




Crosscutting Areas

Strategic Workforce Planning Under Uncertainty

Patrick Jaillet,^a Gar Goei Loke,^b Melvyn Sim^c^aDepartment of Electrical Engineering and Computer Science, Operations Research Center, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139; ^bRotterdam School of Management, Erasmus University, 3062PA Rotterdam, Netherlands; ^cNUS Business School, National University of Singapore, Singapore 119245, SingaporeContact: jaillet@mit.edu,  <https://orcid.org/0000-0002-8585-6566> (PJ); loke@rsm.nl,  <https://orcid.org/0000-0003-1007-4575> (GGL); melvynsim@gmail.com,  <https://orcid.org/0000-0001-9798-2482> (MS)

Received: May 22, 2018

Revised: November 3, 2019; January 10, 2021; June 27, 2021

Accepted: July 27, 2021

Published Online in Articles in Advance: December 29, 2021

Area of Review: Policy Modeling and Public Sector OR

<https://doi.org/10.1287/opre.2021.2183>

Copyright: © 2021 INFORMS

Abstract. The workforce planning problem of hiring, dismissing, and promoting has been the perennial difficulty of Human Resource (HR) management. To cope with uncertain attrition, we propose a new approach of finding a course of action that safeguards against violating organizational target-meeting constraints, such as productivity, budget, headcount, dismissal threshold, and managerial span of control. As such, this approach leads to a tractable conic optimization model that minimizes a decision criterion that is inspired by the riskiness index of Aumann and Serrano, for which its value can be associated with probabilistic and robustness guarantees in meeting constraints under uncertainty. Additionally, our model departs from the literature by considering employees' time-in-grade, which is known to affect resignations, as a decision variable. In our formulation, decisions and the uncertainty are related. To solve the model, we introduce the technique of *pipeline invariance*, which yields an exact reformulation that may be tractably solved. Computational performance of the model is studied by running simulations on a real data set of employees performing the same job function in the Singapore Civil Service. Using our model, we are able to numerically illustrate insights into HR, such as the consequences of a lack of organizational renewal. Our model is also likely the first numerical illustration that lends weight to a time-based progression policy common to bureaucracies.

Funding: Financial support from the Public Service Division, Prime Minister's Office, Singapore; and the Singapore Ministry of Education [Grants MOE2016-SSRTG-059, Service Productivity and Innovation Research Programme (SPIRE)] is gratefully acknowledged.

Keywords: manpower planning • organizational studies • stochastic • programming

1. Introduction

Human Resource (HR) function has recently gained prominence. This is driven by the growing practice of Strategic HR Management or Strategic Workforce Planning (Buyens and De Vos 2001, Ulrich and Dulebohn 2015), where human capital is structured to achieve transformational goals of the organization. Strategic Workforce Planning contrasts against traditional workforce planning; the former, Ulrich et al. (2012) argue, is tied in with the strategic objectives of the organization. Although this is recognized as a nascent field, we are seeing in-depth surveys that stocktake on the practice of Strategic Workforce Planning; for example, a study by KPMG surveyed 37 organizations in the United Kingdom, with at least half having at least 10,000 employees.¹ We also see an offering of consultancy services for Strategic Workforce Planning.² These consultancies often provide data-driven workforce planning tools or methodologies.³ All these point to ambitions to utilize data-driven tools in linking workforce planning to business strategy and the actual practice in industry.

Although practitioners employ different Strategic Workforce Planning processes, there are often three stages, as exemplified by the following description from a global management consultancy firm: "Strategic Workforce Planning enables translation of business drivers into demand for skills and capabilities, assesses expected supply within the organization and in the labor market to anticipate gaps or overages, and identifies how to solve them through reskilling, hiring, and redeployment."² The process begins with identification of the organization's strategic objectives and value-add. As Gartner describes, "The first step in building an effective workforce plan is understanding the organization's business strategy and goals."⁴ By squaring the demand for workforce—that is, the workforce resourcing required to support these strategic goals and objectives—against the supply—that is, the existing workforce in the organization—planners are able to identify critical areas of gap to be filled. At this point, many practitioners adopt analytics tools to project the evolution of the workforce into the future. Such tools can "provide organisations/companies with

visibility into their current workforce (the supply) and how the composition of this workforce is expected to change over time as a result of workforce dynamics relating to employees joining, moving within the organization and leaving.⁵ Finally, strategies are devised to close this gap, such as hiring, promotions, transfers, retraining, outsourcing and cosharing, to name a few.

Part of the allure of Strategic Workforce Planning lies in the promise that it adopts a data-driven approach toward the construction of the workforce plan. This is evidenced by how HR analytics has been introduced in almost every aspect of HR Management (Davenport et al. 2010). Examples include attrition and flight risk, talent and pipeline management, recruitment analytics and employee value proposition, underperformance risks, remuneration and benefits, real-time employee engagement and sentiment analysis, learning and gamification in the workplace, team performance, and social networks, to name only a few.

1.1. Linking Data to Strategy

In practice, there are many challenges with executing a data-driven Strategic Workforce Planning process. Most primarily, it is not straightforward to either determine the gaps in the workforce that emerge over time or to deduce the accurate response to close these gaps. For example, an organization is looking to launch a new service, with targeted outcomes within five years. What is the required workforce to support such a function? Should new officers be hired, or should the organization redeploy and retrain its existing officers? Should more specialized workers or generalists that can be flexibly deployed, but have lower task-specific productivity, be employed? Translating these strategic goals, if they can even be described as productivity targets, into actual workforce figures is notoriously difficult amongst practitioners. Moreover, different HR interventions cannot be assessed separately, as it is their combination that affects employees and their behaviors. For example, an employee's career management can have downstream effects on their flight risk, underperformance risk, engagement levels, etc. As such, one needs to plan for this basket of interventions and to minimize the risk of them not articulating into eventual outcomes of individuals and organizational units (Paul and Mitlacher 2008).

These challenges persist, despite the bountiful data available on employees' resignation patterns, performance and learning records, and engagement indicators. This is because HR analytics continues to struggle to draw the link between human capital and organizational outcomes (Marler and Boudreau 2017). Although data abound, Strategic Workforce Planning does not immediately translate into an analysis frame that seamlessly integrates these data and secondary analyses into trade-offs and risks at the organizational level. As such,

organizational leaders have repeatedly reported the use of data as one of the key hurdles in the Strategic Workforce Planning process.¹

In this paper, we hope to make preliminary steps toward this overarching goal. In particular, we would like to concentrate on the topic of workforce planning—how should a business unit hire, promote, and design its operational structure in order to achieve a targeted productivity level, while constrained by budget, availability of workforce, and managerial span of control? This is not simple; the trade-offs between different HR decisions may not be, at first glance, apparent. For example, the optimal staffing level across different competency bands could depend on both the productivity targets that the organization aspires to meet and the expectation of employees on promotion and remuneration.

More specifically, our goal is to propose a data-driven methodology for the Strategic Workforce Planning process by attempting to address two key *capacity-planning* questions:

1. Given productivity targets on various segments of the workforce, how should *staffing levels*—that is, the *number* of employees that are required within every job function and level—be decided, while constrained by the budget and other operating constraints?
2. Over time, how many officers need to be *hired and/or promoted* in order to achieve these staffing levels, while carefully managing for loss of employees through resignations that cannot be controlled?

In contrast, we will not include in the scope of our problem challenges pertaining to the execution of the workforce plan. Specifically, we avoid questions on which employees to promote, hire, or fire. We are cognizant that these decisions often lie within the decision domain of HR practitioners and that we should not readily encroach upon them. These decisions depend on many other factors and considerations that cannot reasonably be articulated, much less optimized. For example, a possible consideration for recruitment would be the mix of skills and working styles that the new member would bring to the team. Such considerations will depend idiosyncratically on the managers and how they envision operations to be run. Indeed, the HR practitioners with which we have interacted for this project have also communicated that these decisions should lie within their control, although they are happy for the model to advise on the *number* of individuals to be promoted.

1.2. Literature Review

The workforce planning problem is not new. Davis et al. (2018) motivate the need for workforce planning in both the context of service continuity and financial planning. Over the last half-century or so, there have been various approaches, such as a simulation-based

or systems-dynamics approach (to raise a few examples: Park et al. 2008 and Chung et al. 2010), an econometric approach (e.g., Roos et al. 1999 and Sing et al. 2012), and, finally, a mathematical programming approach, which is the focus of this paper.

The most popular approach has been from the perspective of a Markov model. Bartholomew et al. (1991) provide a broad overview. Various improvements over the years have incorporated learning effects and productivity (e.g., Gans and Zhou 2002), interdepartmental flows (Song and Huang 2008), and staff scheduling (such as Abernathy et al. 1973 and Kim and Mehrotra 2015), just to name a few extensions. The primary goal of the Markov model is to set up the transition probabilities through the hierarchy and determine the two central questions of attainability (is it possible to transit from one organization of work to another?) and sustainability (what is the minimum cost to do so?). As explained in Guerry and De Feyter (2012), attainability is not always guaranteed. As such, additional conditions and approximate measures (such as fuzzy sets as in Dimitriou et al. 2013) have been introduced. Many of these models also require the development of a heuristic to obtain tractable solutions (as is the case in Gans and Zhou 2002).

At the broader level, some researchers have moved away from the Markov paradigm and approached the problem via dynamic programming (as in Mehlmann 1980, Flynn 1981, and Rao 1990). In order to balance between competing organizational outcomes, some have adopted a goal-programming paradigm (Price and Piskor 1972, Georgiou and Tsantas 2002). In more modern literature, researchers have applied stochastic programming techniques supported by linearizations and Bender's decomposition, as in Zhu and Sherali (2009), in order to tackle the computational difficulties. A recent work by De Feyter et al. (2017) considers a multiple objective model to control for costs and proximity to the desired organizational structure. Their approach, however, does not consider promotions as part of the decision variables.

Nonetheless, these methods suffer from the curse of dimensionality and become rapidly unscalable with the number of input variables. For example, in Zhu and Sherali's case, the stochastic model only solved three out of 10 times in computational tests. In the age of data analytics, taking as input individual-level machine-learning predictions of flight risk and performance risk would very likely exceed the computational limits of these models. Often, the optimal organizational structure or production model may also not be known (Valeva et al. 2017 designed a learning model in the context of uncertain product demand). Moreover, although uncertainty in resignations, which are known to fluctuate wildly, is accounted for in the Markov structure of the problems, it

is not immediately apparent how these models can be made robust to the wrong estimation of resignation likelihood from the data.

Most critically, time spent by an employee in a grade (time-in-grade, for short) is often ignored, although it is known to be a major contributing factor that shapes employee behavior, such as resignations. Studies drawing the connection between resignations and time-in-organization or time-in-grade are not in scarce supply (e.g., Iverson 1999 and Kuwaiti et al. 2016). The data from our partnering agency also illustrate the connection, and it is not a linear one. Understanding HR decisions along the dimensions of time-in-grade is also an important problem (such as Şenerdem 2001 and the subsequent literature). Incorporating them, however, poses challenges—the uncertainty at each time period will depend on decisions made in the previous period. Hence, techniques to deal with it are few and far between.

One of the earliest attempts was by Bres et al. (1980), who presented a linear goal programming model that decided on the number of promotions where time-in-organization was a factor. Subsequently, Kalamatianou (1987) retained the Markov framework, by cutting up the population of employees into those yet-to-be and those ready-to-be promoted and estimating the transition probabilities based on the age distribution. Nonetheless, this did not directly address the interdependence of decision and uncertainty and was a workaround. Finally, Nilakantan and Raghavendra (2008) attempted a Markov model based on both time-in-grade and time-in-organization, but only under strict assumptions. Unfortunately, they also stopped short of attainability.

We also make a quick note about the literature on learning curves (Shafer et al. 2001). In this stream, a learning curve is assumed that describes the evolution of the productivity of employees with time and then optimized under a productivity and cost model. As described by Nembhard and Bentefouet (2012), nonlinear formulations often arise out of optimization problems structured around learning curves. These models may require heuristics or simplifications to solve. From a different perspective, Arlotto et al. (2014) instead utilized an infinite-armed bandit model to understand the trade-offs between productivity and the opportunity cost of retaining a poorly performing employee. We note the presence of such literature in learning curves; however, we seek to describe productivity in a more general fashion and to be able to consider decisions that relate promotion decisions to time-in-grade in a tractable fashion.

The above literature review identifies a few key gaps in the present literature. First, tractability of present models is a key consideration. In the first place, attainability is not guaranteed in Markov models.

Moreover, many of the models we examined may only be solved under heuristics or only for small instances. It is also not clear whether these heuristics provide guarantees on model performance. As we move toward greater integration of predictive analytics with HR management, where large amounts of data are ingested by the model, such approaches will rapidly become untenable. Second, these models cannot be fundamentally extended to take into consideration time-in-grade without key technical innovations. This is despite the benefits of better accuracy, resulting from greater granularity in the decisions and resignations. This is because decisions and uncertainty become intertwined. Attempts to model time-in-grade are also extremely limited.

This lack of computational tractability and modeling flexibility limits the ability of HR practitioners to implement a recruitment and progression strategy based on time-in-grade. At its base, there isn't even conclusive numerical evidence in support of time-based progression, which is practiced across many bureaucracies. We aim to fill this gap in this paper.

To address these challenges, we propose to consider approaches based on robust optimization. In workforce planning, robust optimization has traditionally been applied to staffing and scheduling problems (e.g., Burke et al. 2004, Lusby et al. 2012, and Yan et al. 2018). However, to the best of our knowledge, we haven't seen any literature on its application to strategic workforce capacity planning.

1.3. Contributions

First and foremost, our model is a novel robust multiperiod optimization framework that considers time-in-grade as a second timescale. In particular, the uncertainty and decision space have a specific interdependent structure, termed "pipeline invariance." This improves on the literature because:

1. Its decision criterion is based on Aumann and Serrano's (2008) index, for which its value can be associated with probabilistic and robustness guarantees in meeting organizational constraints under uncertainty. Specifically, it is able to handle distributional ambiguity in the estimation of resignation probabilities.
2. It can be formulated as an exponential conic optimization problem, whose properties can be exploited to be solved efficiently, despite the interdependence of uncertainty and decisions.
3. It may be reasonably extended to incorporate data at the individualized level, which can take as input the results from various predictive analytics models.

Second, we claim that our model provides a novel application in the domain of HR, by providing a possible means to address the difficulties in Strategic Workforce Planning. Our model can be applied within a variety of contexts and is customizable to different

measures of productivity and organizational structure. As such, we can use our model to illustrate insights into HR, in particular, giving quantitative substantiation for a time-based progression model and the ramifications of a lack of organizational renewal. The model is considered in the Singapore Civil Service.

Our work is also closely related to the literature on inventory models and service management. Specifically, Gans and Zhou (2002) motivated their approach from the context of call centers and explicitly drew the connections to inventory models in their paper; Fry et al. (2006) described their workforce model as an adaptation of the newsvendor problem to a decision-dependent context.

