



Determining the Required Training Capacity Within a Military Establishment

Etienne Vincent¹ · Michelle Straver¹

Received: 22 August 2021 / Accepted: 31 March 2022 / Published online: 26 April 2022
© Crown 2022

Abstract

We address the problem of deciding how many positions to set aside for military recruits undergoing training. Within a cap on the total number of military members, we vary the ratio between positions allocated to the training pipeline versus those required in the trained effective establishment. This is done with the goal of determining the extent to which given ratios are sustainable. We use a Markovian model of the training pipeline, with parameters derived from historical personnel data. Through Monte Carlo simulation, we predict how often a given ratio allows the required trained force to be fully generated, as well as the surplus of trained personnel, it is expected to generate. We extend our previous work in this area by considering an alternative Human Resources policy that uncaps the training pipeline. Our modelling results have informed ongoing initiatives to optimize the force mix and structure of the Canadian Armed Forces.

Keywords Markov Manpower model · Monte Carlo simulation · Training pipeline · Military establishment · Personnel operations research · Workforce analytics

Introduction

This paper describes our solution to a challenging structural problem faced by the Canadian Armed Forces: determining the number of positions to set aside for Regular Force recruits undergoing training. Our solution is based on a stochastic simulation of the training pipeline. We used simulation to estimate how often given numbers of positions reserved for trainees would be sufficient to generate a desired trained force. Simulation also gave us the expected mean number of surplus trained personnel above the requirement. With our simulation results in hand, decision makers are designing a military establishment that balances the risk

of not meeting operational requirements against the costs of pipeline overcapacity.

The approach presented in this paper is simple and effective. It is presented in the hope that it will be appreciated by our peers as a practical application of military Operations Research. Our solution was developed towards designing a future fighting force, but can also be used to study the current force structure. Regularly revisiting the ratio of trainee-to-trained positions is necessary to maintain a sufficiently manned force.

This paper updates and extends our work published in the proceedings of the 10th International Conference on Operations Research and Enterprise Systems [1]. Our results were solidified through consideration of one additional year of personnel data, while the main extension of our study since the original paper is the consideration of a variation on the model allowing an uncapped training pipeline. In December 2021, the Armed Forces Executive Council, which is chaired by the Chief of the Defence Staff, endorsed a recommendation to increase the number of positions allocated to the training pipeline based on results obtained from the simulations presented in this paper.

This article is part of the topical collection “Operations Research and Enterprise Systems” guest edited by Federico Liberatore, Greg H. Parlier and Marc Demange.

✉ Etienne Vincent
Etienne.Vincent@Forces.gc.ca
Michelle Straver
Michelle.Straver@Forces.gc.ca

¹ Director General Military Personnel Research and Analysis, Department of National Defence, 101 Colonel By Drive, Ottawa, Canada

Force Mix and Structure Design

In 2017, the Government of Canada issued Strong Secure and Engaged, its latest national defence policy. Unlike previous policies, it explicitly mandated the capability to conduct a defined set of concurrent operations. These operations include responses to emerging events, as well as planned deployments, and range from time-limited to sustained commitments. To ensure that the Canadian Armed Forces have the right mix of military personnel to fulfil this mandate, an initiative was launched called Force Mix and Structure Design. This initiative is ongoing, and aims to design, from the ground up, the required military establishment.

The first phase of Force Mix and Structure Design aimed to determine the force employment requirement—the number of “boots on the ground” required for the conduct of operations. The second phase then looked at the rest of the force, consisting in those units tasked with generating the ready troops, as well as units with institutional roles. It is in the context of this second phase that the present study was conducted. This study aimed to determine the required size for the combined Basic Training List (recruits in their initial phases of training) and Subsidized University Training List (recruits receiving training in higher education institutions). For simplicity, we will refer to the positions reserved for members on these lists as the *training pipeline*. Regular Force recruits remain in this pipeline until they reach an Operationally Functional Point—the point where they are considered trained and can occupy trained effective positions.

Force Mix and Structure Design is not the first establishment review of the Canadian Armed Forces, and some previous initiatives also benefited from Stochastic Simulation studies. Notably, both Bender [2] and Couillard et al. [3] have used simulation to quantify the extent to which certain mixes of military personnel from various groups of occupations could be expected to meet stochastically determined operational requirements. Filinkov et al. [4] have similarly studied the Australian Army. The present study is most similar to that described by Straver and Christopher [5], who used stochastic simulation to determine the sustainable composition of the Regular Force, including the size of its training pipeline. However, that earlier study is now understood to have been underpinned by data that imperfectly inferred the trained or untrained status of Canadian Armed Forces personnel. Recent developments improving the tracking of personnel within staffing categories (especially the Training Pipeline and Trained Effective Strength) now allow a revision of that work. As such, our present model is a successor to Straver and Christopher’s that focuses exclusively on the training pipeline requirement, and that was designed around the improved staffing categories data.

Method

Many approaches are possible for modelling personnel systems. Wang [6] splits them into four categories: Markov Chain models, Computer Simulation models, Optimisation models and System Dynamics. Here, we take an approach that mixes the first two of these. We perform stochastic simulation on a discrete-time Markovian model.

