

MANAGING LEARNING AND TURNOVER IN EMPLOYEE STAFFING

NOAH GANS

*OPIM Department, The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104-6366,
gans@wharton.upenn.edu*

YONG-PIN ZHOU

*Department of Management Science, School of Business Administration, University of Washington, Seattle, Washington 98195-3200,
yongpin@u.washington.edu*

(Received July 1999; revisions received November 2000, May 2001; accepted October 2001)

We study the employee staffing problem in a service organization that uses employee service capacity to meet random, nonstationary service requirements. The employees experience learning and turnover on the job, and we develop a Markov Decision Process (MDP) model which explicitly represents the stochastic nature of these effects. Theoretical results show that the optimal hiring policy is of a state-dependent “hire-up-to” type, similar to an inventory “order-up-to” policy. For two important special cases, a myopic policy is optimal. We also test a linear programming (LP) based heuristic, which uses average learning and turnover behavior, in stationary environments. In most cases, the LP-based policy performs quite well, within 1% of optimality. When flexible capacity—in the form of overtime or outsourcing—is expensive or not available, however, explicit modeling of stochastic learning and turnover effects may improve performance significantly.

1. INTRODUCTION

We consider the employee staffing problem at a service organization. Suppose there are service requirements the organization must meet, and a forecast of these requirements is available for several periods of time in the future. The forecast may be based on historical data and a projection of the organization’s business path. The organization’s objective is to use the capacity provided by its employees to meet the service requirements in a least-cost fashion. Two essential factors in the problem differentiate it from more traditional capacity-planning problems.

First, the people providing these services are not homogeneous. Different employees will have different service capacities (or skill levels), and these service capacities change over time. When people learn on the job, their service capacities increase; when they turn over, their service capacities are lost. Learning can take place because of either initial training or on-the-job training. Turnover can happen because of a mismatch with the job or a better career opportunity elsewhere.

Second, learning and turnover are typically random elements. Therefore, the numbers of employees in the organization—and the need to hire additional workers—can be difficult to predict. Furthermore, the initial training period may be long. These factors make hiring to meet future service requirements a difficult problem to solve effectively.

1.1. Our Approach


We take a hierarchical approach to this problem that mirrors the hierarchical planning strategies used in a manufacturing

environment. A long-term, high-level staffing problem corresponds to long-term capacity planning. A medium-term, mid-level workforce scheduling problem corresponds to the mid-level planning for which many companies use Material Requirements Planning (MRP). Finally, there exists a moment-by-moment, low-level work-assignment problem that is the analogue of real-time shop-floor control in a manufacturing firm.

Example: A Telephone Call Center. Figure 1 displays three months of daily call volumes for a small retail-banking call center. (These data are more systematically analyzed in Mandelbaum et al. 2000.) The service standards in these types of call centers are something like: “on average a phone call is put on hold for 20 seconds or less” or “80% of the time a phone call is put on hold for 20 seconds or less.” Typically, the problem is disaggregated into short-term, medium-term, and long-term planning components. In the short run, the call center must solve a real-time control problem, assigning incoming calls to available customer service representatives (CSRs). The medium-term problem, workforce scheduling, is typically solved on a weekly basis. In any given week, the numbers of employees of different types are fixed, and a schedule is developed that minimizes overtime and outsourcing costs, subject to the call center’s service level requirements. (“Outsourcing” refers to a common practice in which one call center diverts incoming calls to another on a contract basis. In addition to paying a long-term retainer, the diverting call center typically pays a per-call fee for this service.) The long-term component seeks to hire the right numbers of people and

Subject classifications: Dynamic programming/optimal control, applications: hierarchical model for manpower planning. Organizational studies, manpower planning: MDP staffing model with learning and turnover.

Area of review: SERVICES AND MILITARY.

Figure 1. Call center example.

train them the right way so that when the sequence of weekly scheduling problems is solved, they produce a low-cost—if not least-cost—solution to the global problem.

