Modelling Defence Planning as a Sequential Decision Problem

Carolyn Chen[a], Mark Rempel[b] and Kendall Wheaton
Centre for Operational Research and Analysis, Defence Research and Development Canada,
60 Moodie Dr., Ottawa, Canada
{carolyn.chen, mark.rempel, kendall.wheaton}@forces.gc.ca

Keywords: Defence Planning, Sequential Decision Problem, Mathematical Modelling, Military.

Abstract: Defence planning in a nation's defence organization is a complex process that requires considering hundreds of projects and billions of taxpayer dollars. While a variety of methods, such as integer programming and genetic algorithms, are used in practice to help decision makers select in which projects to invest, their application tends to not account for the fact that these decisions are made sequentially and under uncertainty. In this paper, we present the first steps towards developing a sequential decision model of the Canadian Department of National Defence's multi-gate Project Approval Process that addressed both issues. Our contributions are twofold. First, using the universal modelling framework for sequential decisions we present a mathematical model that accounts for the sequential nature of project selection, arrival of new projects over time, and uncertainty in future budgets. In addition, we extend this model to account for the uncertainty in how a project's cost changes over time when its selection is delayed. Second, we demonstrate how these models may be used to compare the effectiveness of three project selection decision policies, namely a ranked list approach, a knapsack approach, and a knapsack approach that reserves a contingency fund for future projects.

1 INTRODUCTION

Defence planning is a critical activity that aims to help nations achieve both their short- and long-term defence and security objectives. However, selecting the right capabilities, and thus projects, in which to invest is not always a straightforward process. This is due to a variety of factors—financial constraints, regulatory constraints, multiple conflicting objectives, interdependence between projects, cost uncertainty, etc.—that often introduce a high-degree of complexity to the selection process. In addition, the planning process often must consider various project types simultaneously, including those focused on information technology, equipment, infrastructure, as well as support contracts (Rempel and Young, 2017). As a result, the project delivery timelines considered may vary significantly. For example, the Canadian Defence Investment Plan 2018 included projects with both near- and long-term delivery dates ranging from 2022 to 2038 (Government of Canada, 2019).

In order to accommodate the wide range of timelines, planning processes "must consider a temporal dimension including immediate activities to possible demands a few decades into the future", and as a result occur in "an inherently uncertain and often unstable external environment where defence organisations are required to reorganise for, and respond to, unpredicted turns of events" (Filinkov and Dortmans, 2014, p. 76). With this in mind, many defence planning problems may be aptly described as *sequential decision making problems under uncertainty*. Furthermore, when decision makers are seeking a portfolio, the defence planning problem may be described as a *multi-period portfolio optimization problem* (Salo et al., 2024) and modelled as a knapsack problem (Locatelli, 2023).

The classic approach to solve a sequential decision problem is to model it as a Markov Decision Process (MDP) (Puterman, 2005) and use Dynamic Programming (DP) (Bellman, 1957) to find a decision policy—"*a rule (or function) that determines a decision given the information available*" (Powell, 2011, p. 221)—that makes decisions which result in the system performing optimally with respect to a given criterion or objective. However, in many real-world situations this is not feasible due to the curse of dimensionality (Kuo and Sloan, 2005) and the curse of modelling (Bertsekas and Tsitsiklis, 1996).

Given these limitations, alternative approaches such as linear programming (and its variants—integer

[a] https://orcid.org/0000-0003-0209-1666
[b] https://orcid.org/0000-0002-6248-1722

Chen, C., Rempel, M. and Wheaton, K.
Modelling Defence Planning as a Sequential Decision Problem.
DOI: 10.5220/0013248200003893

In *Proceedings of the 14th International Conference on Operations Research and Enterprise Systems (ICORES 2025)*, pages 156-164
ISBN: 978-989-758-732-0; ISSN: 2184-4372

programming, mixed-integer programming, etc.), genetic algorithms, tabu search, and other novel heuristics have been used to provide decision support. This is evidenced by a recent review of 54 application-based articles that described the use of such approaches for portfolio optimization in a defence context (Harrison et al., 2020, see Table 1). Further analysis of these articles reveals that while many address sources of uncertainty in their respective problems, only five articles (Crawford et al., 2003; Tsaganea, 2005; Fisher et al., 2015; Shafi et al., 2017; Moallemi et al., 2018) account for both uncertainty and new arriving projects which are selected for investment over time—thus, modelling their respective problems as sequential decision problems. Inspired by Fisher et al. (2015), who studied a capital investment planning problem in the context of the Royal Canadian Navy, this paper presents the first steps towards developing a sequential decision model of the Canadian Department of National Defence (DND)'s multi-gate Project Approval Process (PAP).