Markov Manpower Models

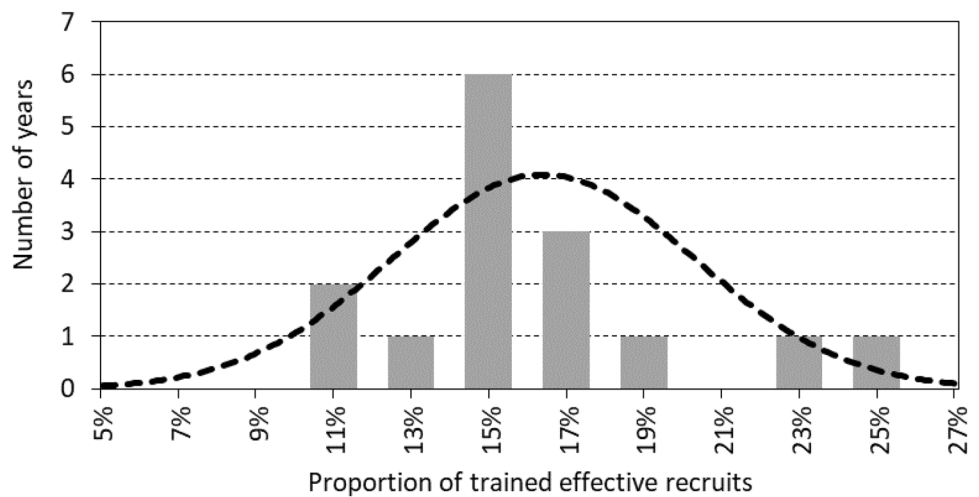
The earliest described application of Markov Chains to a personnel system was developed by Seal [7]. Guerry and De Feyter later put together a practical overview of *Markov Manpower models* [8]. They define them as having the following four properties:

- Markov Manpower Models are memory-less (the usual Markovian assumption);
- Their flow rates are time-independent;
- They are discrete-time, with fixed-increment intervals;
- Their stocks describe homogeneous populations.

For studies of the Canadian Armed Forces personnel system, Discrete Event Simulation has been the most often employed paradigm [9]. Discrete Event Simulation generally satisfies the Markov Manpower Model properties, with the exception that it is based on next-event time progression, rather than fixed time increments. Nevertheless, in the present work, we employ fixed annual time increments as prescribed for Markov Manpower Models.

The Markov Manpower Model property that is most difficult to satisfy is that requiring homogeneous populations. To achieve it, non-homogeneous populations must be broken down into homogeneously behaving sub-populations. This can be done using regression to identify the variables that distinguish the behaviours of some employees from the others. In our case, segmenting our population into homogeneously behaving sub-populations would pose a problem. The Regular Force employs members of various ranks, and who have completed different periods of service—factors that are closely tied to the propensity to release (to leave the Armed Forces). Also, members belong to several distinct occupations, and are recruited through different entry plans (some are unskilled, while others already have earned degrees in their area of specialty). Thus, they require training of substantially different duration. If the workforce was broken down into homogeneously behaving rank/years-of-service/occupation/entry-plan sub-populations, the small size of

Fig. 1 Histogram of the proportion of recruits who either do not require training or who complete their training within the year when they are hired, 2006–2021; a Gaussian fit is also shown as the dashed line



each would make it impossible to accurately estimate their parameters from available historical data.¹

We therefore built a model that diverges from Markov Manpower Models with respect to the homogeneous population assumption. Our stocks group cohorts of members who behave diversely, but who when taken collectively, have aggregate flow probabilities that are fairly consistent from year to year. Then, although the historically observed proportion of members from a stock who flow in a given direction cannot be interpreted as a probability applying to individuals, it can still be interpreted as the expected proportion of individuals who will transition. Instead of having transitions determined by simple transition probabilities, we use distributions that capture the range of observed outcomes over the years.

Transition Probabilities

Anderson and Goodman [10] demonstrate that the maximum likelihood estimator for Markov transition probabilities is given by the numbers of employees that underwent that transition divided by the total number of eligible employees in the relevant sub-periods (e.g. each year, when looking for an annual probability). Under the homogeneous population assumption, the probability for the total number of employees undergoing that transition is then binomial, and so, binomial distributions are commonly used in stochastic interpretations of Markov Manpower Models. This must, however, be reconsidered when the model stocks do not represent homogeneous populations.

¹ We had access to 15 years of staffing categories data, but even if earlier years were available, they may not have been as relevant since the behaviours of this military workforce changes along with society, employment markets and the evolving composition of the force.

In the absence of a priori knowledge of the shape of our flow probability distributions, we fit Gaussians. As an example, Fig. 1 shows a histogram of the annual proportion of recruits who either do not require training (e.g. re-hires or transfers from the Reserve Force) or who complete their training within the year when they were hired. A Gaussian fit is overlaid as a dashed line. Given that only 15 annual data points are available, we cannot determine that a Gaussian is truly the best option to represent this distribution, but we see in Fig. 1 that it is not incompatible with the observed data. The Gaussian is further justified by the fact that the range of possible values is continuous, and that we intuitively expected deviations from the mean to be cumulative and symmetric. In our simulation, we will draw from this distribution to obtain annual values. However, to avoid unrealistically extreme outcomes, we only draw within three standard deviations of the mean.

For some variables, however, a straightforward Gaussian fit proved clearly inadequate. Consider Fig. 2, which shows graduating rates from the training pipeline, among members who were recruited the previous year. We see a cluster of annual values centred near 45%, but also an outlier at 27%. This outlier corresponds to the fiscal year running from 1 April 2020 to 31 March 2021, and is the result of training delays caused by COVID-19 restrictions. Similar outliers for that fiscal year were also observed for the other graduation rates considered in our model.

When a Gaussian distribution is fit to all 15 available annual data points from 2006 to 2021, the results shown as a dotted line in Fig. 2 is a poor match to the data. We prefer to dismiss the outlier and fit with the remaining years, which yields the dashed line in Fig. 2, closely matching the available data other than the outlier. Although external shocks like the COVID-19 pandemic could occur in the future, and as such the deemed outlier is within the range of future possibilities, its removal is justified by the fact

Fig. 2 Histogram of the proportion who graduates within the year, among those who had first appeared in the training pipeline in the previous year, 2006–2021; a Gaussian fit is also shown as the dashed line