1.4. Notation

Given $N \in \mathbb{N}$, let $[N]$ represent $\{1, \dots, N\}$ and denote $[N]_0 := \{0\} \cup [N]$. Let \mathbb{Z}_0^+ be the nonnegative integers. We use bold-faced characters, such as $x \in \mathbb{R}^N$, to represent vectors, while x_i denotes its i -th element. The tilde sign denotes an uncertain or random parameter, such as \tilde{z} , without explicitly stating its probability distribution. We use the convention $\log 0 = \max \emptyset = -\infty$ and $\min \emptyset = \infty$. We shall use $\mathbb{E}_{\mathbb{P}}[\cdot]$ to represent the expectation, with respect to the reference distribution \mathbb{P} , over the uncertainty across all time periods, unless otherwise stated. When the reference distribution is unambiguous, \mathbb{P} is dropped. Where unambiguous, sums are assumed to be over the entire range of the indices.

2. Workforce Planning Under Uncertainty

Traditionally, the workforce planning problem is set up over a finite time horizon $t \in [T]_0$, where $t = T$ is the last time period to be considered. Often, the objective is to attain a known staffing level. Employees are often split into different departments, specializations, or job roles. For now, we assume that the employees belong to just one department or job role. The interested reader is diverted to Appendix B for details on the general setting.

Let the *stock* $(\tilde{s})_l^{t,\tau} := \tilde{s}_l^{t,\tau}$ denote the number of employees at time $t \in [T]$ for all the times up to the last planning time T , having spent $\tau \in [M]_0$ years at grade $l \in [L]$. When $t = 0$, $s_l^{0,\tau}$ represents the known initial data. Each employee in grade l and having spent τ years in the grade is paid wage w_l^τ and generates a return of productivity of r_l^τ . The *organizational structure* is the hierarchy of grades. Similar to existing literature, we categorize individual contributors into skills strata $l \in \mathcal{W} := [\bar{L}]$, where \bar{L} is the highest skills stratum. These contributors are supervised by managers, limited by the maximum number of employees they can manage, called the span of control c_l^τ . In our model, managers occupy the higher grades $l \in \mathcal{M} := \{\bar{L} + 1, \dots, L\}$, where

L is the highest grade in the hierarchy. Promotion is the movement of employees between adjacent strata. For simplicity, assume that promotion only occurs between adjacent grades and ignore complications, such as transfers across departments (see Appendix B for details).

Employees may be lost through *attrition*. In the literature, attrition is often understood as a rate—an annual proportion of stock \tilde{s} . Instead, we hope to model attrition as a random variable depending on the decision variables, so as to capture the interdependence—employees who were not promoted have a different chance of leaving compared with those who were. To do so, we need the following assumption:

Assumption 1.

a. *Different employees make independent resignation decisions; and*

b. *An employee's probability of resignation depends solely on his grade and the time spent in that grade.*

Arguably, Assumption 1 is debatable on both counts. In practice, resignation decisions of different employees might not be independent. For example, employees from the same department might leave due to similar reasons, possibly even influencing each other to leave. Moreover, resignation decisions by employees are also endogenous with the promotion decisions made; for example, it is conceivable that if there are more employees promoted, there may be a perceived impression that there are fewer opportunities for advancement amongst the remaining employees, which might promote higher levels of resignation. In general, part (b) of the assumption would also not be true. It is well known that resignation decisions depend on a wide range of factors, both idiosyncratic to the employee and shared across all employees, such as organizational culture.

Despite these reasons, we have chosen to make this assumption for the reason that it does make the formulation tractable, and, additionally, there are some features that we can build into the model that can alleviate the impact of this assumption. In the discussion after Proposition 1, we will discuss how the definition of the uncertainty set will allow our model to be robust to the independence assumption. Additionally, we also note that Assumption 1(b) is already comparatively relaxed compared with traditional Markov models that assume that the probability of resignation depends solely on the grade.

For both of these assumptions, our model permits an additional modification that can mitigate or reduce their impact. Specifically, our model allows us to group employees into categories $i \in \mathcal{I}$. This would amount to adding an index i , which represents these

groupings, to every state variable, decision variable, and uncertainty. The benefit of this is that Assumption 1 may be relaxed. For example, suppose a clustering was first performed on the employee's likelihood of resignation, using a wide set of predictors and side information, including past evaluations, demographics, reporting lines, and even engagement scores. Then, Assumption 1(b) would now read as "An employee's probability of resignation depends solely on his grade, the time spent in that grade, and the cluster i to which they belong." As such, factors that we believe predict attrition and factors that might lead to the lack of independence previously can all be incorporated into the data set. Based on this, we can construct clusters i and, hence, implicitly capture the effect of these factors in our model. More of this is discussed in Appendix B.

Assumption 1 allows us to model the attrition process via the binomial distribution, where $\text{Bin}(x, q)$ represents the number of successes under x number of trials, each with success probability $q > 0$. In our case, x represents the stock before attrition, and q represents the chance an employee stays within an organization till the next year (also called "retention"). Specifically, define q_l^τ as the probability that an employee that has spent τ time in grade l will voluntarily remain in the organization till the next year, assuming they are not promoted.

2.1. HR Decisions

The workforce planner makes two types of decisions in this process, which articulate into the organizational structure. The first is the promotion of employees. Specifically, let $p_l^{t,\tau} \in [0, 1]$ represent the *decision variable* of the fraction of employees that have spent τ years in grade l at time $t \geq 1$ to be retained in grade l . The remaining fraction $1 - p_l^{t,\tau}$ forms both the promotees and those dismissed.

As such, the Binomial model induces the following dynamics for all $t \in [T]$, $\tau \in [M]$ and $l \in [L]$:

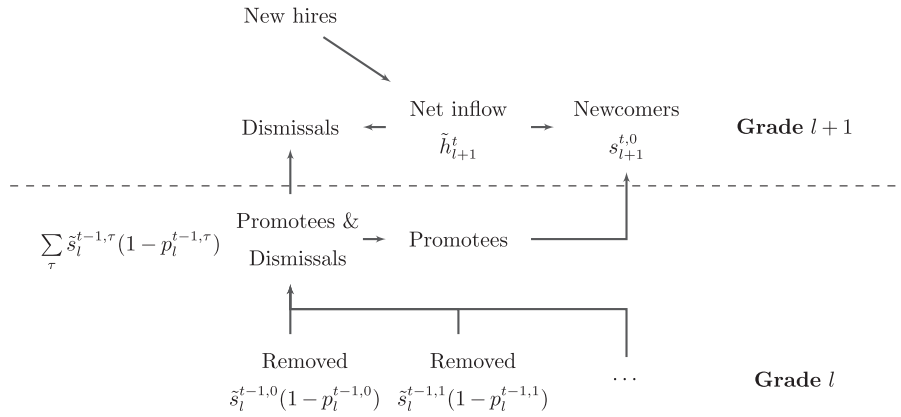
$$\tilde{s}_l^{t,\tau} \sim \text{Bin}\left(\tilde{s}_l^{t-1,\tau-1} p_l^{t-1,\tau-1}, q_l^\tau\right). \quad (1)$$

The sequence of events is as follows: Amongst all of the employees $\tilde{s}_l^{t-1,\tau-1}$, the planner makes the decision to only retain $\tilde{s}_l^{t-1,\tau-1} p_l^{t-1,\tau-1}$ of them in this grade. Attrition then sets in, with each of these employees experiencing a $1 - q_l^\tau$ chance of resigning.

The second type of decision to be made is the number of newcomers to each grade l . Notice that $\tilde{s}_l^{t,0}$ represents the number of employees who have spent zero years at grade l at time $t \geq 1$. This is precisely the newcomers to grade l . Let this be a decision variable.

As such, we are able to determine the net inflow of employees into the organization at grade l at time t . Denote this random variable by \tilde{h}_l^t , and it is given by

Figure 1. Flow Balance Amongst Hiring, Dismissal, and Promotion Decisions



the following expression. Figure 1 illustrates how this expression is derived.

$$\tilde{h}_{l+1}^t := s_{l+1}^{t,0} - \sum_{\tau'} \tilde{s}_l^{t-1,\tau'}(1-p_l^{t-1,\tau'}) \forall t \in [T], \forall l \in [L]. \quad (2)$$

Consequently, whenever $\tilde{h}_{l+1}^t \geq 0$, \tilde{h}_{l+1}^t new hires are made to replenish the stock. Otherwise, $-\tilde{h}_{l+1}^t$ of employees removed from grade l will be dismissed. As an illustration, suppose $s_{l+1}^{t,0} = 4$, and five employees are to be removed from grade l . The interpretation is that out of all the employees from grade l , HR is to choose one to be terminated and four to be promoted. In practice, HR can choose to terminate the worst-performing employee in that grade and promote the best four or to use any other metric they so desire. Observe that at the boundary, we have $s_{L+1}^{t,0} = 0$ to denote the situation when all employees removed from grade L will be dismissed.

2.2. Target-Meeting Constraints

As we motivated in the introduction, the organization plans for various target-meeting constraints during the Strategic Workforce Planning process that must be satisfied. These constraints can be affected by uncertain attrition. These include, inter alia, productivity, budget, headcount, dismissal threshold, and managerial span of control.

1. *Productivity constraint:* $\sum_{l,\tau} \tilde{s}_l^{t,\tau} r_l^\tau \geq P_t, \forall t \in [T]$.

2. *Headcount constraint:* $\sum_{l,\tau} \tilde{s}_l^{t,\tau} \leq H_t, \forall t \in [T]$,

3. *Budget constraint:* $\sum_{l,\tau} \tilde{s}_l^{t,\tau} w_l^\tau \leq B_t, \forall t \in [T]$.

4. *Span of control constraint:* For each $l \in \mathcal{M}$, let $\mathcal{W}_l \subseteq [l-1]$ be the employee grades supervised by the manager. Then,

$$\sum_{\tau} \tilde{s}_l^{t,\tau} c_l^\tau \leq \sum_{\lambda \in \mathcal{W}_l} \tilde{s}_l^{t,\tau}, \forall l \in \mathcal{M}, \forall t \in [T].$$

We can simplify this by letting

$$b_{l,\lambda}^\tau = \begin{cases} -c_l^\tau & \text{if } \lambda = l \\ 1 & \text{if } \lambda \in \mathcal{W}_l \\ 0 & \text{otherwise} \end{cases}$$

then the constraint may simply be written as $\sum_{\lambda,\tau} \tilde{s}_l^{t,\tau} b_{l,\lambda}^\tau \leq 0, \forall t \in [T], \forall l \in \mathcal{M}$.

5. *Dismissal threshold constraint:* $-\tilde{h}_{l+1}^t \leq F_{l+1}^t, \forall t \in [T], \forall l \in [L]$, meaning that no more than $F_{l+1}^t \geq 0$ employees ought to be dismissed. Equivalently, this is

$$\sum_{\tau'} \tilde{s}_l^{t-1,\tau'}(1-p_l^{t-1,\tau'}) - s_{l+1}^{t,0} \leq F_{l+1}^t. \quad (3)$$

Table 1 describes all of the variables and parameters defined above in the model.

2.3. On Productivity Constraints

Although we only consider a single form for the productivity constraint, the model permits the planner to include as many productivity constraints as necessary, and they do not need to be in the same units. The only restriction is that the constraints must be linear in $\tilde{s}_l^{t,\tau}$. This turns out to be reasonably general—many measures of productivity can be described as such. We describe the following examples:

1. *Quantity/quality of completed work:* Suppose an employee having worked for τ years at grade l can finish r_l^τ pieces of work in an allocated time, then the total quantity of work completed is $\sum_{\tau} r_l^\tau \tilde{s}_l^{t,\tau}$. This is in the linear form required.

2. *Time (or average time) to complete tasks:* Again, if an employee having worked for τ years at grade l takes u_l^τ to finish a task, then $\sum_{\tau} u_l^\tau r_l^\tau \tilde{s}_l^{t,\tau}$ is the total time taken to complete all the tasks, and $\sum_{\tau} \frac{u_l^\tau r_l^\tau}{\sum_{\tau'} r_l^{\tau'} \tilde{s}_l^{t,\tau'}} \tilde{s}_l^{t,\tau}$ is the average time to complete tasks. Hence, an average time

Table 1. List of Parameters and Variables

Dimensions	Description
T	Last modeling time
M	Largest possible years-in-grade
L	Largest possible grade
State and decision variables	
$\tilde{s}_l^{t,\tau}$	Random variable of the number of employees having spent $\tau > 0$ years at grade l at time t
\tilde{h}_l^t	Random variable of the net inflow of employees into the organization at grade l at time t
$s_l^{t,0}$	Decision variable of the number of newcomers to grade l at time t
$p_l^{t,\tau}$	Decision variable of the proportion of employees having spent $\tau > 0$ years at grade l to be retained at time t
Parameters	
$s_l^{0,\tau}$	Current employees having spent τ years at grade l
q_l^τ	Retention probability of an officer having spent τ years in grade l
r_l^τ	Productivity rate of having spent τ years in grade l
P_t	Productivity target to be achieved at time t
H_t	Headcount target to be kept within at time t
w_l^τ	Wage of an officer having spent τ years in grade l
B_t	Budget target to be kept within at time t
c_l^τ	Span of control of a manager having spent τ years in grade l
F_{l+1}^t	Target to keep the number of dismissed officers from grade l within at time t

constraint $\sum_\tau \frac{u_l^\tau r_l^\tau}{\sum_{\tau'} r_l^{\tau'} \tilde{s}_l^{t,\tau'}} \tilde{s}_l^{t,\tau} \leq U^t$ has the equivalent expression $\sum_\tau (u_l^\tau - U^t) r_l^\tau \tilde{s}_l^{t,\tau} \leq 0$, which is in the linear form as desired.

3. Less common forms of productivity can also be considered—for example, chance of defective product. Suppose it is necessary to keep the chance of a defective product under some bound ϵ . Suppose every employee has an independent probability $1 - j_l^{(\tau)}$ of creating a defective product. Then, the probability that no defective product is created amongst all the goods is $\prod_\tau j_l^{(\tau) r_l^\tau \tilde{s}_l^{t,\tau}}$. Then, the constraint becomes $\sum_\tau r_l^\tau \log(j_l^{(\tau)}) \tilde{s}_l^{t,\tau} \leq \log(1 - \epsilon)$, which is in the desired form.

In particular, measuring the quality and quantity of finished work could be a daunting task. Assuming that a data set of observed work quantities and employee demographics, containing time-in-grade information, is given, one can, in the simplest case, take an average over all employees in the same grade and having spent the same time in that grade to obtain the curve r_l^τ . As this is subjected to potential noise, an additional smoothing in form of a parametric regression could also be performed. In the numerical simulations later, we shall see how this could be done. Once r_l^τ is obtained, the productivity constraint $\sum_\tau r_l^\tau \tilde{s}_l^{t,\tau} \geq P_t$ could be written.