This paper's main contributions are twofold. First, using the universal modelling framework for sequential decisions (Powell, 2022), two sequential decision models are presented: (i) a model that accounts for the sequential nature of project selection, the arrival of new projects over time, and uncertainty in future budgets; and (ii) an extension of the first model which incorporates the uncertainty in how a project's cost changes over time when its selection is delayed. Both the uncertainty in the future years' budgets and changes in a project's costs were not represented in Fisher et al. (2015), and thus this paper extends the existing research in terms of modelling. Second, we demonstrate how these models may be used to compare the effectiveness of three project selection decision policies, namely a ranked list approach, a knapsack approach, and a knapsack approach that reserves a contingency fund for new projects arriving in future years.

The remainder of this paper is organized as follows. Section 2 presents relevant background information. Section 3 describes the defence planning scenario considered in this paper. The mathematical model that formulates the scenario as a sequential decision problem and its extension are given in Section 4. The results from a series of computational experiments that demonstrate how these models may be used to provide decision support are given in Section 5. Lastly, a conclusion is provided in Section 6.

2 BACKGROUND

Within the Canadian DND/Canadian Armed Forces (CAF), hereafter referred to as DND/CAF, strategic planning is performed by the Chief of Force Development (CFD) on behalf of the Vice Chief of Defence Staff (VCDS) (Government of Canada, 2024b). This process requires the preparation of long-range plans that identify future requirements for defence. Project and force development staff across the DND/CAF then prepare project proposals to address these future requirements. The major capital project proposals (>$10 million CAD) are evaluated through the Capital Investment Program Plan Review (CIPPR) process for continuation through the PAP. This is described in further detail in subsection 2.1.

The time required to bring a proposal through the strategic planning process can vary significantly due to both internal and external reasons. The future requirements upon which proposals are based will evolve over time and a requirement may change significantly while a project proposal is being developed or during the PAP. These uncertainties are discussed in subsection 2.2.

2.1 Project Approval Process

The steps a major capital project must follow throughout its life-cycle within the DND/CAF are depicted in Figure 1.

A project starts with a proposal submitted to the annual CIPPR process by its internal sponsoring organization. The sponsoring organization (henceforth referred to as *sponsors*) can be military (e.g., Army, Navy, Air Force) or a supporting organization (e.g., Materiel, Defence Research and Development, Infrastructure and Environment). In this intake process, projects identify the capability gaps they address, provide cost estimates and tentative project timelines, and the project is evaluated by subject matter experts to determine its overall value (i.e., benefit) to the DND/CAF. At the end of the CIPPR process, a decision is made whether to allow the project to continue through to the first phase of the PAP (Identification), to place the project on a waiting list, or to remove the project from consideration. Decision support tools, such as the Visual Investment Optimization and Revision (VIPOR) or Strategic Portfolio Analyzer with Re-configurable Components (SPARC), are used to support this decision process (Rempel and Young, 2017; Chen and Wheaton, 2024). As the project continues through the PAP, the project is assessed at the end of each phase and a decision is made on the project's continuation by a designated govern-

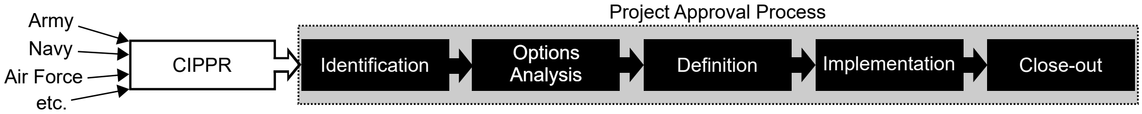

Figure 1: Project pathway from inception within the sponsoring organization, to the intake process (CIPPR), and through the five phases of the PAP.

ing board (Balkaran, 2021). This creates a multi-gate, sequential decision process. The project is complete once it finishes the Close-out phase.

2.2 Uncertainties in Defence Planning

It is important to understand the uncertainties associated with long-term planning in defence as they can have a large impact on the delivery of new capabilities. Uncertainties are caused by many factors—the evolving nature of warfare as new technologies emerge and new tactics are devised (Roncolato, 2022), the plans and policies of Canada's close Allies and the North Atlantic Treaty Organization (NATO) (North Atlantic Treaty Organization, 2022, 2024), and the strategic threats to the peace, security and sovereignty of the nation (Government of Canada, 2024a)—often causing projects to require more time and cost than originally estimated. Inflation alone will change a project's cost if planning takes years longer. Changes in the project scope can also happen, affecting the cost. Finally, it is often the case that the cost estimates for projects tend to be optimistic (U.S. Government Accountability Office, 2020, Ch. 1).

Consider the example of the Maritime Helicopter project. Planning for a new maritime helicopter began in the 1980's, a contract was signed in 1992, and the project was subsequently cancelled in 1993 after a change in the federal government (Rossignol, 1998). New planning in the 1990's led to a contract being signed in 2004. The delivery of the new helicopter commenced in 2015 (Government of Canada, 2022). This shows how political uncertainties caused a significant increase in a project's timeline.