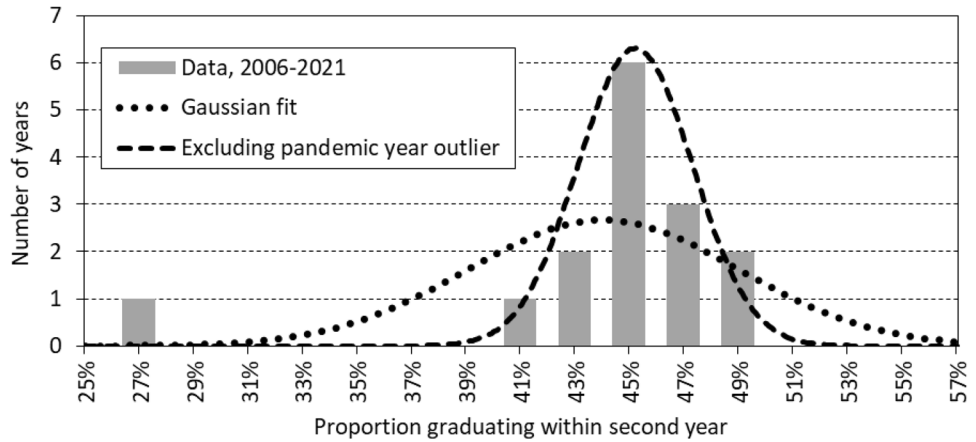
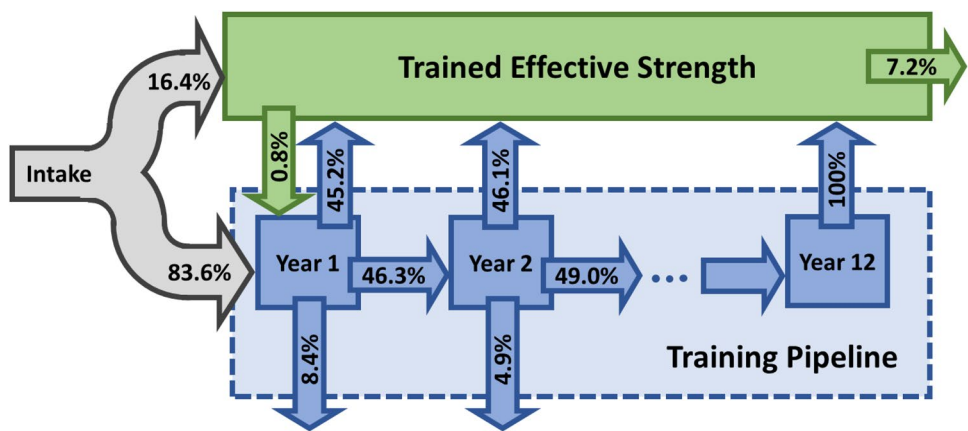


Fig. 3 Illustration of our model of the Canadian Armed Forces training pipeline (modified from [1])



that we are mostly interested in modelling ordinary years, as no force structure could be expected to be resilient to the extreme impacts of shocks of such magnitude.

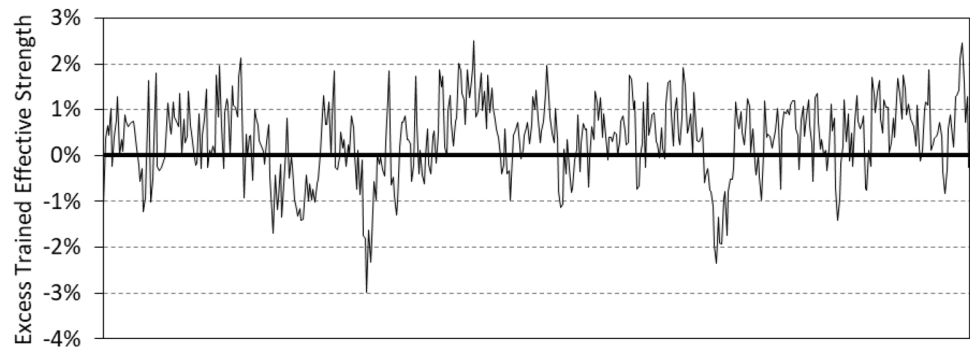
Previous efforts to model the Regular Force training pipeline, such as [2, 5] have sought to explicitly model the specific Human Resource processes applicable to military members. For example, graduations from the training pipeline were not modelled as a single phenomenon, but rather as separately considered graduations of Officers and Non-Commissioned Members, further split according to their entry plan (various categories of trained, semi-trained and untrained recruits). Instead, our approach considers only one flow between each model stock, consolidating a number of Human Resource processes. This greatly simplified the required historical data analysis to estimate model parameters, as it was not necessary to identify and separately consider the data associated with each category of recruit. We also expect that the resulting model will be more reliable, since by directly estimating aggregate flows from historical data, we avoid the error that otherwise results from ignoring covariance between sub-flows.

Training Pipeline Model

Figure 3 depicts our model. The percentages shown on arrows are the mean of the probability distributions derived for the modelled flows. The model has two parts: a stock for the Trained Effective Strength (TES), and a set of stocks for the training pipeline. The TES are those members who have completed their training up to the Operationally Functional Point, and are available to fill a trained effective position (e.g. they do not have a long-term illness or injury that prevents them from working). Members who change occupation usually require re-training—corresponding to the 0.8% arrow flowing towards the training pipeline. Others release or transfer to non-available categories (the 7.2% arrow).

The training pipeline is divided by cohort. From the first year when members appear in the pipeline, they can graduate (45.2%), remain for another year (46.3%), or, release or become unavailable to train due to illness or injury (8.4%). Eleven similar stocks are then included to represent subsequent years in the pipeline, with corresponding proportions for graduating, continuing training or exiting, except in the twelfth year when everyone left graduates. In our historical

Fig. 4 Five hundred years of simulation result showing excess TES—trained effective personnel generated above the set requirement, when the number of training pipeline positions was set to 1600 and the TES to 10,000



record, only one member was in the training pipeline for 12 consecutive years. No occupation requires this much training, but delays can result from pauses in training (such as for parental leave entitlements). Given the compounded graduation and exit rates, only a minute proportion of modelled recruits end up reaching the twelfth year stock.

Markov Chains can be handled deterministically or stochastically. Davies introduced a partially stochastic Markov model [11]. In that model, attrition is considered an uncontrollable flow, and treated as stochastic, whereas promotions are decided by management, and thus treated as deterministic. Our model does not consider promotions, but does set deterministically the total magnitude of intake (recruits and a small number of members returning from ill or injured status), and so is also partially stochastic. We set intake so as to always re-fill the training pipeline with a hard cap on total strength (i.e. the total Regular Force population including both TES and training pipeline). This deterministically set intake reflects reality, where the intake volume is set by recruiters according to their targets. Intake is also the only pull flow in our model. It is generated by vacancies in the destination (pull), rather than arising spontaneously from the source (push), as defined by Bartholomew et al. [12].

