a Simulation Model of the Royal Canadian Air Force (RCAF) Pilot Occupation - an Assessment of the Pilot Occupation Sustainability Under High Student Production and Reduced Flying Rates". In *Proceedings of the 4th International Conference on Operatinal Reserach and Enterprise Systems (ICORES)*, edited by B. Vitoriano and G. H. Parlier, 51–62. Lisbon, Portugal.

Suvorova, S., A. Novak, B. Moran, and T. Caelli. 2019. "The use of Markov Decision Processes for Australian Naval Aviation Training Schedules". *Military Operations Research* 24(2):31–46.

Tanner, M. A., and W. H. Wong. 1983. "The Estimation of the Hazard Function from Randomly Censored Data by the Kernel Method". *The Annals of Statistics* 11(3):989 – 993.

Vincent, E., and S. Okazawa. 2019. "Determining Equilibrium Staffing Flows in the Canadian Department of National Defence Public Servant Workforce". In *Proceedings of the 8th International Conference on Operations Research and Enterprise Systems (ICORES)*, edited by G. H. Parlier, F. Liberatore, and M. Demange, 205–212. Prague, Czech Republic.

Zais, M., and D. Zhang. 2015, 11. "A Markov Chain Model of Military Personnel Dynamics". *International Journal of Production Research* 54:1–23.

AUTHOR BIOGRAPHIES

GEORGE RYAN AMBROSE is a Junior Defence Scientist at the Defence Research Development Canada – Centre for Operational Research and Analysis (DRDC-CORA). He is currently finishing a PhD in physics working with organic electronics at Dalhousie University. He holds a MSc in experimental nuclear physics from the University of Regina. His work with DRDC-CORA is focused on developing simulation tools and results to inform Royal Canadian Air Force decisions. Mr. Ambrose can be contacted via email: ryan.ambrose@forces.gc.ca.

FRANÇOIS-ALEX BOURQUE is a Defence Scientist at the Defence Research Development Canada – Centre for Operational Research and Analysis. He was a Senior Scientist at the NATO Centre for Maritime Research and Experimentation from 2015-2018 and a Visiting Scientist in 2013. Dr. Bourque holds a PhD in theoretical nuclear physics from McGill University. As a Defence Scientist, he has provided decision support to the Canadian Department of National Defence on a variety of projects. Additionally, he has explored the potential defence and security use of unmanned systems. Dr. Bourque can be contacted via email: alex.bourque@forces.gc.ca.