The Canadian Surface Combatant (CSC) project is an example of cost uncertainty. In a recent study (Office of the Parliamentary Budget Officer, 2021), the estimated cost for the CSC was reported to be $26.2 billion in 2008, almost $62 billion in 2017, and to $69.8 billion in 2019. After awarding the contract in 2019, this 2021 study estimated the cost at $77.3 billion, a 295% increase from the initial amount.

3 PROBLEM DEFINITION

The defence planning scenario in this paper is based on Fisher et al. (2015) and the CIPPR process described in subsection 2.1. While the PAP includes multiple decisions through which a project must proceed, by focusing on the CIPPR process, this paper limits the defence planning scenario to a single stage process in which projects are either approved or not approved.

For this scenario, given Canada's defence policy (Government of Canada, 2024a) and recent defence capability plans (Government of Canada, 2019), suppose that at the start of a given fiscal year a set of candidate projects are put forth by sponsors. Each project is seeking multi-year funding from the available defence budget. To help decision makers select which candidate projects to fund, information on each project is collected from its sponsor: its value to defence, and its purchase cost which is distributed over a number of fiscal years. In addition, the available budget for 20 years is provided by the Assistant Deputy Minister (Finance) organization.

With this information, a decision policy (see Section 4) is used by decision makers to select those candidate projects that will be funded, and those that will not. Following this decision, the funding associated with the selected projects is removed from the available budget in both the current and respective future fiscal years, and projects that were not selected are added to the set of candidate projects for the next fiscal year—albeit with their value reduced due to their implementation being delayed. This reduction in value may be caused by a variety of reasons, including improvement of adversarial capabilities rendering the project less effective, a shift in defence policy rendering the project less important, or investment in similar capabilities by other sponsors or Allies reducing the project's value by making it redundant.

At the start of the next fiscal year, sponsors propose new candidate projects, each with their own value and cost, which are added to the set of unselected projects from the previous year. In addition, changes to the remaining uncommitted budget are provided by the Finance organization, which may be due to a variety of factors, such as Government-wide cost cutting measures or the addition of new funds due to changes in the defence policy. Decision makers must again apply a decision policy, like in the previous fiscal year, to select which candidate projects to allocate funding. The cycle then continues as described above.

Given the uncertain number of candidate projects put forth each fiscal year, lack of knowledge of their value and cost beforehand, and that investment decisions are made throughout the planning horizon, this problem is aptly described as a sequential decision making problem under uncertainty. The following section formulates two models using the universal modelling framework for sequential decisions (Powell, 2022). Section 5 presents a hypothetical case study which demonstrates how the models may be used to provide decision support.

4 SEQUENTIAL DECISION PROBLEM FORMULATION

This section formulates the defence planning scenario introduced in the previous section as a sequential decision problem. To do so, the universal modelling framework for sequential decisions is employed. This framework consists of three components (Powell, 2022, pp. 10-14):

- a *sequential decision model* that describes the state variable, decision variables, exogenous information that arrives after a decision is made, a transition function that defines the dynamics of how the system evolves from one state to the next, and an objective function;
- the *stochastic modelling* that describes the uncertain information contained within the problem's initial state and exogenous information; and
- the *decision policies* to be explored.

The remainder of this section presents these three components applied to the problem described in Section 3.

4.1 Sequential Decision Model

Throughout the defence planning horizon, typically 20 years, there are a set of decision epochs $\mathcal{T}^D$ in which decisions are made regarding projects, and a set of budget epochs $\mathcal{T}^B$ in which budget constraints are enforced, where $\mathcal{T}^D \subset \mathcal{T}^B$. For this problem, an epoch occurs every year throughout the planning period. The size of the set $\mathcal{T}^B$ is determined by the length of the funding requirements of the candidate projects and generally exceeds the number of decision epochs. This aims to address end effects associated with defence planning in terms of financial commitments, thus eliminating "outrageous behavior at the end of the planning horizon" (Brown et al., 2004, p. 422).

At each decision epoch $t \in \mathcal{T}^D$ there exists a set of candidate projects $\mathcal{P}_t$. For each project $i \in \mathcal{P}_t$, there is an associated information vector $p_{t,i}$ that contains two elements: the value of the project $v_{t,i}$; and the yearly cost of the project $c_{t,i} = (c_{t,i,t'})_{t' \geq t, t' \in \mathcal{T}^B}$.

The state variable is then defined as

$$S_t = (P_t, B_t), \tag{1}$$

where $P_t = (p_{t,i})_{i \in \mathcal{P}_t}$ with element $p_{t,i}$ being defined as above, and $B_t = (b_{t,t'})_{t' \geq t, t' \in \mathcal{T}^B}$ is a vector of budgets.