Although we set the magnitude of intake deterministically, we vary its proportion going to the TES versus the training pipeline stochastically. This is shown as the split grey arrow in Fig. 3. In the available historical record, 16.4% of intake reached the TES before the end of the year when they were recruited. This determination of TES versus training pipeline intake as a direct proportion of total intake resembles the proportionality constraint that Nilakantan and Raghavendra employ [13]. Their constraint requires that a fixed proportion of vacancies in a given grade be filled externally. Our model is, however, different in that our proportion varies according to a Gaussian distribution fit to the available historical data record.

Simulation

Our model was implemented in a spreadsheet. Each row of that spreadsheet records a year of simulation by tracking

the state of each model stock, the stocks being computed from the values for the previous year (the row above) with random draws from the transition probability distributions. The number of years simulated therefore corresponds to the number of rows in the spreadsheet. We found that the statistics that we tracked across this Monte Carlo simulation were very stable when simulating 200,000 years, requiring a few seconds of computation time for each scenario.

Simulation is normally described in terms of the number of replications—not years. However, note that as we are simulating a personnel system with a structure different from the current one, we cannot use the current state of the Canadian Armed Forces as a reasonable initial condition for replications. The initial conditions for a replication must be derived by simulating, starting from an arbitrary state until the system stabilizes. Once the system stabilizes, each subsequent year is effectively a replication based on a plausible initial state. This justifies the simulation of 200,000 years as our basis for deriving results.

Our goal was to assess the adequacy of various potential structures by considering their trainee-to-trained ratio (the ratio of the numbers of positions allocated to the training pipeline versus the TES). The numbers of positions were thus the model inputs that were varied across scenarios. Each simulated year, as the intake is set so as to re-fill the training pipeline (subject to the cap of total strength), recruits are added to the training pipeline who will graduate over the following years, and we may track to what extent the available training pipeline capacity manages to generate the required TES.

Each simulation begins with a fully manned TES, and a training pipeline full of fresh recruits. The first 200 years of simulation are then used to allow the distribution of the training pipeline across cohorts to thoroughly stabilize. Thereafter, the frequency of TES shortages and surpluses each year are expected to be representative of what could be achieved in steady state. To give an idea of how the simulation proceeded, Fig. 4 shows 500 years of simulation, with a ratio of 1600 training pipeline credits for every 10,000 required TES positions. The graph shows the resulting annual variation in excess TES. We see that the excess TES

tends to vary within a set range. Whenever the excess is too great, the cap on total strength means that the training pipeline cannot be filled to capacity, eventually leading to fewer graduations into the TES, and thereby eventually reducing the excess. Conversely, when there is a shortage, the training pipeline is filled to capacity, allowing for eventual growth of the TES. With the input values that generated Fig. 4, the TES meets (or exceeds) the requirement approximately 75% of the time.

Limitations of the Model

Before launching into modelling results, we will highlight three limitations of our approach that should be kept in mind. The first one has to do with using historical data to estimate the rates of flow out of the training pipeline. In the historical data, we observe delays in training that are the result of insufficient training capacity in some Canadian Armed Forces training institutions, or are the result of the sub-optimal allocation of training resources. This data describes the current training system, but not anticipated future improvements. As training institutions continually look for ways to improve throughput, and as the Force Mix and Structure Design initiative rebalances training resources where they are required, we should expect the training delays to be reduced. Historically observed training throughput may therefore over-estimate future training durations, and consequently, over-estimate the number of future required training pipeline positions.

A second limitation of our model is that it is based on annual-duration iterations (taken as of 31 March). We use these fiscal year end dates because it corresponds to the data readily available. As a result, the annual workforce snapshots used to feed the model are not taken at the point of the annual peak in required training pipeline capacity. That peak would likely be in summer, when more recruits begin their training. Other things being equal, we expect that slightly more training pipeline positions are actually required than determined by our simulation to account for the annual peak.

Finally, the third limitation is that our model only considers an overall TES target, ignoring its composition in terms of ranks and occupations. This masks gaps in trained personnel at certain ranks and in certain occupations. Certain positions can be filled from a range of different occupations, but others cannot, and difficulties in recruiting for certain specialties cannot be addressed by simply adding training positions. Sustainability is also required for the rank structure of each occupation. If the rank structure of an occupation is such that the junior ranks are insufficient to generate the required number of senior rank personnel, there will be gaps in the establishment at senior ranks that cannot be addressed by the size of the training pipeline alone. Filling such gaps would require increasing the number of positions

in intermediate ranks to eventually allow a sufficient pool of members promotable to the problematic senior rank. The rank pyramid of each occupation thus cannot have too-steep a slope, if it is to accommodate expected attrition and the time for members to sufficiently develop professionally before they can be promoted. The sustainability of occupations' rank structure is typically treated as a separate problem, such as by Diener [14].

Fully addressing these three limitations is not possible with the data at hand. For example, given that there are only so many members in each occupation, and that it is only relevant to look back so many years in the data record, accurately estimating occupation-specific training and attrition model parameters is not feasible. However, the rough impact of these limitations on our modelling results is understood by decision makers, and can therefore be considered in redesigning the Canadian Armed Forces' establishment. Overall, we believe that our model outputs slightly under-estimate true training pipeline requirements.

Results

Achieving the Required Trained Effective Strength

Our main reason for studying the required number of positions allocated to the training pipeline is to ensure that it is sufficient to generate the required number of trained personnel. Figure 5 was obtained by simulating with various training pipeline capacities for a fixed TES requirement of 10,000. It shows how often each resulting trainee-to-trained ratio is expected to leave a TES shortage. As the capacity of the training pipeline increases, the frequency of TES shortages diminishes. Simplistically, Canadian Armed Forces decision makers could look at this chart to pick the training pipeline capacity that results in the worst frequency of shortages that they would be willing to accept.