A common approach to parametrically smooth the curve r_l^τ is via learning curves—for example, using the form that appears in Shafer et al. (2001). In learning curves, the productivity of the employee is assumed to vary with the amount of work experience—in other words, time-in-grade. As such, *learning curves*

are fully compatible with our model. Using the notation in our paper,

$$r_l^\tau = \eta_l \left(\frac{\tau + \nu}{\tau + \nu + \psi_l} \right), \quad (4)$$

where τ is the accumulated experience, here in the context of our exposition, understood as the time-in-grade, and η_l , ν , and ψ_l are fitted parameters denoting the maximum asymptotic productivity that can be reached for that grade l , contribution from prior experience, and a learning rate term idiosyncratic to the grade l , respectively. Similar to the previous case, once the parametric learning-curve model is obtained, one simply needs to iterate over discrete times-in-grade $\tau \in [M]_0$ to obtain the vector r_l^τ to compose the productivity constraint $\sum_\tau r_l^\tau \tilde{s}_l^{t,\tau} \geq P_t$. We present a simple case study in Appendix B.3 to illustrate specifically how this is done on a synthetic data set.

Recently, it has been increasingly popular to use wider sets of employee data to predict employee performance. With the addition of an index i to represent employee clusters, our model is able to remain compatible with such data-driven predictive analytics methodologies. More is discussed in Appendix B.

The above approach assumes the expected productivity rate on each employee. More generally, it is also possible to model productivity of individual officers as being independently and identically drawn from a productivity distribution. In this case, we phrase the productivity constraints as

$$\sum_{l,\tau} \tilde{r}_l^\tau \left(\tilde{s}_l^{t,\tau} \right) \geq P_t, \quad (5)$$

where $\tilde{r}_l^\tau(s)$, representing the total random productivity contributions by s employees, each with a random independent and identically distributed productivity rate of \tilde{r}_{li}^τ , is defined as

$$\tilde{r}_l^\tau(s) := \sum_{i \in [s]} \tilde{r}_{li}^\tau.$$

2.4. Decision Criterion

In the literature, one might minimize the costs of maintaining a workforce, maximize the total productivity of employees, or deal with these multiple objectives in the goal-programming sense (for example, in Price and Piskor 1972). However, it could be difficult to prescribe the trade-offs between costs and productivity (e.g., for a maintenance crew), say, in goal programming.

It may also be appropriate in some business contexts neither to maximize output nor minimize operating costs, but to run the least risk of disruption, such as a service center. Without a clear objective function, we instead pursue an optimization model that minimizes this risk. It sounds tempting to minimize the joint probability of constraint violation, similar to the P-model proposed by Charnes and Cooper (1963), which, often and in particular in this case, has intractable formulations. In fact, our goal doesn't necessitate the minimization of the chance of constraint violation per se. Instead, we simply desire a policy that *does not fare too poorly*—in other words, a course of action with some guarantees over the risks of violation.

Aumann and Serrano's (2008) index has this functionality. Let \mathcal{Z} be the set of all random variables on our probability space $(\Omega, \Sigma, \mathbb{P})$. Define the Aumann and Serrano (2008) index as the functional $\mu : \mathcal{Z} \rightarrow \mathbb{R}_0^+ \cup \{\infty\}$:

$$\mu[\tilde{z}] = \inf \{k \geq 0 : C_k[\tilde{z}] \leq 0\}, \quad (6)$$

in terms of the certainty equivalence

$$C_k[\tilde{z}] := \begin{cases} k \log(\mathbb{E}[\exp(\tilde{z}/k)]) & \text{if } k > 0 \\ \mathbb{E}[\tilde{z}] & \text{if } k = \infty \\ \text{esssup } \tilde{z} & \text{if } k = 0. \end{cases}$$

Here, \tilde{z} represents the size of the violation—a positive number constitutes a violation and vice versa. The exponential disutility penalizes ever larger violations.

Proposition 1. *The Aumann and Serrano (2008) index obeys the following properties:*

1. *Satisficing:* $\mu[\tilde{z}] = 0$ if and only if $\mathbb{P}[\tilde{z} \leq 0] = 1$.
2. *Infeasibility:* If $\mathbb{E}[\tilde{z}] > 0$, then $\mu[\tilde{z}] = \infty$.
3. *Convexity:* μ is convex in \tilde{z} .

4. *Probabilistic guarantees:* For $\mu[\tilde{z}] > 0$ and $\phi \geq 0$,

$$\mathbb{P}[\tilde{z} > \phi] \leq \exp(-\phi/\mu[\tilde{z}]).$$

5. *Robustness guarantees:* For any probability measure \mathbb{Q} absolutely continuous in \mathbb{P} and $\mathbb{Q} \neq \mathbb{P}$,

$$\frac{\mathbb{E}_{\mathbb{Q}}[\tilde{z}]}{\mathcal{D}(\mathbb{Q}||\mathbb{P})} \leq \mu[\tilde{z}],$$

where $\mathcal{D}(\mathbb{Q}||\mathbb{P})$ is the Kullback-Leibler divergence of \mathbb{Q} from the reference distribution \mathbb{P} .


Proof. The first four properties are well established (see, for instance, Brown and Sim 2008). The last property arises from the dual representation of the certainty equivalence relating to the Kullback-Leibler (KL) divergence (see, for example, Lim and Shanthikumar 2007) given by

$$C_k[\tilde{z}] = \sup_{\mathbb{Q}} \{\mathbb{E}_{\mathbb{Q}}[\tilde{z}] - k\mathcal{D}(\mathbb{Q}||\mathbb{P})\}. \quad \square \quad (7)$$

The first property states that there is no risk if there is no chance of violation. The second dictates that if violations are always expected, then the risk is always infinite. The third requires convexity, and the fourth is our desired guarantee against constraint violation, which is the consequence of Markov's inequality. Hence, $\mu[\tilde{z}]$ captures the notion of risk—the lower $\mu[\tilde{z}]$ is, the sharper the guarantee against ever larger violations ϕ of the constraint. The last property connects the index with the notion of robust optimization. It implies that even if the true probability distribution were to deviate from \mathbb{P} , the worst-case expectation of the underlying random variable, normalized by its KL divergence from the reference distribution, is bounded below by the index. Intuitively, a lower index is associated with higher tolerance of distributional ambiguity against the impact of constraint violation.

Robustness is critical—any model would naturally be sensitive to the specification of the attrition estimates q_l^τ . In reality, estimating q_l^τ from the data could be subject to large errors (see Figure 2). These errors would arise from a few sources. First, there could be factors affecting resignations that vary over the time span of the data set—for example, the outlook of the economic sector to which the business belongs. Second, by considering the additional dimension of time-in-grade, a greater number of data points is required to achieve the same error per estimate. These errors cannot be fully eradicated, even after parsing the estimates through a smoothing model (be it a Loess regression, such as in Figure 3, or a survival-based model). Without Proposition 1, the model would suffer from similar model-misspecification errors, as experienced by assuming the probability distributions

Figure 2. (Color online) Empirical Retention Rates with Time-in-Grade Observed in the Dataset in the Numerical Study

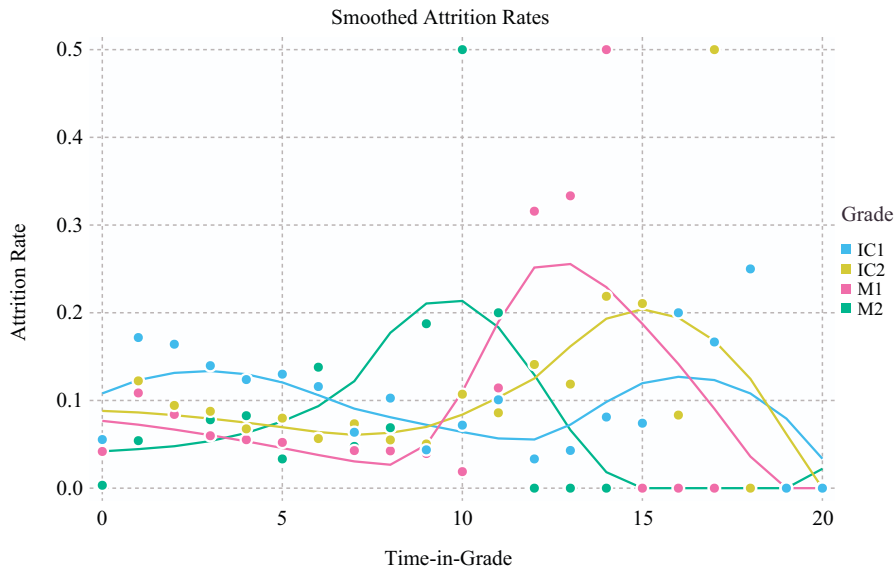


in other stochastic programming approaches in the literature.

Moreover, the robustness guarantee property does also mitigate the independence assumption (Assumption 1) to some extent. In this case, the reference distribution \mathbb{P} would be the assumed empirical distribution and that which we have assumed to be independent across employees. Suppose that the true distribution \mathbb{Q} was indeed dependent, but with marginals that are identical to \mathbb{P} . The Kullback-Leibler divergence is known to capture this difference: $\mathcal{D}(\mathbb{Q}||\mathbb{P}) \geq \mathcal{D}(\mathbb{Q}_x||\mathbb{P}_x) + \mathcal{D}(\mathbb{Q}_y||\mathbb{P}_y)$, if $\mathbb{P} = \mathbb{P}_x\mathbb{P}_y$ is independent, with equality if and only if $\mathbb{Q} = \mathbb{Q}_x\mathbb{Q}_y$ is independent—

that is, the sum of KL-divergences over the marginal distributions is less than the KL-divergence over the full distributions, which is a consequence of the properties of Shannon entropy (Shannon 1948). Hence, by being robust over the set of all distributions that are within some radius in terms of the KL divergence distance from the reference distribution, the model also encompasses dependent distributions that are not too far away from the reference distribution. In other words, the optimal solution obtained would provide a lower bound in terms of the performance to any solution obtained by using any of these distributions.

Figure 3. (Color online) Loess-Smoothed Attrition Rates for Each Grade



Although it might have been possible to pick an uncertainty set that does not require the assumption of independence—for example, a moment uncertainty set (such as one that defines the mean and variance on the number of employees who resigned)—however, it would be difficult to define a tractable formulation. In particular, it would be difficult to decouple the effects of our decisions (promotions) from the definition of the uncertainty set itself (the mean number of resigning employees) or to deal with the multiperiod and sequential decision nature of the problem. Our model, by assuming independence, allows us to decompose these two components after some careful reformulation. In contrast, we have paid a small cost in exchange for this tractability, because the robustness property of the model alleviates having confined ourselves to independent distributions.

Aumann and Serrano’s (2008) index fits well to our multiobjective setting, where we have to assess the combined risk of violating any of the operational constraints. Specifically, given a set of linear constraints, $\{\tilde{x}_j \leq G_j, j \in \mathcal{J}\}$, we evaluate the combined risk under uncertainty via the following decision criterion,

$$\begin{aligned} \mu^* &= \max_{j \in \mathcal{J}} \left\{ \mu \left[\frac{\tilde{x}_j - G_j}{\theta_j} \right] \right\} \\ &= \inf \left\{ k \geq 0 : C_k \left[\frac{\tilde{x}_j - G_j}{\theta_j} \right] \leq 0 \quad \forall j \in \mathcal{J} \right\}, \end{aligned} \quad (8)$$

which picks the value of $\mu[\cdot]$ arising from the worst-performing constraint j . This criterion gives rise to the probabilistic guarantees,

$$\mathbb{P} \left[\frac{\tilde{x}_j - G_j}{\theta_j} \geq \phi \right] \leq \exp(-\phi / \mu^*) \quad \forall \phi \geq 0,$$

and robustness guarantees,

$$\frac{\mathbb{E}_{\mathbb{Q}}[\tilde{x}_j - G_j] / \theta_j}{\mathcal{D}(\mathbb{Q} \parallel \mathbb{P})} \leq \mu^* \quad \forall \mathbb{Q},$$

for all constraints, $j \in \mathcal{J}$, as a result of Proposition 1.

Across constraints, $\theta_j > 0$ are the normalization parameters that calibrate the uncertainty aversion of violating each constraint—for example, say, to emphasize that the budget constraint is more critical than the headcount or across time, such as a stronger aversion to earlier time violation than in the future, as with discounting. In practice, the constraint would be normalized by the target—for instance, $\theta_j = |G_j|$ —and, hence, violations are understood as proportional to the target G_j , making it comparable across different constraints and across different units of measurement.

2.5. Model Formulation

We now state our proposed Strategic Workforce Planning Under Uncertainty (SWPU) model, arising from (8):

$$\begin{aligned} &\inf k \\ \text{s.t. } &C_k \left[\frac{1}{\theta_t^1} \left(\sum_{l,\tau} \tilde{s}_l^{t,\tau} - H_t \right) \right] \leq 0 \quad \forall t \in [T] \\ &C_k \left[\frac{1}{\theta_t^2} \left(\sum_{l,\tau} \tilde{s}_l^{t,\tau} w_l^\tau - B_t \right) \right] \leq 0 \quad \forall t \in [T] \\ &C_k \left[\frac{1}{\theta_t^3} \left(P_t - \sum_{l,\tau} \tilde{s}_l^{t,\tau} r_l^\tau \right) \right] \leq 0 \quad \forall t \in [T] \\ &C_k \left[\frac{1}{\theta_{l,\lambda}^4} \sum_{\lambda,\tau} \tilde{s}_\lambda^{t,\tau} b_{l,\lambda}^\tau \right] \leq 0 \quad \forall t \in [T], \forall l \in \mathcal{M} \\ &C_k \left[\frac{1}{\theta_{l,\lambda}^5} \left(\sum_{\tau} \tilde{s}_l^{t-1,\tau-1} (1 - p_l^{t-1,\tau-1}) - s_{l+1}^{t,0} - F_{l+1}^t \right) \right] \leq 0 \quad \forall t \in [T], \forall l \in [L] \\ &k \geq 0, s_l^{t,0} \geq 0, s_{L+1}^{t,0} = 0, 0 \leq p_l^{t,\tau} \leq 1 \quad \forall t \in [T], \forall l \in [L], \forall \tau \in [M], \end{aligned} \quad (9)$$

where the random variables have the decision-dependent marginal distributions presented in (1). Note that it is not immediately clear whether we can formulate Problem (9) as a tractable optimization problem, because the problem is not convex in the decision variables, even if k is fixed.