In this paper, we focus on a common special case of the long-term hiring problem. In this case, every employee progresses through the same sequence of learning or training states, and the organization’s primary decision is the number of new employees to hire in each period.

We formulate the problem as a discrete-time, continuous-state-space Markov Decision Process (MDP), where the state variable vector represents the numbers of people at different levels on the learning curve. This approach allows us to model naturally the randomness in the system and to prove the existence of desirable structural properties of the optimal hiring policy.

To suppress the other two lower-level problems, we will assume that there exists an “operating cost” function, $O(\cdot)$, which, given the numbers of employees at different skill levels and a forecast for service requirements, will tell us how much the total operating costs will be. This operating cost function should be based on an efficient, if not optimal, solution to the scheduling problem.

For example, in call centers managers usually use commercial software to do workforce scheduling each week. This software solves a large-scale mixed integer program, and it returns a schedule which shows the amount of overtime and outsourcing used for the week. Using this software, one may generate $O(\cdot)$ as a response function.

1.2. Overview of Results

We show that under a discounted-cost criterion, convexity of the operating cost function (along with certain other costs, such as hiring and wages) is propagated through MDP value iteration. Therefore, when $O(\cdot)$ is convex,

the optimal hiring policy can be characterized as a state-dependent “hire-up-to” policy. This holds in both finite and infinite horizon cases, with general, nonstationary service requirements.

We also develop results that offer a more detailed characterization of the optimal policy in two special cases. In particular, we show that when service requirements are stationary or increasing, and when there is (1) no learning, no training lead time, and stochastic turnover; or (2) no learning, positive training lead time, and deterministic turnover, then a “myopic” policy is optimal. A myopic policy optimizes a one-period static problem for each period, rather than the dynamic multiperiod problem of the MDP. In both special cases, when demands are k -periodic, then a k -period analogue of a myopic policy is optimal. That is, it is sufficient to solve a k -period MDP with appropriate end-of-horizon cost function.

Using numerical examples, we also compare a linear programming (LP) based heuristic—which ignores the stochastic nature of the state evolution—with the optimal hiring policy. The examples have stationary, deterministic service requirements, and we find that in most cases the cost incurred using the LP heuristic falls within 1% of optimal. This suggests that, in many cases, one may be able to develop effective staffing plans without explicitly modeling random variation in learning and turnover rates.

In some cases the performance of the LP heuristic lags that of the optimal policy, however. These are instances in which: (1) there is a (hiring or training) lead time between the time the employer seeks to hire an employee and the time the employee becomes productive, and (2) flexible capacity, in the form of overtime and outsourcing, is limited. In these cases, the LP heuristic does not recognize the need for a “buffer” of excess staff that the optimal policy provides, and the explicit modeling of random variation in learning and turnover behavior may be warranted.

As is clear from our MDP formulation and the structure of optimal hiring policies, our results connect closely with those of inventory and production control. In our system, the “inventory” is the population of employees, and “demand” corresponds to turnover, however. In turn, the dynamics in our model differ a bit from more traditional inventory models in that, unlike the demand for inventory, the demand for service in any period does not directly affect how the system state evolves.

This paper is organized as follows. Section 2 reviews the relevant literature. In §3 we present our model and in §4 we develop analytical results concerning the structure of optimal hiring policies. Numerical examples are presented in §5. Finally, in §6 we discuss our results and offer directions for future research.

2. LITERATURE REVIEW

Holt et al.’s (1960) seminal manpower planning model and its linear hiring rules have inspired a long stream of research papers. For example, see Orrbeck et al. (1968), Ebert (1976), Gaimon and Thompson (1984), and Khoshnevis and Wolfe (1986), all of which use a mathematical programming approach.

Akşin (1999) also uses a mathematical programming approach. The paper considers infinite-horizon problems and develops characterizations of hiring policies for cases of stationary, increasing, decreasing, and periodic demands. These are analogues of the conditions under which we investigate myopic policies.