Surplus Trained Personnel

However, there are other considerations in selecting a training pipeline capacity. Chief among them is the Human Resources cost associated with a larger training pipeline. The principal component of that cost is the cost of maintaining the training capacity, and should be roughly proportional to the size of the training pipeline. However, another component of that cost results from the surplus trained personnel likely to be generated by a larger training pipeline. These surplus trained personnel are in excess of the number needed to fill all trained effective positions. They will still contribute to the Canadian Armed Forces' mission, but beyond the capability for which the establishment was designed. As such, the Human Resources cost involved in their upkeep is

Fig. 5 Fraction of simulated years when a TES shortage is observed, as a function of the number of training pipeline credits, when 10,000 trained effective positions require filling

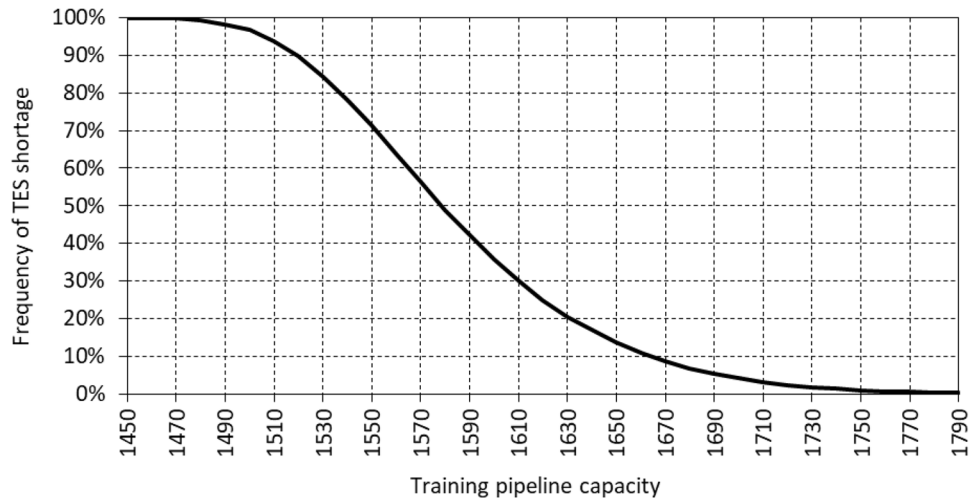
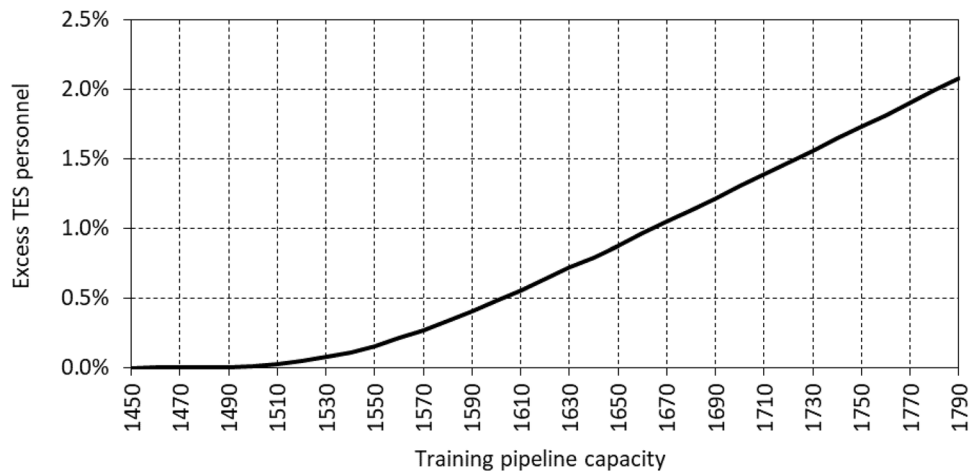


Fig. 6 Mean TES above the requirement in those simulated years when the TES requirement is exceeded, as a function of the number of training pipeline credits, when 10,000 trained effective positions require filling



in excess of that for the required capability. Figure 6 shows how the mean TES surplus increases with pipeline capacity. When selecting a training pipeline capacity, this can be a secondary consideration to the pipeline’s ability to fulfil its requirement shown in Fig. 5.

Modelling Error

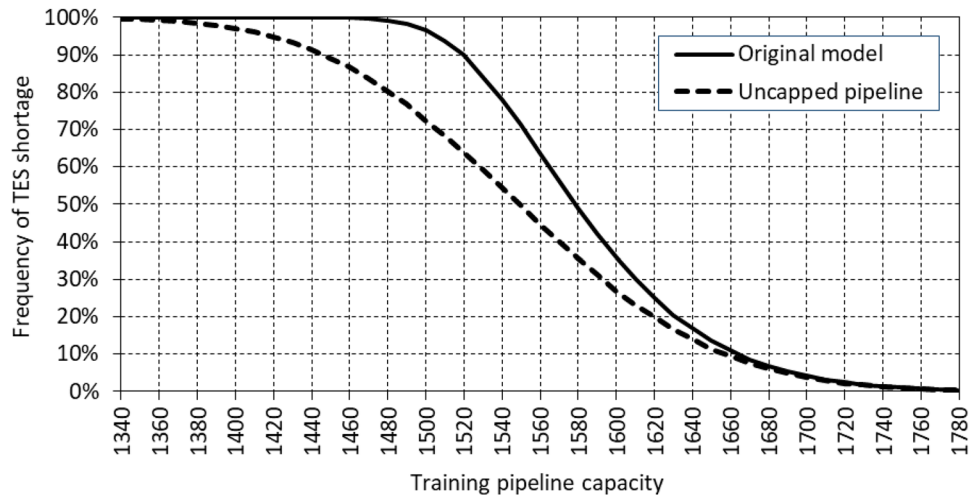
We will now briefly discuss modelling error. Guerry and De Feyter distinguish three types of error that apply to Markov Manpower models: statistical, estimation and specification error [8]. Statistical error results from the stochastic nature of models and may be reduced by increasing the number of simulation iterations. In our case, we expect this error to be small, given that the 200,000 simulated years per scenario resulted in fairly stable statistics for the frequency with which scenarios achieve a required TES and the mean excess generated. Given that the 200,000 simulated years only required seconds of computational time, this could be determined empirically, by concluding that the output

statistics minimally varied over several successive independent runs, without having to resort to statistically deriving a required minimum number of simulated years, which would certainly have been far lower than 200,000.