At each decision epoch t, the decision regarding which candidate projects are selected is defined as $x_t = (x_{t,i})_{i \in \mathcal{P}_t}$: $x_{t,i} = 1$ if project i is selected, and 0 otherwise. The decision vector at epoch t is constrained by the available budgets,

$$\sum_{i \in \mathcal{P}_t} x_{t,i} c_{t,i,t'} \leq B_{t'}, \; t' \geq t, \, t' \in \mathcal{T}^B, \tag{2}$$

collectively labelled as $\mathcal{X}(S_t)$.

The state transition function is defined as $S_{t+1} = S^M(S_t, x_t, W_{t+1})$, where W_{t+1} is exogenous information that arrives after the decision x_t is made. This includes new candidate projects $\hat{\mathcal{P}}_{t+1}$, where for each project $i \in \hat{\mathcal{P}}_{t+1}$ there is an associated vector $\hat{p}_{t+1,i}$ defined as $(\hat{v}_{t+1,i}, \hat{c}_{t+1,i})$ with the components being defined as those in $p_{t,i}$. Thus, W_{t+1} is given as

$$W_{t+1} = (\hat{P}_{t+1}, \hat{B}_{t+1}), \tag{3}$$

where $\hat{P}_{t+1} = (\hat{p}_{t+1,i})_{i \in \hat{\mathcal{P}}_{t+1}}$ is a vector of new project information vectors as described above, and $\hat{B}_{t+1} = (\hat{b}_{t+1,t'})_{t' \geq t+1, t' \in \mathcal{T}^B}$ is a vector of changes in the remaining uncommitted budgets in upcoming years. It should be noted that the exogenous information depends on the decision at time t; that is, W_{t+1} depends on the post-decision state $S_t^x = S^x(S_t, x_t)$ (Powell, 2022, p. 580).

Given x_t and W_{t+1}, the transition function then: reduces the value of projects not selected in year t by 10%; combines the set of unselected projects with the newly arrived projects to determine the set of candidate projects in year $t+1$; and updates the budgets in year $t+1$ onward.

When a decision x_t is made, a contribution (or reward) is received and is defined as

$$C(S_t, x_t) = \sum_{i \in \mathcal{P}_t} v_{t,i} x_{t,i}. \tag{4}$$

Given this contribution function, the objective is then to maximize the expected cumulative value of the selected projects. The objective is given as

$$\max_{\pi \in \Pi} \mathbb{E}\left(\sum_{t \in \mathcal{T}^D} C(S_t, X^{\pi}(S_t))|S_0\right), \quad (5)$$

where: π is a label that carries information about a decision policy; Π is the set of all decision policies considered; $X^{\pi}(S_t)$ represents the implementation of a decision policy that returns the decision x_t for each state $S_t \in \mathcal{S}$ and is bounded by the constraints $\mathcal{X}_t(S_t)$; S_0 is the initial state that includes the initial budgets B_0 and projects P_0; and $\mathbb{E}$ is an expectation operator that is over all uncertainties within the problem.

This model, labelled as Model-I, describes the problem discussed in Section 3. However, this formulation can be extended to include a range of additional uncertainties, such as: injection of additional funding beyond yearly fluctuations due to political uncertainty, new defence policies, etc.; cancellation of projects from the set of previously selected projects due to advancement of adversarial capabilities, geopolitical changes, etc.; and updates to the cost of previously unselected projects due to inflation, changes in project scope, etc. These additions require changes to the state transition function and exogenous information, and in the case of the cancellation of projects the state variable as well.

In this paper we focus on extending Model-I by updating the costs of projects not selected at epoch t and using these updates in epoch $t+1$. This model, labelled as Model-II, requires that W_{t+1} to be modified such that

$$W_{t+1} = (\hat{P}_{t+1}, \hat{C}_{t+1}, \hat{B}_{t+1}), \quad (6)$$

where $\hat{C}_{t+1} = (\hat{c}_{t+1,i,t'})_{t' \geq t+1, t' \in \mathcal{T}^B}$ is a vector of updated yearly costs for the set $\{i | i \in \mathcal{P}_t, x_{t,i} = 0\}$. In addition, the transition function must be modified to update the costs of projects not previously selected.

4.2 Stochastic Modelling

As described in the previous section, in Model-I the exogenous information that arrives after a decision is made includes two stochastic components: the arrival of new projects, and changes in the future years' budgets.

In this study, $\hat{P}_{t+1}$ is based on a data set of 198 projects collected in the 2022 CIPPR process and includes the project value, total cost, and duration. Three variations of each project were added to the original data set (total of 792 sample projects), each with their own value, total cost, and duration such that each was scaled using a random value from a uniform distribution ranging between $\pm 20\%$. This set was then randomly sampled with replacement to create the new projects arriving in both the initial state S_0 and exogenous information W_{t+1}.