3. Tractable Conic Optimization Model

To convexify Problem (9), we perform a change of variables to obtain the following formulation:

$$\begin{aligned} &\inf k \\ \text{s.t. } &C_k \left[\frac{1}{\theta_t^1} \left(\sum_{l,\tau} \tilde{s}_l^{t,\tau} - H_t \right) \right] \leq 0 \quad \forall t \in [T] \\ &C_k \left[\frac{1}{\theta_t^2} \left(\sum_{l,\tau} \tilde{s}_l^{t,\tau} w_l^\tau - B_t \right) \right] \leq 0 \quad \forall t \in [T] \\ &C_k \left[\frac{1}{\theta_t^3} \left(P_t - \sum_{l,\tau} \tilde{s}_l^{t,\tau} r_l^\tau \right) \right] \leq 0 \quad \forall t \in [T] \\ &C_k \left[\frac{1}{\theta_{l,\lambda}^4} \sum_{\lambda,\tau} \tilde{s}_\lambda^{t,\tau} b_{l,\lambda}^\tau \right] \leq 0 \quad \forall t \in [T], \forall l \in \mathcal{M} \\ &C_k \left[\frac{1}{\theta_{l,\lambda}^5} \left(\sum_{\tau} \tilde{s}_l^{t-1,\tau-1} (1 - d_l^{t,\tau} / d_l^{t-1,\tau-1}) - d_{l+1}^{t,0} - F_{l+1}^t \right) \right] \leq 0 \quad \forall t \in [T], \forall l \in [L] \\ &k \geq 0, d_{L+1}^{t,0} = 0, 0 \leq d_l^{t,\tau} \leq d_l^{t-1,\tau-1}, d_l^{0,\tau} = s_l^{0,\tau} \quad \forall t \in [T], \forall l \in [L], \forall \tau \in [M], \end{aligned} \quad (10)$$

where the underlying random variables have the following dynamics

$$\tilde{s}_l^{t,\tau} \sim \text{Bin} \left(\tilde{s}_l^{t-1,\tau-1} \frac{d_l^{t,\tau}}{d_l^{t-1,\tau-1}}, q_l^\tau \right) \quad \forall t \in [T], \forall \tau \in [M], \forall l \in [L]. \quad (11)$$

We use the convention that $d_l^{t,\tau} / d_l^{t-1,\tau-1} = 0$ whenever $d_l^{t-1,\tau-1} = 0$. We call the optimal k^* the risk level associated with this specification of constraints.

Proposition 2. *Models (9) and (10) are equivalent. In particular, given an optimal solution to Problem (10), we can obtain the corresponding solution to Problem (9) by letting $s_{l+1}^{t,0} = d_{l+1}^{t,0}$ and $p_l^{t-1,\tau-1} = d_l^{t,\tau} / d_l^{t-1,\tau-1}$, with $p_l^{t-1,\tau-1} = 0$ whenever $d_l^{t,\tau} = 0$.*

Proof. Consider the feasible solution in Problem (9), as $d_{l+1}^{t,0} = s_{l+1}^{t,0}$ and $d_{l+1}^{0,\tau} = s_{l+1}^{0,\tau}$, we let

$$d_l^{t,\tau} = p_l^{t-1,\tau-1} d_l^{t-1,\tau-1},$$

for all $t \in [T], l \in [L], \tau \in [M]$. Observe from (1), whenever $p_l^{t-1,\tau-1} = 0$, then $\xi_l^{t-1,\tau-1} = 0$ almost surely for all $t' \geq 0$. Equation (11) indicates that this is also true for the decision-dependent random variables in Problem (10). Therefore, the solution would also be feasible in Problem (10). Conversely, consider a feasible solution in Problem (10), and let $s_{l+1}^{t,0} = d_{l+1}^{t,0}$ and $p_l^{t-1,\tau-1} = d_l^{t,\tau} / d_l^{t-1,\tau-1}$, with $p_l^{t-1,\tau-1} = 0$ whenever $d_l^{t-1,\tau-1} = 0$. By inspection, this solution would be feasible in Problem (9). \square

The decision variable $d_l^{t,\tau}$ has the convenient interpretation that it is the number of employees that have stayed for τ years in grade l in the absence of any attrition:

Proposition 3. For all $t \in [T], l \in [L], \tau \in [M]$, we have that $d_l^{t,\tau} = \text{esssup } \tilde{s}_l^{t,\tau} = \mathbb{E}[\tilde{s}_l^{t,\tau}] / \gamma_l^\tau$, where $\gamma_l^\tau = \prod_{t \in [\tau]} q_l^t$.

Proof. The results follows easily from (11). \square

Hence, under Proposition 3, the feasible set of Problem (10) is a polyhedron whenever $k = 0$ or $k = \infty$. To obtain nontrivial solutions, we assume that the constraints are such that $k \in (0, \infty)$ —that is, there does not exist a solution that satisfies all constraints with certainty—and that there exists a solution such that all the constraints can be met in expectation. Organizations operating in the former regime are overly nonchalant in setting targets, while those operating in the latter are deemed unrealistic. Subsequently, we will show that for a given $k > 0$, the feasible set of Problem (10) is convex in \mathbf{d} . As we have explained, quite apart from other approaches, the decision criterion based on the Aumann and Serrano (2008) index, which is associated with robustness guarantees, permits modest divergence from the above assumptions, while ensuring the organizational constraints are satisfied as well as possible under distributional ambiguity.

3.1. Pipeline Invariance

This model turns out to be tractable. We first notice a useful property about the dynamics we have defined:

Property 1 (Pipeline Invariance). Let $q \in (0, 1)$ be fixed. Let $\tilde{s}(x) \sim \text{Bin}(x, q)$ be a family of Binomial-distributed random variables with parameter $x \in \mathcal{X} \subset \mathbb{Z}_0^+$. We say that they are pipeline invariant when

$$\mathbb{E}[\exp(y\tilde{s}(x))] = \exp(x \cdot \rho(y)), \forall x \in \mathcal{X}, \forall y \in \mathbb{R}, \quad (12)$$

and $\rho(\cdot)$, which we call the relay function, is given by the expression,

$$\rho(y) = \log(1 - q + qe^y). \quad (13)$$

Pipeline invariance preserves the exponential functional form under the action of taking expectations. It turns out that pipeline invariance is satisfied by distributions other than the Binomial, such as the Poisson random variable $\text{Pois}(x)$ with rate parameter x or the Chi-squared distribution $\chi^2(df)$ with degrees of freedom df . Moreover, their relay functions are convex over the domain.

As the consequence of pipeline invariance, the constraints in (10) have convex reformulations.

Theorem 1 (Pipeline Reformulation). If integrality of \tilde{s}^t is relaxed, then for any $y \in \mathbb{R}$,

$$\begin{aligned} C_k[y\tilde{s}_l^{t,\tau}] &= \inf_{\xi} k \xi_l^{t-\tau+1,1} \\ \text{s.t. } d_l^{t,\tau} \rho_l^\tau(y/k) &\leq \xi_l^{t,\tau} \\ d_l^{t-t',\tau-t'} \rho_l^{\tau-t'}(\xi_l^{t-t'+1,\tau-t'+1} / d_l^{t-t',\tau-t'}) &\leq \xi_l^{t-t',\tau-t'} \\ \forall t' \in [\min\{t, \tau\} - 1] \end{aligned} \quad (14)$$

where $\rho_l^\tau(y) := \log(1 - q_l^\tau + q_l^\tau e^y)$.

Proof of Theorem 1. We present the proof in Appendix A. \square

Remark 1. Relaxing integrality of \tilde{s} is common in the literature (for example, it also appears in Gans and Zhou 2002). When integrality is relaxed, the random variable $\tilde{z} \sim \text{Bin}(x, q)$ is understood as being defined by the corresponding moment-generating function $\mathbb{E}[\exp(\tilde{z}t)] = ((1 - q) + q \exp(t))^x$ and that $\text{esssup } \tilde{z} = x$. When x is large, integrality is less a concern—inaccuracies arising from the approximation are minimal.

Remark 2. The constraints in Problem (14) have the form

$$d \log(1 - q + qe^{\zeta/d}) \leq \xi,$$

when it is defined on $d > 0$. At $d = 0$, observe that $\lim_{d \downarrow 0} d \log(1 - q + qe^{\zeta/d}) = \max\{0, \zeta\}$, and the constraint should interpreted as $\zeta \leq \xi, 0 \leq \xi$ at $d = 0$.

Proposition 4 (Independence of Pipelines). Under Assumption 1, any two state variables $\tilde{s}_l^{t,\tau}$ and $\tilde{s}_{l'}^{t,\tau'}$, $l \neq l'$ and $\tau \neq \tau'$ in the same time t are independent, conditional on decisions $\{d_i^{t',\tau'} : t \in [T], \tau \in [M]_0\}$ in the previous time periods $t' \leq t$.

Proof. We relegate the proof to Appendix A. \square

Remark 3.

a. Notice that this result does not require independence across modeling time t . This is neither true in general nor required in Theorem 2.

b. Assumption 1 alone is insufficient for this proposition. The specific definition of $d_l^{t,0}$ as a decision variable is required.

Theorem 1 depends on the repeated application of pipeline invariance. The idea is that the functional form $\exp(\cdot)$ is preserved within the expectation, hence enabling us to evaluate $\mathbb{E}[\exp(\cdot)]$ repeatedly over time. In this process, it creates a nested series of relay functions ρ , which, being convex, can be represented as auxiliary variables ξ in epigraph form. This is illustrative of the concept of *pipelines*, which the stock in each grade l is aligned in:

$$\begin{aligned} \mathcal{P}_l^{0,0} &= \{s_l^{0,0}, \tilde{s}_l^{1,1}, \tilde{s}_l^{2,2}, \dots\} \\ \mathcal{P}_l^{0,1} &= \{s_l^{0,1}, \tilde{s}_l^{1,2}, \tilde{s}_l^{2,3}, \dots\} \\ &\vdots \\ \mathcal{P}_l^{1,0} &= \{s_l^{1,0}, \tilde{s}_l^{2,1}, \tilde{s}_l^{3,2}, \dots\} \\ \mathcal{P}_l^{2,0} &= \{s_l^{2,0}, \tilde{s}_l^{3,1}, \tilde{s}_l^{4,2}, \dots\} \\ &\vdots \end{aligned}$$

An employee belonging to a particular pipeline remains in the same pipeline across time. Attrition erodes the stock in the pipelines over time, and promotion redistributes across pipelines. Such an interpretation also explains why the independence result in Proposition 4 works—each state variable is a stochastic function of its predecessor in its pipeline—that is, retraced to its ancestor, either an initial condition or a decision variable.

The following results illustrate that considering individual variations in productivity as in (5) can be accepted under the Model Formulation (10).

Proposition 5. *The general productivity constraints in (5), where individual variations in productivity is considered, $C_k \left[\frac{1}{\theta_t^3} \left(P_t - \sum_{l,\tau} \tilde{r}_l^\tau(\tilde{s}_l^{t,\tau}) \right) \right] \leq 0$, has the equivalent form as a normal productivity constraint,*

$$C_k \left[\frac{1}{\theta_t^3} \left(P_t - \sum_{l,\tau} \tilde{s}_l^{t,\tau} r_l^\tau(k) \right) \right] \leq 0,$$

where

$$r_l^\tau(k) = \theta_t^3 k \log \left(\mathbb{E} \left[\exp \left(\frac{\tilde{r}_{li}^\tau}{\theta_t^3 k} \right) \right] \right).$$

Proof of Proposition 5. By Proposition 4, $\tilde{r}_l^\tau(\tilde{s}_l^{t,\tau})$ are independent random variables for different grade l and time-in-grade τ . As such, we have

$$C_k \left[\frac{1}{\theta_t^3} \left(P_t - \sum_{l,\tau} \tilde{r}_l^\tau(\tilde{s}_l^{t,\tau}) \right) \right] = P_t / \theta_t^3 - \sum_{l,\tau} C_k \left[\tilde{r}_l^\tau(\tilde{s}_l^{t,\tau}) / \theta_t^3 \right].$$

Observe that

$$\begin{aligned} C_k \left[\tilde{r}_l^\tau(\tilde{s}_l^{t,\tau}) / \theta_t^3 \right] &= k \log \mathbb{E} \left[\exp \left(\sum_{i \in [\tilde{s}_l^{t,\tau}]} \frac{\tilde{r}_{li}^\tau}{k \theta_t^3} \right) \right] \\ &= k \log \mathbb{E} \left[\exp \left(\theta_t^3 k \log \left(\mathbb{E} \left[\exp \left(\frac{\tilde{r}_{li}^\tau}{k \theta_t^3} \right) \frac{\tilde{s}_l^{t,\tau}}{\theta_t^3 k} \right] \right) \right) \right] \\ &= C_k \left[r_l^\tau(k) \tilde{s}_l^{t,\tau} / \theta_t^3 \right], \end{aligned}$$

and hence the results follow. \square

Theorem 2. *If integrality of \tilde{s}^t is relaxed, then constraints of the form,*

$$C_k \left[\left(\sum_{l,\tau} \tilde{s}_l^{t,\tau} u_l^\tau - U_t \right) / \theta \right] \leq 0, \quad (15)$$

may be reformulated as the convex set of constraints

$$\sum_l s_l^{t,0} u_l^0 + k \sum_{1 < l' \leq t} \xi_{l'}^{t',1} + k \sum_{\tau \geq t} \xi_l^{1,\tau-t+1} \leq U_t, \quad (16)$$

$$d_l^{t,\tau} \rho_l^\tau(u_l^\tau / k \theta) \leq \xi_{l'}^{t',\tau}, \quad \forall \tau \in [M],$$

$$d_l^{t',\tau} \rho_l^\tau(\xi_{l'}^{t'+1,\tau+1} / d_l^{t',\tau}) \leq \xi_{l'}^{t',\tau}, \quad \forall t' \in [t-1], \tau \in [M-t+t'].$$

Proof. Independence as a result of Proposition 4 allows the sum to be taken out of the certainty equivalence operator $C_k[\cdot]$, which can be evaluated using Theorem 1. \square

The remaining challenge is to deal with the dismissal threshold constraint. Thankfully:

Proposition 6 (Redistribution Constraint). *Under the same assumptions as Theorem 2, for fixed l , the constraint*

$$C_k \left[\left(\sum_{\tau} \tilde{s}_l^{t-1,\tau-1} \frac{d_l^{t-1,\tau-1} - d_l^{t,\tau}}{d_l^{t-1,\tau-1}} - d_{l+1}^{t,0} - F_{l+1}^t \right) / \theta \right] \leq 0, \quad (17)$$

is equivalent to the set of equations

$$\begin{aligned} d_l^{t-1,0} - d_l^{t,1} + k \sum_{1 < l' < t} \xi_{l'}^{t',1} + k \sum_{\tau \geq t-1} \xi_l^{1,\tau-t+2} &\leq F_{l+1}^t + d_{l+1}^{t,0} \\ d_l^{t-1,\tau} \rho_l^\tau \left(\frac{d_l^{t-1,\tau} - d_l^{t,\tau+1}}{k \theta d_l^{t-1,\tau}} \right) &\leq \xi_l^{t-1,\tau} \quad \forall \tau \in [M] \\ d_l^{t',\tau} \rho_l^\tau \left(\frac{\xi_{l'}^{t'+1,\tau+1}}{d_l^{t',\tau}} \right) &\leq \xi_l^{t',\tau} \quad \forall t' \in [t-2], \tau \in [M-t+t'+1]. \end{aligned} \quad (18)$$

Proof. The proof is similar to the proof of Theorem 2; as such, it is omitted. Again, note that for fixed k , the problem remains convex. \square

3.2. Solving the Model

To solve the SWPU model, one can perform a bisection search on k . For a fixed k , the feasible set of Problem (10) can be expressed as a conic optimization problem involving exponential cones.