Estimation error is related to the accuracy to which parameters are estimated. In our case, we relied on 15 years of historical personnel data. Older data, or data from other organisations could also have been considered, but would have been less representative of the current system, and therefore unlikely to yield better estimates for our model parameters. Therefore, although our estimation error is likely substantial, there is no viable path to reducing it.

Lastly, specification error results from the model being an imprecise representation of reality. Reducing specification error eventually involves increasing the complexity of the model with additional parameters. For example, we could have introduced alternative training paths within the pipeline for semi-trained recruits versus completely untrained ones, or we could have acknowledged different training durations for officers and non-commissioned members, but each such

Fig. 7 Fraction of simulated years when a TES shortage is observed, as a function of the number of training pipeline credits, when 10,000 trained effective positions require filling, using the original model and the modified model with an uncapped training pipeline



parallel pipeline streams would also result in doubling the number of pipeline parameters and in introducing another parameter to determine the split between streams. The resulting additional parameters would reduce specification error, but involve a trade-off with estimation error, since estimation error increases with the number of parameters that must be estimated. We have sought to achieve an appropriate balance between the two errors in our model design. We believe that further increasing the number of parameters to reduce specification error would necessitate further splitting the historical dataset, leaving less data to estimate each parameter, leading to worse estimates of these parameters, and thereby increasing overall modelling error.

Modified Model with an Uncapped Training Pipeline

To derive the results that we presented above and in [1], our model determined the number of recruits hired each year as that which filled the training pipeline (unless the cap on total strength was the limiting factor). This is how things should happen in theory—trainees are only accepted when training pipeline credits are available for them. However, in practice, the Canadian Armed Forces currently over-fill their training pipeline, as long as the hard cap on total strength is not exceeded. This has the advantage of allowing TES to recover more quickly by allowing slots allotted to it to be used for recruits in training when no trained members are available to fill them.

We have thus decided to explore a modification to our model that uncaps the training pipeline. In this modified model, the number of recruits hired each year is simply set to fill any space available under the total strength cap, no matter the number of training pipeline credits available. This conforms better to how the Canadian Armed Forces

have been managed in recent years, and as we will see, has advantageous consequences.

Results with Uncapped Training Pipeline

Figure 7 presents the TES shortage frequencies observed from the modified model as a function of the size of the training pipeline. Even if the number of recruits is not directly limited by training pipeline capacity, the number of positions allotted to the training pipeline contributes to the cap in total strength, which is what still limits the number of accepted recruits. Thus, the frequency of TES shortages still changes as a function of the number of positions allocated to the training pipeline. The full line in Fig. 7 is the same as that shown in Fig. 5 from the original model. We see that the graph for the modified model, shown as a dashed line, achieves fewer TES shortages for the same pipeline size, and is thus advantageous from the perspective of this performance metric.

Figure 8 shows the mean TES surplus obtained with the modified model, as a function of pipeline size. The surplus observed from the original model is shown as the full line (repeated from Fig. 6), while that from the modified model is shown as a dashed line. We see that the TES excesses seen when the training pipeline is uncapped are slightly larger than with the original model for a given pipeline size. However, this is more than compensated by the lower frequency of shortages seen in Fig. 7. For example, to fill the TES 50% of the time, the original model requires approximately 1580 training pipeline credits for each 10,000 TES positions, while the mean surplus when the TES requirement is exceeded is of 0.34%. However, to achieve the same 50% fill, the modified model requires only a 1550 pipeline capacity, while the resulting mean surplus is of only 0.28%. The modified model with uncapped pipeline is thus superior with respect to both metrics simultaneously.

Fig. 8 Mean TES above the requirement in those simulated years when the TES requirement is exceeded, as a function of the number of training pipeline credits, when 10,000 trained effective positions require filling, using the original model and the modified model with an uncapped training pipeline

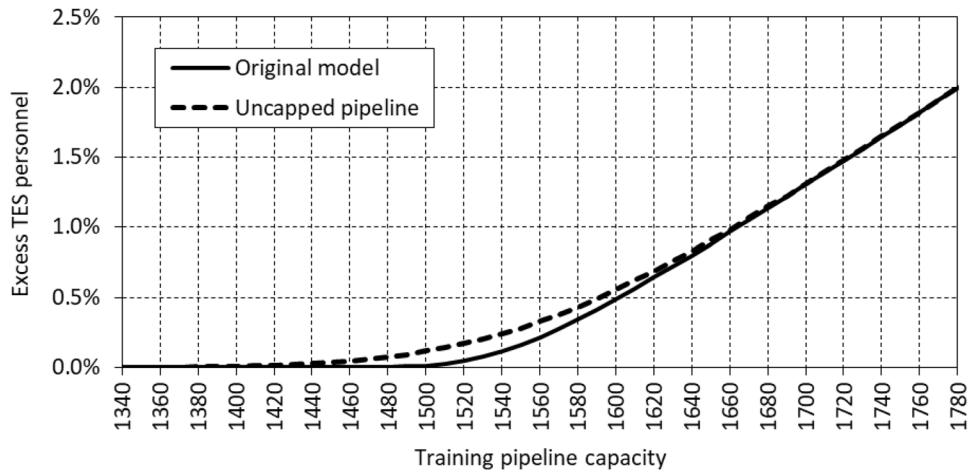
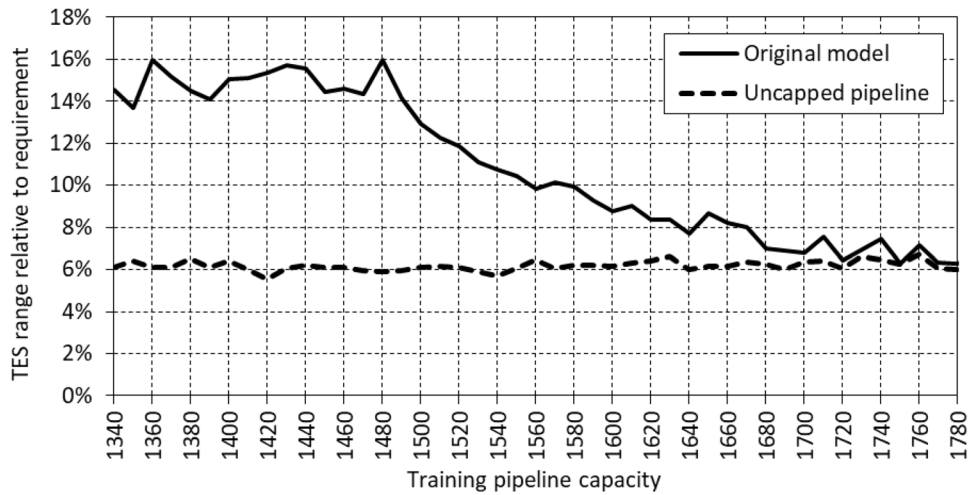