The number of new projects arriving in a given year was modelled using a binomial distribution, with a maximum of 20 projects and each having a 50% probability of arriving. Distribution parameters are notional. While project data was available for years 2018-2022, inconsistencies in annual data collection methods made it difficult to determine the number of new projects arriving. Each project's total cost was then split over the duration of the project using a triangular distribution. The peak of the triangular distribution occurs at half of the duration of the project. In the case of odd numbered durations, the peak occurs at half of the duration rounded down to the closest year.

Lastly, $\hat{B}_{t+1}$ is notionally determined. After projects are selected using a policy $X^{\pi}(S_t)$ and their costs removed from the current and future years' budgets, each remaining future year's budget is scaled by a random value from a uniform distribution ranging between $\pm 10\%$.

Regarding Model-II, the costs of unselected projects ($x_{t,i} = 0$) that are propagated to the next epoch $t+1$ have their costs adjusted prior to doing so. Each project's annual costs are scaled by a single random value which is sampled from a uniform distribution ranging from zero to 3.7%. The upper bound was calculated as a yearly inflation percentage to achieve a 20% increase in total cost over a five year period. The random value is then multiplied by: -1 (cost decrease) or +1 (cost increase), with a 10% probability of being -1, and a 90% probability of being +1, therefore creating a higher likelihood of a project's costs increasing over time.

4.3 Decision Policies

The purpose of the objective function in Equation 5 is to find the decision policy that maximizes the expected value of the projects selected. In this paper we limit the policies considered to three myopic project selection policies.

The first policy, labelled as the *Ranked list policy* and given as $X^{Rl}(S_t)$, is a Policy Function Approximation (PFA) that "directly [returns] an action given a state, without resorting to any form of imbedded optimization, and without using any forecast of future information" (Powell, 2011, p. 221). This policy ranks candidate projects by value, and selects the top-ranked projects that fit within the available budget. The second policy, labelled as the *Knapsack policy* and given as $X^K(S_t)$, is a Cost Function Approximation (CFA) that aims to maximize the value of selected projects in a given year and "[does] not make any effort at approximating the impact of a decision now on the future" (Powell, 2022, p. 535). The third policy, labelled as the *Contingency fund policy* and given as $X^{Cf}(S_t, \theta)$, is similar to the *Knapsack policy* with the exception that it is parameterized by a single scalar θ (ranging between zero and one) that limits the amount of budget that can be used in future years when making a decision at epoch t. For example, $\theta = 0.4$ represents a 40% contingency fund. This policy requires an adjustment to the constraints $\mathcal{X}(S_t)$ listed in Equation 2, specifically

$$\sum_{i \in \mathcal{P}_t} x_{t,i} c_{t,i,t} \leq B_t, \quad (7)$$

$$\sum_{i \in \mathcal{P}_t} x_{t,i} c_{t,i,t'} \leq (1-\theta) B_{t'}, \forall t' \geq t+1,\ t' \in \mathcal{T}^B, \quad (8)$$

where Equation 7 enables the full budget to be used in epoch t and Equation 8 ensures that only a portion of the future budgets are used. Note that the Knapsack policy is a special case of the Contingency fund policy when $\theta = 0$. As a Knapsack policy is commonly known, the two policies are included separately to distinguish the results.

5 RESULTS

In this section, we demonstrate how the mathematical models presented in the previous section may be used to provide decision support. While many questions may be posed by both planners and decision makers during the defence planning process, in this section we consider the following:

- Given Model-I, how do the decision policies described in subsection 4.3 perform relative to the Ranked list policy as a function of the planned annual budget?
- How does the inclusion of project cost uncertainty (for projects that are not selected and are propagated to the next decision epoch) in Model-II impact the results when compared to Model-I?

To answer these questions, the two models and the decision policies were implemented in Python v3.12.6. Given either Model-I or Model-II, an initial state S_0, and decision policy, 50 trials were executed and the expected cumulative portfolio value (see Equation 5) and associated 95% Confidence Interval (CI) were computed. All trials were executed on an AMD Ryzen 5 4500U CPU with Radeon Graphics running at 2.38 GHz with 16 GB RAM and using the Windows 10 64-bit operating system.

The aim of the first question is to provide decision makers with insights on: which of the policies studied maximizes the expected cumulative portfolio value given an initial yearly budget for a 20-year period (B_0); and the robustness the recommended policy to changes in the initial budgets. The results of the trials are depicted in Figure 2a and the 95% CIs are listed in Table 1.