Proposition 7. *The constraint of the form*

$$\begin{aligned} d \log(1 - q + qe^{\zeta/d}) &\leq \xi && \text{if } d > 0 \\ \zeta &\leq \xi, 0 \leq \xi && \text{if } d = 0, \end{aligned}$$

is equivalent to the following constraints:

$$\begin{aligned} (1 - q)y_1 + qy_2 &\leq d \\ (y_1, d, -\xi) &\in \mathcal{K}_E \\ (y_2, d, \zeta - \xi) &\in \mathcal{K}_E, \end{aligned} \tag{19}$$

for some $y_1, y_2 \in \mathbb{R}$, where the exponential cone is defined as

$$\mathcal{K}_E := \{(x_1, x_2, x_3) : x_1 \geq x_2 \exp(x_3/x_2), x_2 > 0\} \cup \{(x_1, 0, x_3) : x_1 \geq 0, x_3 \leq 0\}.$$

Proof. For $d > 0$, the nonlinear constraint can be expressed as

$$(1 - q)d \exp(-\xi/d) + qd \exp(\zeta - \xi/d) \leq d,$$

or equivalently as

$$\begin{aligned} (1 - q)y_1 + qy_2 &\leq d \\ d \exp(-\xi/d) &\leq y_1 \\ d \exp(\zeta - \xi/d) &\leq y_2, \end{aligned}$$

for some $y_1, y_2 \in \mathbb{R}$. We can also check that when $d = 0$, the constraints of (19) require $y_1 = y_2 = 0$ and $\xi \geq 0, \xi - \zeta \geq 0$. \square

It is well known that exponential cones constraints can be approximated with a series of second-order cones (see, for example, Ben-Tal and Nemirovski 2001). More recently, there have also been advances in the efficient computation of exponential cones, especially using interior point methods. A commercial solver, MOSEK ApS (2019), is among the first to include support for exponential cones. Solvers are already available in MATLAB (CVX Research 2012) and also in Julia/JuMP (e.g., Miles et al. 2016), extending to mixed-integer convex programs. Our model does not compromise tractability—the number of constraints does not grow exponentially with time horizon T or any of the other parameters, such as grades L or maximum time-in-grade M . Indeed, in Theorem 1, for each $t \in [T]$, the number of exponential cone constraints required to reformulate one linear constraint is of order $O(LMT)$. Hence, in total, $O(LMT^2)$ exponential cone constraints are required.

3.1.1. Cutting-Plane Approach. Even in the absence of nonlinear solvers, we can use the cutting-plane approach to solve the conic optimization problem. This approach has the advantage of keeping the model

linear, which has the benefit of having greater availability of solvers and incorporating discrete decision variables. In fact, we use this approach to compute the solutions.

We observe that for $q \in [0, 1]$, the function $\delta(d, \zeta) := d \log(1 - q + qe^{\zeta/d})$ is jointly convex and differentiable on $d > 0$, hence, for all $d > 0$ and $\zeta \in \mathbb{R}$. Hence, we can replace the nonlinear function by a maximum of an infinite set of affine functions,

$$\delta(d, \zeta) = \max_{d_0 > 0, \zeta_0} \{\delta(d_0, \zeta_0) + \partial\delta_1(d_0, \zeta_0)(d - d_0) + \partial\delta_2(d_0, \zeta_0)(\zeta - \zeta_0)\},$$

where $\partial\delta_1$ and $\partial\delta_2$ respectively, denote the partial derivatives of δ , with respect to its first and second argument. These affine functions can be introduced on the fly in a standard cutting-plane implementation. This approach works surprisingly well in our computational studies. Observe that $\delta(d, \zeta) \approx d \log(1 - q)$ as $\zeta \rightarrow -\infty$ and $\delta(d, \zeta) \approx \zeta + d \log q$ as $\zeta \rightarrow \infty$. As such, the behavior of $d \log(1 - q + qe^{\zeta/d}) \leq \xi$ is asymptotically linear with respect to ζ , alluding to why the cutting-plane method works well in practice.

Lastly, we comment that one can adapt the model, by using a different k_j for each constraint, indexed in a set $j \in \mathcal{J}$ and then performing a lexicographic minimization on $k := (k_j)_{j \in \mathcal{J}}$ (see Waltz 1967 on how to execute this procedure). This methodology may be employed if the decision maker is agnostic to the relative risk aversions of each constraint and would prefer the strongest performance achievable. In this paper, we hope to use θ_j to control the tightness of each constraint and to gather insights from how the *cost* of greater risk aversion in one constraint would affect other constraints. As such, we do not perform the lexicographic minimization in this paper. We shall see this at work later when we examine the flexibility of public-sector agencies in dismissing employees in Section 4.

4. Strategic Workforce Planning in a Firm

In this section, we illustrate the SWPU model using real data of > 5,000 employees in the Singapore Civil Service, who can be safely assumed to have similar job characteristics and backgrounds, tracked over six years. These data are collected periodically at the individualized level, and, for this illustration, we are able to summarize it into the form of the inputs for our model. This includes their attrition, performance, and wage patterns—personnel data that are similarly collected by most large organizations. Because of confidentiality, we are unable to reveal more about the nature of the data, although in the subsequent description, we will illustrate some features as far as we are able to share. In this illustration, we shall look at a five-year time window, $T = 5$.

We model $L = 4$ grades in this organization, two “individual contributor” grades labeled IC1 and IC2, which generate a large part of the productivity, and two manager grades, denoted M1 and M2. Progression occurs in this order, and skipping of grades is disallowed. We truncate the maximum number of years that an employee may remain in any grade to $M = 20$, where thereafter the employee is assumed to have retired. At each grade l , we assume that employees are paid a base wage ω_l with an annual fixed increment ι_l . Hence, $w_l^\tau = \omega_l + \tau\iota_l$. The parameters ω_l and ι_l were statistically estimated from wage data by means of a linear regression and rounded. Because of its sensitivity, we are unable to disclose these estimates.

To prescribe the productivity curve of the employees, we directly obtained the average productivity curve from the data themselves. Specifically, we grouped employees into each grade and (discrete) time-in-grade and then averaged their performances to obtain r_l^τ across time-in-grade. In this study, the performance data are measured as a composite of both quality and quantity of work done by the employee. They are ranked and also normalized across the workforce. Figure 4 shows the mean productivity by years-in-grade, as obtained from the data. It rises with more years-in-grade, a reflection of the accumulation of experience, before dipping with increasing employee boredom and disengagement. Manager span of control is also computed similarly.

Retention rates q_l^τ were estimated from the data. Figure 2 illustrates the estimates. Where the data were sparse, fluctuations were severe. Nonetheless, Proposition 1 provides the guarantee that even if we were to wrongly estimate the retention rates, we are still robust, as long as the true distribution is not far off from

Table 2. Specification of Constraints in Numerical Study

Constraint	Equation	Target	Specification
Headcount	$\sum_{l,\tau} s_l^{l,\tau} \leq H_t$	$H_t = g_h^t H_0$	$g_h = g$
Budget	$\sum_{l,\tau} s_l^{l,\tau} w_l^\tau \leq B_t$	$B_t = g_b^t B_0$	$g_b = g$
Productivity	$\sum_{l,\tau} s_l^{l,\tau} r_l^\tau \geq R_t$	$R_t = g_p^t R_0$	$g_p = 1 + 1.05(g - 1)$

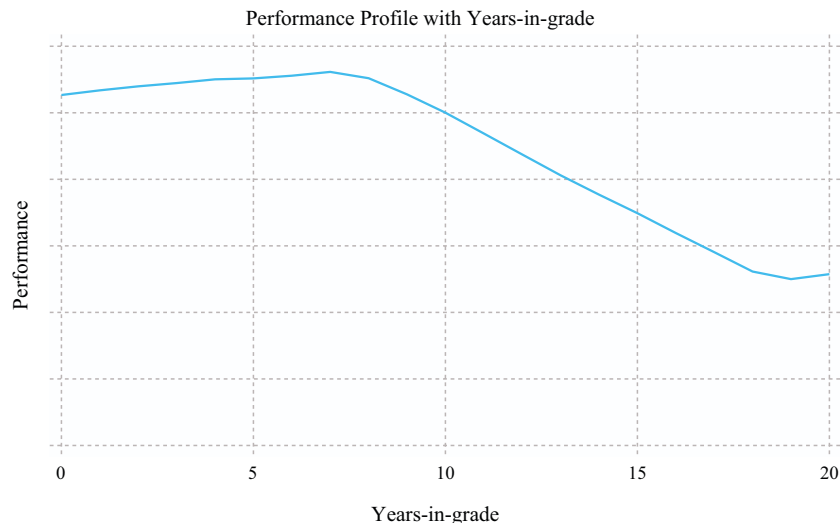
our estimate. Later, when analyzing the robustness of the model, we shall explore this further.

Finally, we specify the constraint targets. From here on, the targets shall always be fixed as a geometric rate of growth g from the initial state at time $t = 0$. We vary these rates of growth in different simulations. Table 2 summarizes this. We also require that the productivity target grows at a slightly faster rate than the headcount and budget targets. We set $F_l^t \equiv 0$ —that is, that zero dismissals is preferred.

Because the certainty equivalence C_k is not scale-invariant, we normalized all constraints so as to ensure equitable comparisons (without having to calibrate θ_j separately for each constraint). In other words, the model penalizes the proportional violation of targets equally across constraints.

With this specification, the model seeks to minimize the risk level, k . It returns k , in addition to optimal solutions for the decision variables of newcomers $s_l^{t,0}$ (from which we compute net inflow h_l^t) and promotions $d_l^{t,\tau}$. To simulate the uncertainty and test the model, for each analysis, we ran 1,000 simulations with the random outcomes of employees’ retention drawn from a binomial distribution of estimated retention rates, q_l^τ , as the success probability.

Figure 4. (Color online) Profile of Performance with Time-in-Grade



The model was solved by using the cutting-plane algorithm, as detailed in Algorithm 1. Because of the asymptotically linear structure of the conic constraints, the algorithm reaches within high accuracy very quickly. In computational tests, the model always solves within five minutes on an Intel® i7-6650U dual-core processor, with the worst constraint requiring no more than seven cutting planes to get within 10^{-5} accuracy of estimating the constraint.

4.1. Robustness

We first examine the robustness properties of the model. By design, the model provides guarantees against constraint violation. To illustrate this, we compare our model against a deterministic model. The deterministic model is not intended as a comparative benchmark to ascertain the strength of our model. Because our robust model recovers the deterministic model as $k \rightarrow \infty$, the latter is a guidepost for us to understand the degree of performance traded off for robustness. For the description of the deterministic model, please refer to Appendix A.3.

We first ran the model for growth rate $g = 1.02$ —that is, the organization is allowed to grow by 2% annually. Our model seeks the minimum risk level k^* . In this case, $k^* \approx 35$, which yields the exponential envelope of the probability of constraint violation.

In Figure 5, we plot, for the headcount constraint, the actual materialized deviation from target $H_t - \sum_{l,\tau} s_l^{t,\tau}$ based on the uncertainty. A positive figure indicates that the headcount target was not exceeded, and its magnitude gives the slack; a negative value

Table 3. Comparison of Constraint Violation in Robust and Deterministic Models

Deviation from H_5	Robust	Deterministic
Median	8.07	-4.40
Mean	7.18	-5.57
1 st Quartile	-14.29	-27.42

indicates constraint violation and its magnitude, the extent. The line represents the Markov guarantee, where the probability of constraint violation should be no more than one-third. As Figure 5 illustrates, this guarantee is very loose—only 2% of the simulations exceeded this bound.

We now compare this against the deterministic model. The simulated deviations from the headcount target for each model are compared in Table 3. Our model provides guarantees against constraint violation, and if violations occur, they do so with a smaller magnitude.

However, one can expect that the gains in the guarantees may not be universal for different specification of the targets. To illustrate this, let us vary the productivity target P_t (via g_p), while keeping all other targets fixed. Intuitively, there should be a monotone relationship between P_t and k^* —the higher P_t , that is, the higher the productivity target that must be met, the more difficult it is to do so, and hence, the risk level k^* of failing should be expected to rise. We try this for three configurations: $g_p = 1.023$ (where k^* is large), $g_p = 1.021$ (an intermediate region), and $g_p = 1.015$ (where k^* is small). Table 4 summarizes the statistics

Figure 5. (Color online) Simulated Violation of Headcount Target in Year $t = 1$

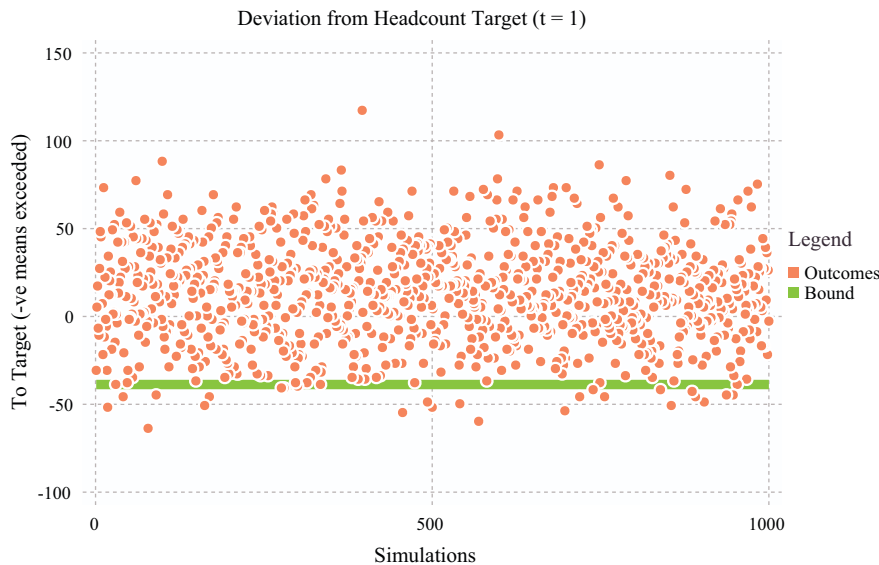


Table 4. Different Regimes of Tightness of Targets

Violation of constraint	Tougher target $g_p = 1.023(k^* \approx 232)$		Intermediate $g_p = 1.021(k^* \approx 35)$		Easier target $g_p = 1.015(k^* \approx 10)$	
	Robust	Deterministic	Robust	Deterministic	Robust	Deterministic
Growth rate (risk level)						
Deviation from P_5						
Median deviation	-4.09	5.43	30.19	96.70	127.10	379.36
Mean deviation	-5.16	4.14	29.13	97.63	123.16	379.07
1 st Quartile	-43.7	-28.60	-1.74	65.59	88.39	344.40
Deviation from H_5						
Median deviation	17.95	15.81	41.64	18.39	111.18	14.51
Mean deviation	18.91	16.59	42.02	18.45	113.48	13.92
1 st Quartile	-14.30	-14.19	12.64	-12.61	82.18	-16.49

for these three regimes, under a comparison between the robust and deterministic models.