There also exist papers devoted to manpower planning that use alternative approaches. Gerchak et al. (1990) develop a partial differential equation model to solve the recruitment rate required to maintain a fixed capacity in an organization. Grinold and Stanford (1974) develop a dynamic programming model based on linear costs, linear capacity (or budget) constraints, and deterministic fractional flows of employees. Moreover, hiring is allowed for all employee types. As a result, it is optimal to adopt a linear hiring policy: In any period, hire only the employee type that has the lowest cost/capacity ratio.

In all of the above papers, the flow of employees among different types, including learning and turnover, is modeled as deterministic. In the model we propose, however, we allow learning and turnover to be stochastic and use a distinct MDP approach. This approach allows us to more carefully model the dynamics of learning and turnover among employees.

There is a separate stream of research that develops Markov chain models for human resource planning problems. Bartholomew et al. (1991) provide an excellent summary of research in this area. In this paper, we adopt a model closely related to the “mixed-exponential” model found in the human resource literature. Little work in this area has been devoted to the control aspect of the recruitment process, however. Control of the hiring process is the central question for us.

Two recent exceptions which consider aspects of control are papers by Bordoloi and Matsuo (2001) and Pinker and Shumsky (2000). The first paper derives steady-state performance measures for (heuristic) linear control rules that are applied in a manufacturing environment that is similar to ours. It does not, however, consider the nature of optimal control policies. The second paper considers the improvement in quality which comes with worker experience. This focus on quality is complementary to the capacity analysis of this paper.

As is noted in the Introduction, there exists a close connection between our results and those of classic inventory theory. In addition to papers by Karlin (1960a, 1960b) and Zipkin (1989), our work is most closely related to the work on inventory systems with spoilage by Iglehart and Jaquette (1969).

There is also a close connection between our results and those for capacitated production systems. In these systems, demand for a product is stochastic and a system controller varies production from period to period in order to meet demand through inventory, as well as through current production. For example, see Aviv and Federgruen (1997) and Kapuściński and Tayur (1998) and the references therein.

3. MODEL

We assume that in our discrete-time, continuous-state-space MDP the number of planning periods under consideration is T ($T = \infty$ in the infinite planning horizon case), and we index time periods by $t = 0, 1, \dots, T$. The length of time included in one period may be a week, a month, a quarter, or a year, depending on the application.

3.1. State Space and Control

In any period, an employee may have attained one of m discrete skill levels, $i = 1, 2, \dots, m$. These skill levels may correspond to service speeds or acquired skills. In the former case, all employees do the same job, but their service speeds increase with i . In the latter case, there are many job types. Employees start at Level 1, with the most basic job skills, and progressively acquire more skills to become capable of handling more types of jobs. A combination of both cases is also possible: Employees not only acquire new skills as they progress, they also become faster at these skills.

Turnover may occur at any level i , $1 \leq i \leq m$, but hiring is made only at the entry level, Level 1. Thus, when new employees are hired they start at Level 1 and progressively go through a fixed sequence of skill levels until they turn over. Furthermore, the employer exercises control over the workforce solely through entry-level staffing decisions.

This model is appropriate for the call center environments with which we are familiar. In them, all CSRs go through the same initial training when hired. After training, CSRs then go through a common learning curve as they

gain experience and speed at handling calls. More generally, the model should fit environments in which relatively unskilled people are hired and trained to perform a standardized set of tasks.

Formally, we use a vector $(n_{1,t}, \dots, n_{i,t}, \dots, n_{m,t})$ to denote the numbers of employees at levels $i = 1, \dots, m$ before hiring at the beginning of period t . Because only Type-1 employees will be hired, the vector becomes $(y_t, n_{2,t}, \dots, n_{m,t})$ after x_t people are hired, where $y_t = n_{1,t} + x_t$. We treat $n_{i,t}$ as real variables. This relaxation of integrality makes our proofs, which rely on convexity properties, less burdensome. For large operations with many employees, the continuous approximation should be reasonable.