The results demonstrate that all policies outperform the benchmark Ranked list policy, and the Contingency fund ($\theta = 0.6$) policy performs the best across all initial budgets studied. While the Contingency fund policy with $\theta = 0.8$ performed similar to same policy with $\theta = 0.6$, the latter consistently outperformed the former as listed in Table 2. However, it is worth noting that as the initial yearly budget increases, the difference between the two policies becomes negligible.

The project selection results for each policy were also examined to investigate why the Contingency fund policies with higher reserves ($\theta = 0.6, \theta = 0.8$) outperformed the other policies. It was noted that these policies selected fewer high cost projects since the available future budgets were limited by the reserved funding. The total costs in the sample project data ranged from 10s of millions to 10s of billions of dollars. Therefore, the selection of a single high cost project significantly reduces the number of lower cost projects that can be selected in future years. Additionally, the project values range less significantly when compared to the total costs and are only moderately correlated to the logarithm of the total cost. As such, the selection of a high cost project does not necessarily result in a significant increase in the cumulative portfolio value. While all tested policies are myopic and do not consider future effects, the Contingency fund policy begins to mimic forward-looking policies by forcing the reservation of funding for future projects.

It may therefore be concluded that, for the scenario and policies studied, using a knapsack-based policy that reserves 60% of the available budget in future years will provide the most value to defence. While our scenario differs from that studied by Fisher et al. (2015)—different project values, different cost profiles, uncertainty in future budgets—our results generally agree with their conclusion that holding back budget results in a higher cumulative portfolio value.

The aim of the second question is to evaluate the impact that incorporating project cost uncertainty, of delayed projects (Model-II), has on the policy recom-

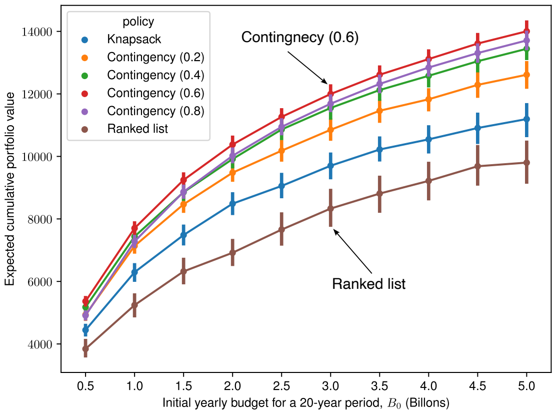

(a) Model-I.

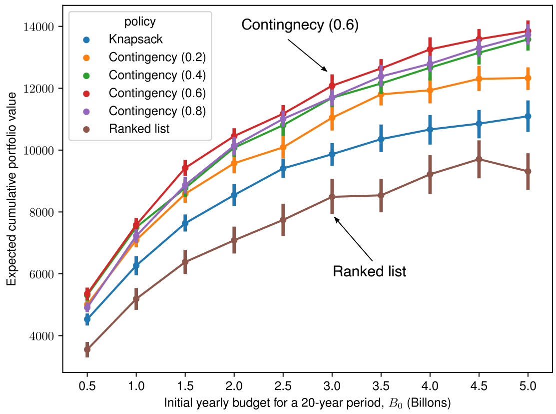

(b) Model-II.

Figure 2: Expected cumulative portfolio value and 95% CI as a function of initial yearly budget for a 20-year period (B_0) for each decision policy described in subsection 4.3.

Table 1: Summary of 95% CIs for the percentage difference between each policy and the benchmark Ranked list policy when using Model-I. **Key**: Rl = Ranked list, K = Knapsack, Cf (θ) = Contingency fund with a specific value of θ. All values are rounded to two significant figures.

Policy	Initial yearly budget for a 20-year period B_0 (Billions)									
	0.5	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0
Rl	-	-	-	-	-	-	-	-	-	-
K	[14, 30]	[18, 31]	[16, 28]	[20, 30]	[16, 28]	[15, 25]	[14, 25]	[12, 21]	[11, 19]	[12, 23]
Cf (0.2)	[27, 44]	[35, 53]	[31, 47]	[35, 49]	[32, 49]	[29, 43]	[28, 43]	[27, 40]	[25, 41]	[27, 43]
Cf (0.4)	[33, 51]	[40, 60]	[38, 53]	[41, 56]	[41, 59]	[37, 55]	[36, 52]	[35, 50]	[33, 50]	[36, 54]
Cf (0.6)	[38, 56]	[45, 66]	[44, 62]	[48, 65]	[46, 66]	[42, 61]	[41, 59]	[40, 57]	[39, 57]	[41, 62]
Cf (0.8)	[26, 44]	[37, 58]	[38, 55]	[43, 59]	[42, 62]	[40, 58]	[38, 56]	[38, 54]	[36, 54]	[39, 59]

Table 2: Summary of 95% CIs for the percentage difference between the Contingency fund ($\theta = 0.6$) policy and the Contingency fund ($\theta = 0.8$) policy when using Model-I. All values are rounded to one significant figure.