In Table 4, we compare two constraints. The first is the productivity constraint at time T . This was the objective in the deterministic model, and, hence, we should reasonably expect the deterministic model to outperform the robust model in all instances. Their difference can be understood as the price to pay for robustness. In the second, we compare the headcount constraint. Across the different regimes, because the deterministic model is agnostic to the risk of constraint violation, the distribution in the deviation from the headcount target is approximately the same. Here, however, we can see the action of the robust model.

We observe three kinds of scenarios:

1. When k^* is very large (first regime), the problem is near infeasible. In this case, the guarantees on constraint violation erode away, and the robust model approximates the deterministic model. The guarantees are so minimal that it is effectively suboptimal. In other words, when the system is near its limits of operability, robustness is a luxury that cannot be afforded.

2. When k^* is very small, we are in the third regime. Here, the guarantees are very sharp—so sharp, it is overconservative. As seen in Table 4, the loss in productivity is sizeable. On the other hand, the deterministic model is not without reproach—a huge trade-off between headcount and productivity was made, by virtue of the fact that productivity was the objective. A reasonable course of action at this point is to tighten the target.

3. There is an intermediate region where the trade-off is balanced to some extent. In the second regime, the robust model does not incur a large cost to productivity, yet provides reasonable guarantees against constraint violation.

In the last segment of this analysis on robustness, we examine whether the optimal solution is robust to the input parameters of attrition rates. To do so, we

smooth the attrition rates (one minus the retention rates in Figure 2) using a Loess regression and prune negative values. The smoothed attrition rates are shown in Figure 3. Here, points represent the raw estimates and lines the smoothed outcome. We then perform the same analysis we have done before.

With smoothing, the risk level rises to $k^* \approx 44$ from the previous $k^* \approx 35$. Additionally, we also examine the optimal policies for promotion (in Figure 6, which can be compared against the optimal policy without smoothing in Figure 7) and hiring (in Figure 8, where the original policy is in points and the smoothed version is lined). We can see that there are only slight differences between the optimal policies suggested by the two models.

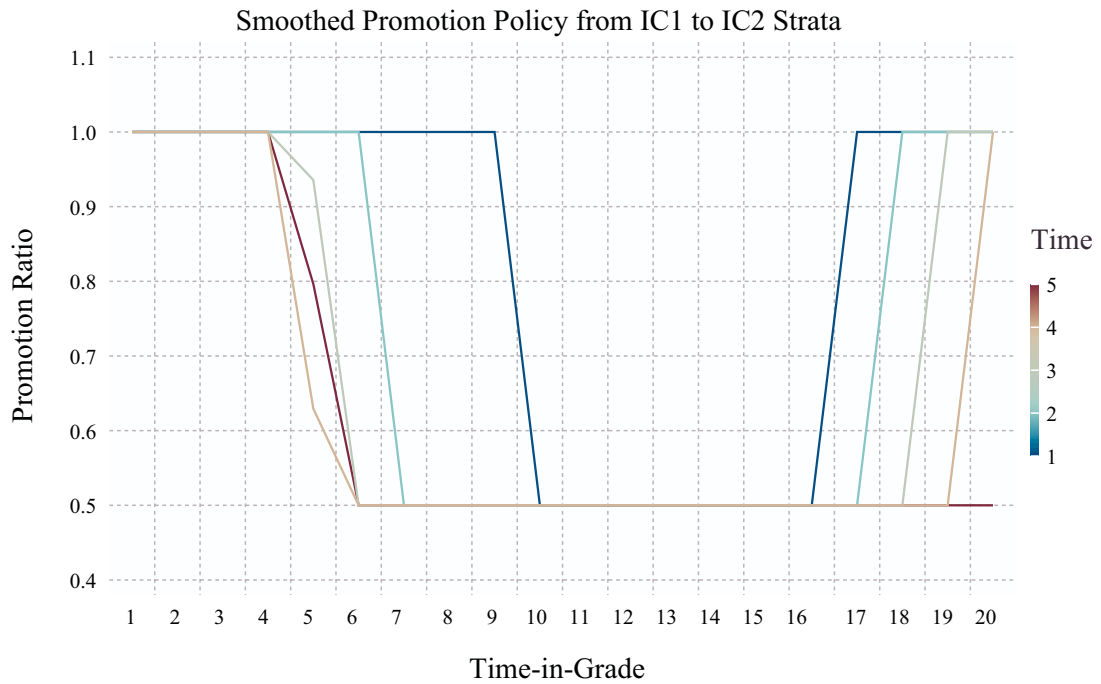
4.2. Time-Based Progression

In this section, we examine insights that can be gleaned for HR. In the first instance, we are interested in the question: When is it optimal to promote employees? In other words, how long should I keep an employee at a particular grade before promoting them?

We study $p_l^{t,\tau}$. Recall that $p_l^{t,\tau}$, which is equivalent to the ratio $d_l^{t+1,\tau+1}/d_l^{t,\tau}$, is the proportion of employees at time t whom we retain at grade l for an additional year, having already spent τ years at this grade l . As such, the closer this ratio is to one, the fewer employees we are promoting. For the purposes of fairness and continuity, HR would set a limit to the maximum proportion of employees at a grade that may be promoted in any year. Our partners do not wish for their limit to be shared. As such, for illustrative purposes, we have chosen the bar of 50%: $d_l^{t,\tau}/d_l^{t-1,\tau-1} \geq 0.5$. Figure 7 shows the policy for progressing employees at grade IC1 to grade IC2, as prescribed by the SWPU model.

The prescribed policy is a threshold—the model believes that employees should not progress to the next grade until they have accumulated a minimum number of years, after which they should be promoted with

Figure 6. (Color online) Policy for Progressing from IC1 to IC2 Under Smoothing



haste. There is a certain logic in this. In the early years, the productivity of employees rises with time spent in that grade due to the learning curve (Figure 4). As such, promoting employees too early incurs an opportunity cost of potential productivity. The model avoids this. After some point, remaining for too long at the same grade can have a disengaging effect on employees, and

they may leave the organization (Figure 2). The model also avoids this, by expediting their promotion after some time. In other words, the model seeks a balance between the productivity an employee brings and the risk of losing the employee. This finding lends numerical support not just to the choice of “time-based progression,” but also its rationale.

Figure 7. (Color online) Policy for Progressing from IC1 to IC2 at Optimal Solution


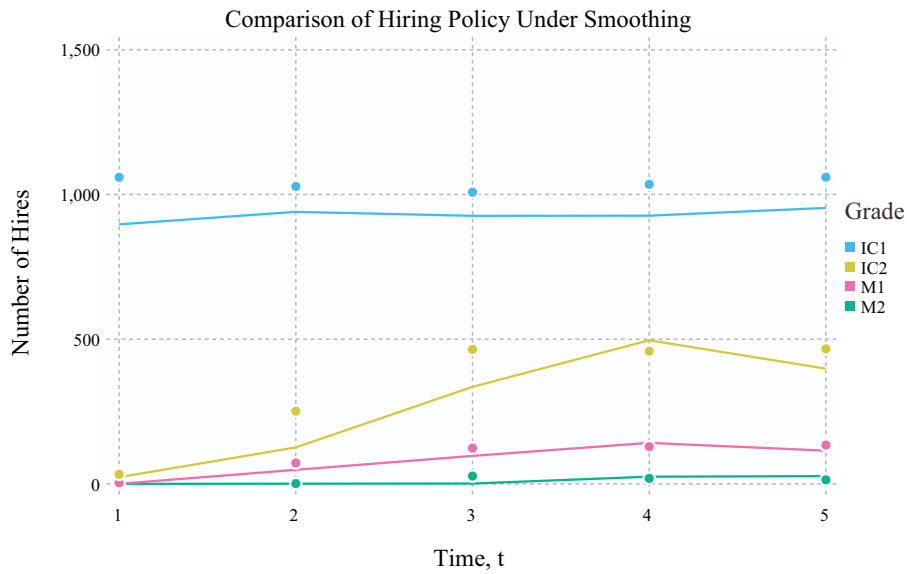


Figure 8. (Color online) Policy for Hiring Across Grade and Time Under Smoothing



4.3. Factors Affecting the Risk Level

In this second piece of analysis, we shall examine the impact that the growth rate g has on the optimal risk level k^* . As before, we fixed the allowed growth rates of headcount, budget, and productivity to be a function of g . Now, we vary g . Figure 9 plots the relationship.

From Figure 9, we infer that there is a higher risk level when the growth rate is smaller. This mirrors common wisdom that it is easier to grow firms than to downsize. The explanation the model gives is this: Risk originates from the uncertainty—resignations.

The higher the growth rate g , the greater the number of new recruits. Recruitment of employees fills the vacancies created by those who left and, thus, mitigates the uncertainty. As such, the more employees that can be recruited, the larger the capacity of HR to manage the risks arising from resignation, and, thus, the lower the risk on the overall.

The simple consequence of this is that there are inherent operational risks to a lack of organizational renewal. Yet, this is not necessarily a straightforward question to address. For example, in an organization

Figure 9. (Color online) Risk Level k^* at Different Growth Rates g


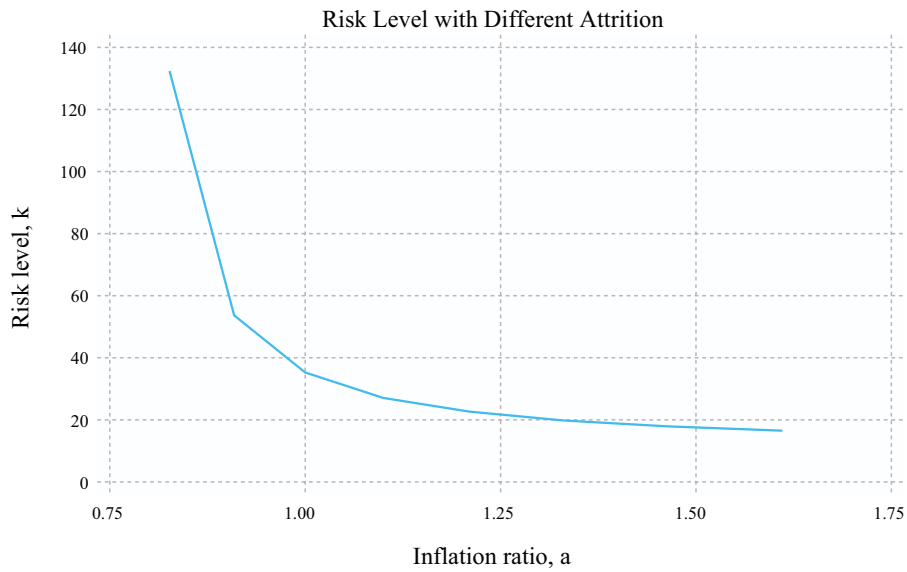


Figure 10. (Color online) Risk Level k^* with Different Scaling of Attrition a



with a higher attrition level, we can expect two competing forces at work. One, the higher the attrition, the greater the uncertainty, and, hence, the higher the risk. Two, the higher the attrition, the greater the capacity to hire because there are more vacancies to replace; hence, the lower the risk. We study which effect really plays out in our data set.

In our model, q_l^τ represents the retention rate of officers having spent time τ at grade l . Hence, the attrition rate is $\alpha_l^\tau = 1 - q_l^\tau$. We now artificially suppress or boost the attrition rate by a factor of a , via $\bar{\alpha}_l^\tau = a \cdot \alpha_l^\tau$.

If $a < 1$, the attrition rate is suppressed, and vice versa. As such, we have a new $\bar{q}_l^\tau = 1 - \bar{\alpha}_l^\tau$.

Figure 10 plots what happens to the risk level k^* as we vary a . With lower attrition, the risk level k^* rises. This is a grim consequence for advanced economies, where an aging population is beginning to take hold. As older employees are often less employable in the workforce, they tend to move between organizations less frequently than younger employees. As such, with an aging population, firms can expect to see attrition rates fall across the board. Instead, they will be

Figure 11. (Color online) Risk Level k^* with Tightening of Hiring Constraint

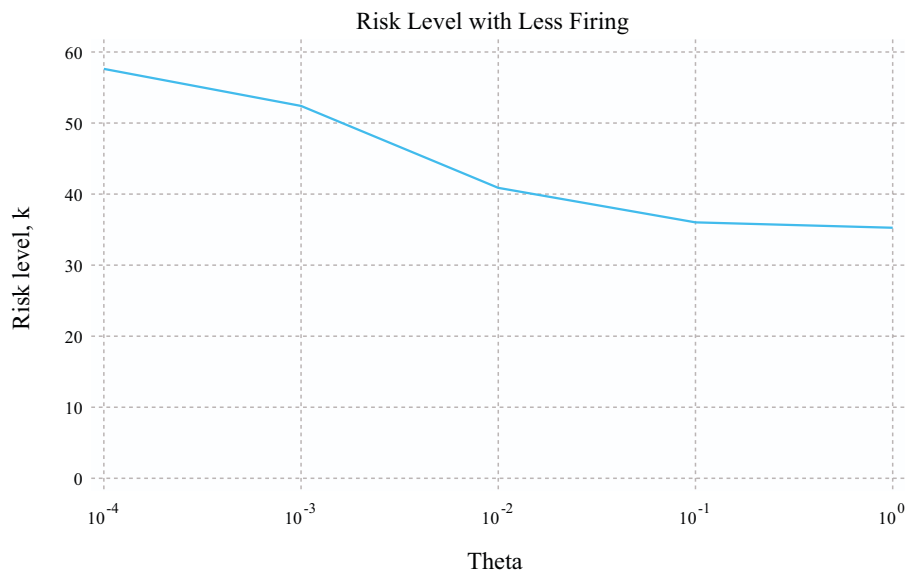



Figure 12. (Color online) Largest Number of Employees Released in any of the 1,000 Simulations



faced with ever-rising challenges in managing their workforce. This is not to mention that the shrinking workforce would force many firms to reduce their growth rates, which further heightens the risk.

For the final insight, we look at varying the tightness of a constraint. Specifically, we shall examine the importance of organizational renewal. For our partners, they were interested to see whether the inherent difficulty of public-sector organizations to lay off their employees—and, hence, a limited capacity for organizational renewal—would result in greater difficulties in managing their workforce and, if so, how large these difficulties are. To elucidate this, we perform the following analysis. Recall that we could calibrate θ_j to dictate the tightness of the bounds for the corresponding constraint j . Now, we do so for the dismissal threshold constraint. The lower the value of θ_j , the more averse the model is to releasing employees. Figures 11 and 12 tell us the consequences of this.