Fig. 9 Range between the maximum and minimum TES observed over 200,000 modelled years, expressed as a percentage of the required TES, graphed as a function of the number of training pipeline credits, given a 10,000 TES requirement



Another advantage of uncapping the training pipeline is that it reduces the range of outcomes in terms of the TES generated. Figure 9 illustrates this by showing, for both models, the range of TES achieved across the 200,000 years simulated for each scenario. With the original model, we see that this range is as high as 16% of the required TES (1600 members per 10,000 required), but it never exceeds 7% of the required TES for the modified model with an uncapped pipeline. The range is lower for the original model when the size of the pipeline is larger, and recruitment is therefore more often limited by the cap on total strength, but increases as the pipeline becomes smaller and the cap on the pipeline itself comes into play more often. The narrower range of outcomes achieved with an uncapped pipeline is an advantage, because it requires less flexibility in the establishment, to accept large fluctuations in the number of trained effective personnel.

Study Recommendations

Given the advantages of allowing an uncapped training pipeline highlighted in the previous subsection, this is the Human Resources policy that we recommend. It is also how the Canadian Armed Forces are currently managed in practice. We will now look at how decision makers can use the results generated by this study to select a preferred trainee-to-trained ratio for the Canadian Armed Forces establishment. As of the time of writing, a final decision had not yet been made by the chain of command, but the options contemplated were being informed by our model.

A first realization is that the cost of achieving a very low expected frequency of TES shortages is high. As such, military decision makers requested that we explore options that would result in more frequent shortages, but while ensuring that those shortages are not too great. Figure 10 augments

Fig. 10 Fraction of simulated years when various magnitudes of TES shortages are observed, as a function of the number of training pipeline credits, when 10,000 trained effective positions require filling, with an uncapped training pipeline

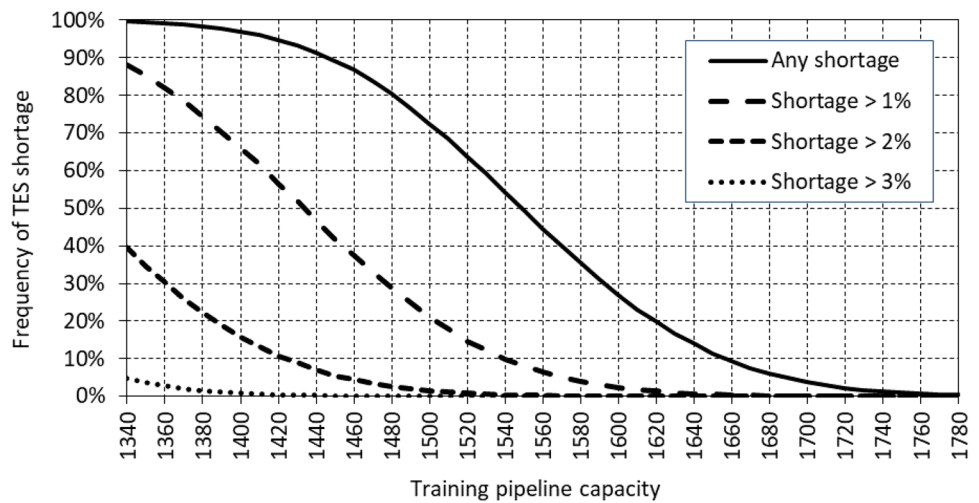


Table 1 Potential trainee-to-trained ratio choices, with corresponding metrics obtained from simulation

Trainee-to-trained ratio	1549:10,000	1570:10,000	1592:10,000	1620:10,000	1657:10,000
Frequency of TES shortage	50%	40%	30%	20%	10%
Frequency of TES shortage above 1%	8.3%	5.1%	3.0%	1.4%	0.39%
Frequency of TES shortage above 2%	0.29%	0.12%	0.054%	0.018%	0.0045%
Excess TES personnel	0.28%	0.38%	0.51%	0.69%	0.96%

the frequency of shortages that had previously been shown in Fig. 7 with curves showing the frequencies of shortages that are greater than 1%, 2% and 3% of the TES. We see that even when TES shortages are frequent, shortages greater than 1% (100 missing members among 10,000 positions) are much less frequent. This may provide grounds for justifying a smaller training pipeline that will result in more frequent shortages, but of a manageable magnitude.

Table 1 gathers modelling results for selected points shown in Figs. 8 and 10, to illustrate how decision makers could go about selecting a preferred trainee-to-trained ratio.