Initial yearly budget for a 20-year period B_0 (Billions)									
0.5	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0
[8, 10]	[5, 7]	[3, 5]	[3, 5]	[2, 4]	[2, 4]	[2, 3]	[1, 3]	[1, 3]	[1, 3]

mendation. The results for the trials are depicted in Figure 2b. Although a detailed analysis is not presented here, the results (similar to those in Figure 2a) suggest that the Contingency fund ($\theta = 0.6$) policy outperforms the benchmark Ranked list policy across all initial yearly budgets studied and performs better than when other values of θ are used. The Model-II results provide further evidence that for the scenario and policies studied, the Contingency fund ($\theta = 0.6$) policy is robust with respect to different initial yearly budgets, changes in future budgets, as well as changes in costs of those projects there are delayed.

6 CONCLUSION

This paper presented initial work on the development of a sequential decision model of the DND/CAF multi-gate PAP. The objective was to investigate the impact of different decision policies in a sequential decision model that explicitly includes uncertainties in budgets and the project information in each new epoch.

The results from Model-I support the conclusion that a Contingency fund policy that reserves 60% of the available budget for future years performs the best. A policy that reserves 80% of the available budget performs almost as well when the initial yearly budgets are large. Regardless, the Contingency fund policies outperform both the benchmark Ranked list policy and Knapsack policy. Model-II demonstrates that the same decision policy, the Contingency fund policy that reserves 60%, remains robust to cost uncertainty in delayed projects and continues to outperform the other tested policies.

This initial model demonstrates the potential of a sequential decision model to determine better decision policies for defence planning. Three areas of future development have been identified. The first area covers improvements to the sequential model to better reflect the real-world dynamics of the DND/CAF PAP. The initial model only covered the first step of the PAP, therefore the model can be expanded to cover the multiple gates of the PAP. Additionally, the awarding of the value can be delayed to when the project is delivered, rather than when the project is selected. The modelled uncertainties can also be expanded to capture changes in project value and scheduling over time, as well as uncertainties that may force the selection or cancellation of a project. The second area of development covers the decision policies. The set of policies can be expanded to include future-looking policies, and reinforcement learning can be explored as a method for determining an optimal decision policy. The last area of development aims to improve the project dataset. The current dataset is limited to a single year of projects. This dataset can be expanded to include additional years of historical data. Tools, such as natural language processing, can be explored for extracting data from historical project documentation. Data augmentation techniques can also be explored for supplementing the real project dataset.

The results presented here are encouraging, and the completion of this investigation should provide more information to guide the management of projects under uncertainty in defence planning.

ACKNOWLEDGMENTS

The authors would like to acknowledge the contributions of Ammar Lakdawala, who implemented an initial version of the sequential decision models in Python and conducted extensive testing in spring 2024.

ACRONYMS

CAF Canadian Armed Forces

CI Confidence Interval

CIPPR Capital Investment Program Plan Review

CFA Cost Function Approximation

CFD Chief of Force Development

DND Department of National Defence

DP Dynamic Programming

MDP Markov Decision Process

NATO North Atlantic Treaty Organization

PAP Project Approval Process

PFA Policy Function Approximation

SPARC Strategic Portfolio Analyzer with Reconfigurable Components

VCDS Vice Chief of Defence Staff

VIPOR Visual Investment Optimization and Revision

REFERENCES

Balkaran, R. R. (2021). Process alignment: An assessment of dnd project approval process in the 2019 project approval directive. Master's thesis, Canadian Forces College, Toronto, Ontario. Retrieved from: https://www.cfc.forces.gc.ca/259/ 290/23/286/Balkaran.pdf.

Bellman, R. (1957). *Dynamic Programming*. Princeton University Press.

Bertsekas, D. and Tsitsiklis, J. (1996). *Neuro-Dynamic Programming*. Athena Scientific, Belmont, Massachusetts, First edition.

Brown, G. G., Dell, R. F., and Newman, A. M. (2004). Optimizing military capital planning. *Interfaces*, 34(6):415–425.

Chen, C. and Wheaton, K. (2024). Investigating the Impact of Project Dependencies on Capital Investment Decisions in Defence. In *17th NATO Operations Research and Analysis (OR&A) Meeting Proceedings*. NATO Science and Technology Organization. STO-MP-SAS-OCS-ORA-2023.

Crawford, J., Do, J., Leduc, C., Malik, A., Gormley, K., Luebke, E., and Scherer, W. (2003). SWORD: the latest weapon in the project selection arsenal. In *IEEE Systems and Information Engineering Design Symposium, 2003*, pages 95–100.

Filinkov, A. and Dortmans, P. J. (2014). An enterprise portfolio approach for defence capability planning. *Defense & Security Analysis*, 30(1):76–82.