In Figure 11, we can see that the difference between not allowing any and allowing some dismissals is an almost doubling of the risk level. Figure 12 illustrates the trade-off. We plot here the largest number of employees released amongst the 1,000 simulations. If this number is negative, it means that in all the simulations, there wasn't a single case where an employee was released. At risk level $k^* \approx 35$ and $\theta = 1$, about 40 employees were released in total across the grades. If the decision maker is to refrain from releasing any of these employees, then θ must be decreased to 10^{-3} . This would incur an almost 50% increase in the risk level.

This quantifies the natural challenges faced by public-sector organizations compared with their private-sector

counterparts. In this regard, it is therefore paramount that public agencies find new and innovative ways to rejuvenate and renew their workforce.

5. Conclusions

We have presented a tractable model for workforce planning. Although we illustrate our model on data from a public agency, the model can still be utilized in some profit-seeking firms. We have also illustrated HR insights and provided numerical quantification of such risks that firms can face, such as the need for organizational renewal.

At its root, the SWPU model is an application of the concept of pipeline invariance under the context of multiperiod optimization. The general intuition is that, although it is difficult to perform multiperiod optimization, we may alleviate these difficulties if we declare a formal structure (here, pipeline invariance) on how the decisions and the uncertainty are related and, hence, exploit this structure to gain tractable formulations. On this note, we hope, in the future, to construct a formal framework for using pipeline invariance in multiperiod optimization problems.

Acknowledgments

Special thanks go out to Mr. Low Xin Wei and Mr. Eugene Chua for the inception, practitioner insights, and the data set, without which the computational tests and analyses would have been impossible. At the behest of Singapore Ministry of Education, the authors qualify that the analyses in this paper are not a reflection of the position, intent, or opinions of the Singapore Government, but are solely the opinions of the authors of the paper. The authors also thank the review team for their insightful comments, which have greatly improved their paper.

Appendix A. Proofs Omitted from the Main Text

A.1. Proof of Theorem 1

The idea is to show by induction on $\tau' = 1, \dots, \min\{t, \tau\}$ that the following relationship holds:

$$\begin{aligned} C_k[ky\tilde{s}_i^{t,\tau}/d_i^{t,\tau}] &= \inf_{\xi} C_k\left[k\xi_i^{t-\tau'+1, \tau-\tau'+1}\tilde{s}_i^{t-\tau', \tau-\tau'}/d_i^{t-\tau', \tau-\tau'}\right] \\ \text{s.t. } d_i^{t-\tau'+1, \tau-\tau'+1}\rho_i^\tau(y/d_i^{t-\tau'+1, \tau-\tau'+1}) &\leq \xi_i^{t-\tau'+1, \tau-\tau'+1} \forall i \in [\tau']. \end{aligned} \quad (\text{A.1})$$

Here, we have abused the notation slightly, as the C_k operator on the left-hand side is an expectation over all uncertain $\tilde{s}_i^{v,\tau}$ until $v \leq t$, whereas the right-hand side is only up till times $v \leq t - \tau'$.

First, notice that the induction step going from τ' to $\tau' + 1$ is inherent from the form of (A.1)—simply take $y = \xi_i^{t-\tau'+1, \tau-\tau'+1}$. As such, it suffices to prove only the step $\tau' = 1$. We evaluate as follows:

$$\begin{aligned} C_k[ky\tilde{s}_i^{t,\tau}/d_i^{t,\tau}] &= k \log \mathbb{E}\left[\exp\left(y\tilde{s}_i^{t,\tau}/d_i^{t,\tau}\right)\right], \\ &= k \log \mathbb{E}_{\leq t-1}\left[\mathbb{E}_t\left[\exp\left(y\tilde{s}_i^{t,\tau}/d_i^{t,\tau}\right)\right]\right], \quad (\text{A.2}) \\ &= k \log \mathbb{E}_{\leq t-1}\left[\exp\left(\tilde{s}_i^{t-1, \tau-1} \frac{d_i^{t,\tau}}{d_i^{t-1, \tau-1}} \rho_i^\tau\left(\frac{y}{d_i^{t,\tau}}\right)\right)\right]. \end{aligned} \quad (\text{A.3})$$

Here, we have used iterated expectations in (A.2) and then pipeline invariance in (A.3). At this point, notice that $d_i^{t,\tau}\rho_i^\tau(y/d_i^{t,\tau})$ is jointly convex in both y and $d_i^{t,\tau}$, as ρ_i^τ is a convex function. As such, we may represent it in the epigraph format $d_i^{t,\tau}\rho_i^\tau(y/d_i^{t,\tau}) \leq \xi_i^{t,\tau}$. Hence, this proves the $\tau' = 1$ case.

When $\tau' = \min\{t, \tau\}$, we achieve the result in the theorem. \square

A.2. Proof of Proposition 4

Observe that two state variables $\tilde{s}_i^{t,\tau}$ and $\tilde{s}_{i'}^{t,\tau}$, $l \neq l'$ and $\tau \neq \tau'$ in the same time t can be associated with different sets of employees that do not overlap. Hence, under Assumption 1, the random states should also be independent. We shall do so by induction on $t \geq 0$. When $t = 0$, this is trivially true, because s^0 are the initial conditions. Suppose any two different $\tilde{s}_i^{t,\tau}$ and $\tilde{s}_{i'}^{t,\tau}$ are independent. First, consider $\mathbb{E}\left[\mathbb{1}_{\{\tilde{s}_i^{t+1,\tau} \leq i\}} \mathbb{1}_{\{\tilde{s}_{i'}^{t,\tau} \leq j\}}\right]$. If $\tau = 0$, then $\tilde{s}_i^{t+1,\tau}$ is a decision variable; hence, this is trivially true. Suppose now that $\tau \geq 1$ and that if $l' = l$, then $\tau' \neq \tau - 1$,

$$\mathbb{E}\left[\mathbb{1}_{\{\tilde{s}_i^{t+1,\tau} \leq i\}} \mathbb{1}_{\{\tilde{s}_{i'}^{t,\tau} \leq j\}}\right] = \mathbb{E}\left[\mathbb{E}\left[\mathbb{1}_{\{\tilde{s}_i^{t+1,\tau} \leq i\}} \mathbb{1}_{\{\tilde{s}_{i'}^{t,\tau} \leq j\}} \middle| \mathbb{1}_{\{\tilde{s}_i^{t,\tau-1} \leq i'\}}\right]\right], \quad (\text{A.4})$$

$$= \mathbb{E}\left[\mathbb{1}_{\{\tilde{s}_{i'}^{t,\tau} \leq j\}} \mathbb{E}\left[\mathbb{1}_{\{\tilde{s}_i^{t+1,\tau} \leq i\}} \middle| \mathbb{1}_{\{\tilde{s}_i^{t,\tau-1} \leq i'\}}\right]\right], \quad (\text{A.5})$$

$$= \mathbb{E}\left[\mathbb{1}_{\{\tilde{s}_{i'}^{t,\tau} \leq j\}}\right] \mathbb{E}\left[\mathbb{E}\left[\mathbb{1}_{\{\tilde{s}_i^{t+1,\tau} \leq i\}} \middle| \mathbb{1}_{\{\tilde{s}_i^{t,\tau-1} \leq i'\}}\right]\right], \quad (\text{A.6})$$

$$= \mathbb{E}\left[\mathbb{1}_{\{\tilde{s}_i^{t+1,\tau} \leq i\}}\right] \mathbb{E}\left[\mathbb{1}_{\{\tilde{s}_{i'}^{t,\tau} \leq j\}}\right]. \quad (\text{A.7})$$

The line (A.5) follows because of the independence between $\tilde{s}_{i'}^{t,\tau}$ and $\tilde{s}_i^{t,\tau-1}$, as assumed in the induction hypothesis; and

the Equation (A.6) follows because $\mathbb{E}\left[\mathbb{1}_{\{\tilde{s}_i^{t+1,\tau} \leq i\}} \mathbb{1}_{\{\tilde{s}_i^{t,\tau-1} \leq i'\}}\right]$ is just a function of $\tilde{s}_i^{t,\tau-1}$ due to the dynamics defined in (11), and, thus, independence again allows the splitting of expectations. Now, we perform the next step. Again, similar logic applies if $\tau = 0$; otherwise,

$$\mathbb{E}\left[\mathbb{1}_{\{\tilde{s}_i^{t+1,\tau} \leq i\}} \mathbb{1}_{\{\tilde{s}_{i'}^{t+1,\tau'} \leq j\}}\right] = \mathbb{E}\left[\mathbb{E}\left[\mathbb{1}_{\{\tilde{s}_i^{t+1,\tau} \leq i\}} \mathbb{1}_{\{\tilde{s}_{i'}^{t+1,\tau'} \leq j\}} \middle| \mathbb{1}_{\{\tilde{s}_i^{t,\tau-1} \leq i'\}}\right]\right], \quad (\text{A.8})$$

$$= \mathbb{E}\left[\mathbb{1}_{\{\tilde{s}_{i'}^{t+1,\tau'} \leq j\}} \mathbb{E}\left[\mathbb{1}_{\{\tilde{s}_i^{t+1,\tau} \leq i\}} \middle| \mathbb{1}_{\{\tilde{s}_i^{t,\tau-1} \leq i'\}}\right]\right], \quad (\text{A.9})$$

$$= \mathbb{E}\left[\mathbb{1}_{\{\tilde{s}_{i'}^{t+1,\tau'} \leq j\}}\right] \mathbb{E}\left[\mathbb{E}\left[\mathbb{1}_{\{\tilde{s}_i^{t+1,\tau} \leq i\}} \middle| \mathbb{1}_{\{\tilde{s}_i^{t,\tau-1} \leq i'\}}\right]\right], \quad (\text{A.10})$$

$$= \mathbb{E}\left[\mathbb{1}_{\{\tilde{s}_i^{t+1,\tau} \leq i\}}\right] \mathbb{E}\left[\mathbb{1}_{\{\tilde{s}_{i'}^{t+1,\tau'} \leq j\}}\right], \quad (\text{A.11})$$

where (A.9) follows because of the independence between $\tilde{s}_{i'}^{t+1,\tau'}$ and $\tilde{s}_i^{t,\tau-1}$, as proven in the previous step, and similarly for (A.10).

A.3. Description of the Deterministic Model

We write the deterministic model below. We shall take productivity in the last time period (P_T) as the objective. Note that from Proposition 3, the deterministic model is obtained from the robust formulation in the limit $k \rightarrow \infty$.

$$\begin{aligned} \max \quad & \sum_{l,\tau} \gamma_l^\tau d_i^{t,\tau} r_l^\tau \\ \text{s.t.} \quad & \sum_{l,\tau} \gamma_l^\tau d_i^{t,\tau} \leq H_t \quad \forall t \in [T] \\ & \sum_{l,\tau} \gamma_l^\tau d_i^{t,\tau} w_l^\tau \leq B_t \quad \forall t \in [T] \\ & \sum_{l,\tau} \gamma_l^\tau d_i^{t,\tau} r_l^\tau \geq P_t \quad \forall t \in [T] \\ & \sum_{\lambda,\tau} \gamma_\lambda^\tau d_\lambda^{t,\tau} b_{l,\lambda}^\tau \leq 0 \quad \forall t \in [T], \forall l \in \mathcal{M} \\ & \sum_{\tau} \gamma_l^{\tau-1} (d_i^{t-1, \tau-1} - d_i^{t,\tau}) \leq d_{i+1}^{t,0} + F_{i+1}^0 \quad \forall t \in [T], \forall i \in [L-1] \\ & d_{L+1}^{t,0} = 0, 0 \leq d_i^{t,\tau} \leq d_i^{t-1, \tau-1}, d_i^{0,\tau} = s_i^{0,\tau} \quad \forall t \in [T], \forall i \in [L], \forall \tau \in [\mathcal{M}] \end{aligned} \quad (\text{A.12})$$

Appendix B. Extensions and Generality of the Model

In this subsection, we discuss how the model could be applied to two aspects, the first regarding its connection to employee archetypes and clustering and the second regarding departmental transfers. These extensions are made possible because our model permits the categorization of employees under some index set $i \in \mathcal{I}$, by which we meant that we can append the index i to all of our state variables $\tilde{s}_{i,i}^{t,\tau}$, decision variables $d_{i,i}^{t,\tau}$, etc. This is akin to building many copies of the organizational structure that do not intersect, but where they operate under common organizational target-meeting constraints—for example, headcount $\sum_i \sum_{l,\tau} \tilde{s}_{i,i}^{t,\tau} \leq H^t$.

In the most general setting, i could represent each employee, both existing and potential future hires. In this setting, however, all of the parameters would have to be known. These include the likelihood of resignations $1 - q_{i,i}^\tau$ and productivities $r_{i,i}^\tau$. This might be more consistent with

existing paradigms in predictive analytics, where one is able to establish predictions for the likelihood of resignation at different times-in-grade for every employee with different covariates—for example, given a data set of side information of employees X_i , one constructs the predictive model $q_{i,i}^{t,\tau} \sim q_i^t | X_i$, such as via a survival model, or $q_{i,i}^{t,\tau} \sim q_i | X_i, \tau_i$, through, say, a classification methodology, like a random forest. Such side information X_i could contain information of employees on their demographics, work environment and outcomes, or job nature, to name a few. Similar may be said about the productivities, where machine learning could be applied to predict the productivity trajectories of employees, or via a learning curve (which we shall illustrate in the later subsection).

The drawback of such an approach would be the high dimension of the optimization problem to be solved, as introduced by the index i , which now increases the complexity by the scale of the number of employees. Also, there are some subtleties with dealing with how the inflow could be modeled.

Instead, the more reasonable approach might be to group employees into clusters. Suppose now that employees are instead labeled as $e \in \mathcal{E}$; then, we can consider a partitioning of the employees into subsets $\mathcal{E} = \cup_i \mathcal{E}_i$. These subsets can be constructed either as a demarcated grouping—for example, departments—or via data-driven methods—for example, clustering on side information of the employees $i | X_e$.

B.1. Clustering on Employee Archetypes

Many organizations understand their employees along the lines of employee archetypes. In this first application, we can consider $i \in \mathcal{I}$ to represent an employee archetype. In a broader sense, the concept of using archetypes to understand retention, performance, and hiring preferences is well established. Moving forward, there will be greater application of data-driven methods to do so, where the archetypes of employees are constructed through employee-side information. Already, the practice of utilizing latent class analysis to construct archetypes in the domain of HR has existed (e.g., in Perelman et al. 2019). With the onset of analytics, it is increasingly popular to perform clustering (or any dictionary learning algorithms) in order to construct the archetypes $i \in \mathcal{I}$. These archetypes can have a very high accuracy in predicting the retention of employees $q_{i,i}^t$ or their level of performance $r_{i,i}^t$. Our model fully supports such approaches.