The minimum starting point for decision makers might be a value that results in TES shortages 50% of the time (considering more frequent shortages is unlikely to be considered, because there is little point in assigning responsibilities within the establishment to positions that are likely to go unfilled more than half the time). To expect shortages no more than 50% of the time, 1549 training pipeline credits are required for every 10,000 required TES. Table 1 then tells us that shortages greater than 1% (100 trained members for every 10,000 positions) would occur 8.3% of the time. This might be seen as leaving too big a gap in the force structure too often, so decision makers could proceed to look at greater pipeline capacities. With 1592 training credits for every 10,000 required TES, shortages greater than 1% would now be seen only 3% of the time, which might be seen as an acceptable downside. The decision makers could also see

that shortages greater than 2% would be very infrequent, and could observe that the cost in terms of surplus TES would average 0.51% (or 51 members for every 10,000 positions). The final decision will reside with the Canadian Armed Forces chain of command, but this discussion illustrates the thought process that follows from this study.

Conclusion

This paper described an application of Monte Carlo simulation on a model designed similarly to Markov Manpower Models, but where the transitions were drawn from Gaussian distributions. The model was employed to estimate the ability of the Canadian Armed Forces’ Regular Force training pipeline to train enough military members to meet requirements. The results of this simulation are being provided to military staff tasked with redefining the organization’s force structure. An advantage of the method employed in deriving our results is that it was straightforward and easy to follow, and thus readily trusted and appreciated by study clients.

Although the Regular Force training pipeline has been undersized in recent years, efforts are being made to correct this situation. Our model has supported those continuing efforts. This paper also improves on our previous work by introducing the uncapped pipeline variation on our original model, which suggests a Human Resources policy that

is advantageous in terms of reducing the pipeline capacity requirement, the cost of surpluses, and the expected range of achieved outcomes.

Given a hard cap on total strength and the need to improve sustainability, some trained effective positions need to be converted into training capacity, but this comes at the expense of nameplate military capability. Alternatively, the current TES requirement can be maintained if the total strength cap is increased to allow more training pipeline capacity. As of the time of writing, both approaches are being developed into options to be presented to Canadian Armed Forces senior decision makers.

We should also note that a number of parallel mechanisms for the continued optimization of the personnel system continue to be investigated. For instance, incentives and other measures can be developed to increase retention, leading to lower requirements for training new recruits within the training pipeline. For example, Ueno et al. [15] have investigated the use of machine learning to better target women recruits, but similar approaches are also considered to target recruits who will be associated with higher retention rates. Another line of investigation seeks to optimize the training pipeline throughput through the identification and rectification of bottlenecks (e.g. by Lalbakhsh et al. [16] with respect to aircrew training).

Our model was developed to support the currently ongoing Force Mix and Structure Design initiative. In that context, we are informing decisions on the future force structure. Subject to the limitations presented in this paper, our simulation results have allowed initiative staff to predict the sustainability of considered force structures in terms of their ability to achieve sufficient TES.

Acknowledgements The authors would like to thank our colleague, Drago Calitoui, for helpful modelling advice.

Funding The authors are employees of the Canadian Department of National Defence, which requested and funded this work.

Availability of data and material This work is derived entirely from Human Resources Management System data, which is protected and may not be disclosed, as per the Privacy Act.

Code availability The code involved in this work was developed by Government of Canada employees, and as such, belongs to Her Majesty the Queen in Right of Canada.

Declarations

Conflicts of interest/Competing interests No conflicts of interest or competing interests are involved in the publication of this work.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License, which permits any non-commercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative

Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc/4.0/>.

References

1. Vincent E, Straver M. Determining the required size of a military training pipeline. In: International Conference on operations research and enterprise systems. 2021; p. 358–65.
2. Bender P. Towards a sustainable CF: a risk analysis model (DRDC CORA TM 2005–10). Ottawa: Defence Research and Development Canada Centre for Operational Research and Analysis Technical Memorandum; 2005.
3. Couillard M, Arseneau L, Eisler C, Taylor B. Force mix analysis in the context of the Canadian armed forces. In: International Symposium on military operational research. 2015.
4. Filinkov A, Richmond M, Nicholson R, Alshansky M, Stewien J. Modelling personnel sustainability: a tool for military force structure analysis. *J Oper Res Soc.* 2011;62(8):1485–97.
5. Straver M, Christopher G. Estimating the sustainable Canadian Armed Forces Trained Effective Establishment (DRDC-RDDC-2015-R173). Ottawa: Defence Research and Development Canada; 2015.
6. Wang J. A review of operations research applications in workforce planning and potential modelling of military training (DSTO-TR-1688). Canberra: Defence Science and Technology Group; 2005.
7. Seal HS. The mathematics of a population composed of K strata each recruited from the stratum below and supported at the lowest level by a uniform annual number of entrants. *Biometrika.* 1945;33(3):226–30.
8. Guerry M-A, De Feyter T. Markovian approaches in modeling workforce systems. *J Curr Issues Finance Bus Econ.* 2009;2(4):1–20.
9. Okazawa S. A discrete event simulation environment tailored to the needs of military human resource management. In: Winter Simulation Conference. 2013; p. 2784–95.
10. Anderson TW, Goodman LA. Statistical inferences about Markov chains. *Ann Math Stat.* 1957;28(1):89–110.
11. Davies GS. Control of Grade size in partially stochastic Markov manpower model. *J Appl Probab.* 1982;19:439–43.
12. Bartholomew DJ, Forbes AF, McClean SI. Statistical techniques for manpower planning. Chichester: Wiley; 1991.
13. Nilakantan K, Raghavendra BG. Control aspects in: proportionality markov manpower systems. *Appl Math Model.* 2005;29:85–116.
14. Diener R. A solvable model of hierarchical workforces employed by the Canadian armed forces. *Mil Oper Res J.* 2018;23(3):47–58.
15. Ueno R, Bryce RM, Calitoui D. Ranking clusters of postal codes to improve recruitment in the Canadian Armed Forces. In: International Conference on machine learning and applications. 2019; p. 1192–7.
16. Lalbakhsh P, Mak-Hau V, Séguin R, Nguyen V, Novak A. Capacity analysis for aircrew training schools—estimating optimal manpower flows under time varying policy and resources constraints. In: Winter Simulation Conference. 2018; p. 2285–96.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.