Fisher, B., Brimberg, J., and Hurley, W. (2015). An approximate dynamic programming heuristic to support non-strategic project selection for the Royal Canadian Navy. *The Journal of Defense Modeling and Simulation*, 12(2):83–90.

Government of Canada (2019). Defence Investment Plan 2018. Retrieved 22 October, 2024, from https://www.canada.ca/en/ department-national-defence/corporate/reports-publications/defence-investment-plan-2018.html.

Government of Canada (2022). Status report on transformational and major crown projects. Retrieved 23 October, 2024 from https://www.canada.ca/en/department-national-defence/corporate/reports-publications/departmental-plans/departmental-plan-2022-23/supplementary-information/report-crown-project1.html#toc14.

Government of Canada (2024a). Our North, Strong and Free: A Renewed Vision for Canada's Defence. Retrieved 15 October, 2024, from https://www.canada.ca/en/ department-national-defence/corporate/reports-publications/north-strong-free-2024.html.

Government of Canada (2024b). Vice Chief of the Defence Staff (VCDS). Retrieved 22 October, 2024, from https://www.canada.ca/en/ department-national-defence/corporate/organizational-structure/vice-chief-defence-staff.html.

Harrison, K. R., Elsayed, S., Garanovich, I., Weir, T., Galister, M., Boswell, S., Taylor, R., and Sarker, R. (2020). Portfolio optimization for defence applications. *IEEE Access*, 8:60152–60178.

Kuo, F. and Sloan, I. (2005). Lifting the curse of dimensionality. *Notices of the American Mathematical Society*, 52(11):1320–1328. Retrieved from: https://www.ams.org/journals/notices/200511/ fea-sloan.pdf?adat=December%202005&trk=200511fea-sloan&cat=feature&galt=feature.

Locatelli, A. (2023). Optimization methods for knapsack and tool switching problems. *4OR*, 21:715–716.

Moallemi, E. A., Elsawah, S., Turan, H. H., and Ryan, M. J. (2018). Multi-objective decision making in multi-period acquisition planning under deep uncertainty. In *2018 Winter Simulation Conference (WSC)*, pages 1334–1345.

North Atlantic Treaty Organization (2022). NATO 2022 Strategic Concept. Retrieved 23 October, 2024, from https://www.act.nato.int/wp-content/uploads/2023/05/290622-strategic-concept.pdf.

North Atlantic Treaty Organization (2024). Washington Summit Declaration. Retrieved 23 October, 2024, from https://www.nato.int/cps/en/natohq/ official_texts_227678.htm.

Office of the Parliamentary Budget Officer (2021). The cost of Canada's surface combatants: 2021 update and options analysis. Retrieved 22 October, 2024, from https://www.pbo-dpb.ca/en/publications/RP-2021-040-C–cost-canada-surface-combatants-2021-update-options-analysis–cout-navires-combat-canadiens-mise-jour-2021-analyse-options.

Powell, W. (2011). *Approximate dynamic programming: Solving the curses of dimensionality*. Wiley, Hoboken, New Jersey, Second edition.

Powell, W. (2022). *Reinforcement learning and stochastic optimization: A unified framework for sequential decisions*. Wiley, Hoboken, New Jersey.

Puterman, M. (2005). *Markov Decision Processes: Discrete stochastic dynamic programming*. Wiley, Hoboken, New Jersey.

Rempel, M. and Young, C. (2017). VIPOR: A visual analytics decision support tool for capital investment planning. Scientific Report DRDC-RDDC-2017-R129, Defence Research and Development Canada.

Roncolato, G. (2022). The character of war is constantly changing: Organizations and people who can rapidly and effectively adapt are more likely to prevail. *U.S. Naval Institute Proceedings*, 148/5/1.431. Retrieved from: https://www.usni.org/magazines/proceedings/2022/may /character-war-constantly-changing.

Rossignol, M. (1998). Replacement of shipborne and rescue helicopters. Retrieved 23 October, 2024, from https://www.publications.gc.ca/Collection-R/LoPBdP/CIR/943-e.htm.

Salo, A., Doumpos, M., Liesiö, J., and Zopounidis, C. (2024). Fifty years of portfolio optimization. *European Journal of Operational Research*, 318(1):1–18.

Shafi, K., Elsayed, S., Sarker, R., and Ryan, M. (2017). Scenario-based multi-period program optimization for capability-based planning using evolutionary algorithms. *Applied Soft Computing*, 56:717–729.

Tsaganea, D. (2005). Appropriation of funds for anti-ballistic missile defence: a dynamic model. *Kybernetes*, 34(6):824–833.

U.S. Government Accountability Office (2020). Cost estimating and assessment guide: Best practices for developing and managing program costs. Technical Report GAO-20-195G, U.S. Government Accountability Office.