B.2. Modeling Departmental Transfers

Suppose i represents the department that the employee is in. Recall in our Dynamics (3), we had used $\tilde{s}_i^{t-1,\tau-1} \frac{d_i^{t-1,\tau-1} - d_i^{t,\tau}}{d_i^{t-1,\tau-1}}$ to model

the number of employees removed from grade l , where $s_{i+1}^{t,0}$ represents the number promoted and the difference those fired (after subtracting for the new hires). Let us introduce the new notation $\beta_{i,j,l}^{t,\tau} / (d_i^{t-1,\tau-1} - d_i^{t,\tau})$ to represent the proportion of employees who are removed from grade l stipulated for transfer from department i into another department j . Table B.1 illustrates the count of all employees under this notation.

Notice from Table B.1 that the number of employees retained and those transferred from department i to j are in the tractable form required for Theorem 2. Similarly, we have the nonpositive constraint applied on the officers fired. When evaluating this under the entropic risk operator, by independence, each component of the sum will split, and we will arrive at a tractability result similar to Proposition 6. This illustrates how the model may be extended to departmental flows without losing tractability.

B.3. Describing Productivity Using Learning Curves

In this subsection, we illustrate specifically how learning curves may be used within the framework of our model.

Let each employee be denoted by the index $e \in \cup_l \mathcal{E}_l$, where \mathcal{E}_l is the collection of all employees in the grade l . Assume that we have longitudinal past performance data of each employees at time point t and grade l , denoted $R_e(l, t)$. Then, the learning-curve approach involves fitting the observed data $R_e(l, t)$ to the following form (assuming that of Shafer et al. 2001):

$$r_{(e)}^\tau = \eta_l \left(\frac{\tau + v_e}{\tau + v_e + \psi_e} \right), \quad (\text{B.1})$$

where τ is the time-in-grade; η_l is the maximum asymptotic productivity that can be reached for the grade l , in which employee $e \in \mathcal{E}_l$ is in; v_e is employee e 's prior experience; and ψ_e a term idiosyncratic to employee e .

In this situation, we can compile the productivity estimate for the grade by taking the average estimated productivities:

$$r_l^\tau = \frac{1}{|\mathcal{E}_l|} \sum_{e \in \mathcal{E}_l} r_{(e)}^\tau = \frac{1}{|\mathcal{E}_l|} \sum_{e \in \mathcal{E}_l} \eta_l \left(\frac{\tau + v_e}{\tau + v_e + \psi_e} \right). \quad (\text{B.2})$$

If in the case, where we are assuming that each employee segment i is precisely the employees themselves, such as in the earlier discussion at the start of this section, then the learning curves may be directly used: $r_{i,e}^\tau := r_{(e)}^\tau$. As a brief remark, we note that in order to supply the model enough information, we would also require the counterfactual estimates of the employee's performance at higher grades, $l' > l$. This can be obtained in two ways. The first is to find a group of closely matching employees, in terms of their side information, and to consider a weighted sum

Table B.1. Illustration of the Flows under a Departmental Transfer Model

From department i to j	Promoted	Retained	Fired	Total
$\tilde{s}_{i,l}^{t-1,\tau-1} \frac{\beta_{i,j,l}^{t,\tau}}{d_i^{t-1,\tau-1}}$	$s_{i,l+1}^{t,0} - h_{i,l+1}^t - \sum_j \tilde{s}_{j,l+1}^{t-1,\tau-1} \frac{\beta_{j,i,l+1}^{t,\tau}}{d_{l+1}^{t-1,\tau-1}}$	$\tilde{s}_i^{t-1,\tau-1} \frac{d_i^{t,\tau}}{d_i^{t-1,\tau-1}}$	$\tilde{s}_{i,l}^{t-1,\tau-1} \frac{d_i^{t-1,\tau-1} - d_i^{t,\tau} - \sum_j \beta_{i,j,l}^{t,\tau}}{d_i^{t-1,\tau-1}} - \left(s_{i,l+1}^{t,0} - h_{i,l+1}^t - \sum_j \tilde{s}_{j,l+1}^{t-1,\tau-1} \frac{\beta_{j,i,l+1}^{t,\tau}}{d_{l+1}^{t-1,\tau-1}} \right)$	$\tilde{s}_{i,l}^{t-1,\tau-1}$

of their productivities as a means of imputing this estimate. The second is to examine the history of past employees and to estimate the amount of change that the accumulated prior experience v_e and idiosyncratic term ψ_e would change under promotion. Whichever the case, as our model is constructed to be robust to wrong estimation in the parameters, the errors in estimation would be mitigated by the optimization procedure.

Endnotes

¹ Beames C (2015) Workforce strategy audit survey report, KPMG. Accessed January 4, 2021, <https://assets.kpmg/content/dam/kpmg/pdf/2015/08/workforce-strategy-audit-survey-report.pdf>.

² Talent | Organization | McKinsey & Company. Accessed October 5, 2020, <https://www.mckinsey.com/business-functions/organization/how-we-help-clients/talent>.

³ Mercer Strategic workforce planning. Accessed October 5, 2020, <https://www.mercer.com/content/dam/mercer/attachments/global/Talent/Forecast-brochureStratWrkfcPlan.pdf>.

⁴ Gartner. Strategic workforce planning. Accessed October 5, 2020, <https://www.gartner.com/en/human-resources/insights/workforce-planning>.

⁵ Deloitte. Strategic workforce planning. Accessed October 5, 2020, https://www2.deloitte.com/content/dam/Deloitte/za/Documents/audit/za_strategic_workforce_planning_012019.pdf; available from <https://www2.deloitte.com/za/en/pages/audit/solutions/strategic-workforce-planning.html>.

References

- Abernathy WJ, Baloff N, Hershey JC, Wandel S (1973) A three-stage manpower planning and scheduling model—a service-sector example. *Oper. Res.* 21(3):693–711.
- Arlotto A, Chick SE, Gans N (2014) Optimal hiring and retention policies for heterogeneous workers who learn. *Management Sci.* 60(1):110–129.
- Aumann RJ, Serrano R (2008) An economic index of riskiness. *J. Polit. Econom.* 116(5):810–836.
- Bartholomew DJ, Forbes AF, McClean SI (1991) *Statistical Techniques for Manpower Planning*, 2nd ed. (Wiley, Chichester, UK).
- Ben-Tal A, Nemirovski A (2001) *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications*, MOS-SIAM Series on Optimization (Society for Industrial and Applied Mathematics, Philadelphia).
- Bres ES, Burns D, Charnes A, Cooper WW (1980) A goal programming model for planning officer accessions. *Management Sci.* 26(8):773–783.
- Brown DB, Sim M (2008) Satisficing measures for analysis of risky positions. *Management Sci.* 55(1):71–84.
- Burke EK, De Causmaecker P, Berghe GV, Van Landeghem H (2004) The state of the art of nurse rostering. *J. Scheduling* 7: 441–499.
- Buyens D, De Vos A (2001) Perceptions of the value of the HR function. *Human Resource Management J.* 11(3):70–89.
- Charnes A, Cooper W (1963) Deterministic equivalents for optimizing and satisficing under chance constraints. *Oper. Res.* 11(1): 18–39.
- Chung SH, Jung DC, Yoon SN, Lee DH (2010) A dynamic forecasting model for nursing manpower requirements in the medical service industry. *Service Bus.* 4:225–236.
- CVX Research (2012) CVX: Matlab software for disciplined convex programming, version 2.0. Accessed October 5, 2020, <http://cvxr.com/cvx>.
- Davenport TH, Harris J, Shapiro J (2010) Competing on talent analytics. *Harvard Bus. Rev.* 88(10):52–58.
- Davis MJ, Lu Y, Sharma M, Squillante MS, Zhang Bo (2018) Stochastic optimization models for workforce planning, operations, and risk management. *Service Sci.* 10(1):40–57.
- De Feyter T, Guerry M-A, Komarudin (2017) Optimizing cost-effectiveness in a stochastic Markov manpower planning system under control by recruitment. *Ann. Oper. Res.* 253(1):117–131.
- Dimitriou VA, Georgiou AC, Tsantas N (2013) The multivariate non-homogeneous Markov manpower system in a departmental mobility framework. *Eur. J. Oper. Res.* 228:112–121.
- Flynn J (1981) A dynamic programming model of fractional flows with application to maintenance and replacement problems. *Math. Oper. Res.* 6(3):405–419.
- Fry MJ, Magazine MJ, Rao US (2006) Firefighting staffing including temporary absences and wastage. *Oper. Res.* 54(2):353–365.
- Gans N, Zhou Y-P (2002) Managing learning and turnover in employee staffing. *Oper. Res.* 50(6):991–1006.
- Georgiou AC, Tsantas N (2002) Modelling recruitment training in mathematical human resource planning. *Appl. Stochastic Models Bus. Indust.* 18(1):53–74.
- Guerry M-A, De Feyter T (2012) Optimal recruitment strategies in a multi-level manpower planning model. *J. Oper. Res. Soc.* 63: 931–940.
- Iverson RD (1999) An event history analysis of employee turnover: The case of hospital employees in Australia. *Human Resource Management Rev.* 9(4):397–418.
- Kalamatianou AG (1987) Attainable and maintainable structures in Markov manpower systems with pressure in the grades. *J. Oper. Res. Soc.* 38(2):183–190.
- Kim K, Mehrotra S (2015) A two-stage stochastic integer programming approach to integrated staffing and scheduling with application to nurse management. *Oper. Res.* 63(6):1431–1451.
- Kuwaiti AA, Raman V, Subbarayalu AV, Palanivel RM, Prabaharan S (2016) Predicting the exit time of employees in an organization using statistical model. *Internat. J. Sci. Tech. Res.* 5(12):213–217.
- Lim AEB, Shanthikumar JG (2007) Relative entropy, exponential utility, and robust dynamic pricing. *Oper. Res.* 55(2):198–214.
- Lusby R, Dohn A, Range TM, Larsen J (2012) A column generation-based heuristic for rostering with work patterns. *J. Oper. Res. Soc.* 63(2):261–277.
- Marler JH, Boudreau JW (2017) An evidence-based review of HR analytics. *Internat. J. Human Resource Management* 28(1):3–26.
- Mehlmann A (1980) An approach to optimal recruitment and transition strategies for manpower systems using dynamic programming. *J. Oper. Res. Soc.* 31(11):1009–1015.
- Miles L, Emre Y, Russell B, Vielma JP (2016) Extended formulations in mixed-integer convex programming. Louveaux Q, Skutella M, eds. *Integer Programming Combinatorial Optim. 18th Internat. Conf. IPCO 2016 Proc.* (Springer International Publishing, Cham, Switzerland), 102–113.
- MOSEK ApS (2019) MOSEK Fusion API for Python 9.0. Accessed October 5, 2020, <https://docs.mosek.com/9.0/pythonfusion/index.html>.
- Nembhard DA, Bentefouet F (2012) Parallel system scheduling with general worker learning and forgetting. *Internat. J. Production Econom.* 139(2):533–542.
- Nilakantan K, Raghavendra BG (2008) Length of service and age characteristics in proportionality Markov manpower systems. *IMA J. Management Math.* 19(3):245–268.
- Park S-H, Lee SM, Yoon SN, Yeon S-J (2008) A dynamic manpower forecasting model for the information security industry. *Indust. Management Data Systems* 108(3):368–384.
- Paul C, Mitlacher L (2008) Expanding risk management systems: Human resources and German banks. *Strategic Change* 17(1-2): 21–33.

- Perelman BS, Dorton SL, Harper S (2019) Identifying consensus in heterogeneous multidisciplinary professional teams. *2019 IEEE Conf. Cognitive Comput. Aspects Situation Management CogSIMA* (IEEE, Piscataway, NJ), 127–133.
- Price WL, Piskor WG (1972) The application of goal programming to manpower planning. *INFOR Inform. Systems Oper. Res.* 10(3): 221–231.
- Rao PP (1990) A dynamic programming approach to determine optimal manpower recruitment policies. *J. Oper. Res. Soc.* 41(10): 983–988.
- Roos NP, Fransoo R, Bogdanovic B, Carriere KC, Frohlich N, Friesen D, Patton D, Wall R (1999) Needs-based planning for generalist physicians. *Med. Care* 37(6):206–228.
- Şenerdem BH (2001) The determination of optimal time-in-grade for promotion at each rank in the Turkish Army. Unpublished master's thesis, Bilkent University, Ankara, Turkey.
- Shafer SM, Nembhard DA, Uzumeri MV (2001) The effects of worker learning, forgetting, and heterogeneity on assembly line productivity. *Management Sci.* 47(12):1639–1653.
- Shannon CE (1948) A mathematical theory of communication. *Bell System Tech. J.* 27:379–423, 623–656.
- Sing C-P, Love PED, Tam CM (2012) Multiplier model for forecasting manpower demand. *J. Construction Engrg. Management* 138(10):1161–1168.
- Song H, Huang H-C (2008) A successive convex approximation method of multistage workforce capacity planning problem with turnover. *Eur. J. Oper. Res.* 188:29–48.
- Ulrich D, Dulebohn J (2015) Are we there yet? What's next for HR? *Human Resource Management Rev.* 25(2):188–204.
- Ulrich D, Brockbank W, Younger J, Ulrich M (2012) *HR from the Outside In: Six Competencies for the Future of Human Resources* (McGraw-Hill, New York).
- Valeva S, Hewitt M, Thomas BW (2017) A matheuristic for workforce planning with employee learning and stochastic demand. *Internat. J. Production Res.* 55(24):7380–7397.
- Waltz F (1967) An engineering approach: Hierarchical optimization criteria. *IEEE Trans. Automatic Control* 12(2):179–180.
- Yan Z, Gao SY, Teo CP (2018) On the design of sparse but efficient structures in operations. *Management Sci.* 64(7):3421–3445.
- Zhu X, Sherali HD (2009) Two-stage workforce planning under demand fluctuations and uncertainty. *J. Oper. Res. Soc.* 60(1): 94–103.

Patrick Jaillet is the Dugald C. Jackson Professor in the Department of Electrical Engineering and Computer Science and a member of the Laboratory for Information and Decision Systems at MIT. He is also Co-Director of the MIT Operations Research Center. His research interests include online optimization and learning, machine learning, and decision making under uncertainty. He is a Fellow of INFORMS.

Gar Goei Loke is an assistant professor in the Department of Technology and Operations Management in the Rotterdam School of Management, Erasmus University. His research interests include optimization under uncertainty and its application in queueing settings and its incorporation in machine learning techniques.

Melvyn Sim is a professor in the Department of Analytics and Operations, NUS Business School, National University of Singapore. He has broad interests in prescriptive analytics, particularly in decision making and optimization under uncertainty with applications in operations, transportation, healthcare, and service systems.