3.2. Costs and the Objective Function

In any period t , three types of costs are incurred. First, a fixed cost of h is incurred to hire a new employee. This cost typically includes advertising for, interviewing, and testing of job applicants, when appropriate. It may also include one-time training costs that are independent of wages.

Second, we let W_i be the wage cost of a type- i employee. This is the “fully loaded” compensation cost for one period and typically includes wages, benefits, and other direct personnel costs. Note that these costs may vary with the skill level attained by the employee.

All other costs that arise in period t are captured in the *operating cost function*, $O_t(y_t, n_{2,t}, \dots, n_{m,t}; \vec{D}_t)$. These are “variable” costs that change with the level of demand for service, \vec{D}_t . They typically include employee overtime and the variable cost of outsourcing.

The elements of the vector \vec{D}_t define the operation’s service requirements for different sub-intervals within the larger, discrete time period, t . For example, in a telephone call center, hiring may take place on a weekly or monthly basis, while call volume forecasts and service standards are specified in 15- or 30-minute intervals. If the planning period is one week, there are \mathcal{J} job types, and volume forecasts are made every 15 minutes, then the vector \vec{D}_t will have a dimension of $\mathcal{J} \times 672$.

Furthermore, the elements of \vec{D}_t need not be scalars at all. For example, by letting the elements of \vec{D}_t be probability distributions, we may explicitly model the uncertainty inherent in demand forecasts.

We intentionally let the definition of $O_t(\cdot)$ remain a bit vague at this point. In the context of our hierarchical approach, it reflects the (often difficult) work assignment and workforce scheduling problems that must be solved in period t . It must also incorporate service-level constraints that ensure that “adequate” capacity is obtained to “reasonably” serve demand. Because the nature of the work assignment and workforce scheduling problems—and the resulting form of $O_t(\cdot)$ —may vary substantially from one setting to another, we leave the function undefined. Our only technical requirement is that $O_t(y_t, n_{2,t}, \dots, n_{m,t}; \vec{D}_t)$ be jointly convex in $(y_t, n_{2,t}, \dots, n_{m,t})$. In §3.4 we offer a

simple, general example of $O_t(\cdot)$ in which this convexity property holds.

Finally, we use α to denote the one-period discount factor in the MDP and minimize the expected total discounted cost. Therefore, the objective is to minimize the function

$$\min_{x_0, y_0, \dots, x_T, y_T} \mathbf{E} \left\{ \sum_{t=0}^T \left[\alpha^t \left(hx_t + W_1 y_t + \sum_{i=2}^m W_i n_{i,t} + O_t(y_t, n_{2,t}, \dots, n_{m,t}) \right) \right] \right\}, \quad (1)$$

subject to the system dynamics (2), (3), (4), and (5), which are defined below.

Note that for simplicity we only consider costs incurred during the planning horizon in (1), therefore the “end-of-horizon” cost function is zero. Other types of end-of-horizon cost functions can also be incorporated.

3.3. Learning, Turnover, and System Dynamics

Learning and turnover are accounted for only at the end of each time period. We assume that, of the $n_{i,t}$ level- i employees in period t , $\tilde{Q}_{i,t}(n_{i,t})$ will quit and $\tilde{L}_{i,t}(n_{i,t})$ will learn and move to skill level $i+1$ in period $t+1$. Here “ $\tilde{\cdot}$ ” is used to denote random numbers, indicating the fact that learning and turnover occur randomly. Note that the i and t subscripts mean that learning and turnover can have different patterns at different levels in different periods.

Our formulation assumes that learning and turnover are Markovian. That is, we assume an employee’s probability of learning or turning over in any period depends only on that person’s state i and not on how many periods he or she has been in that state.

The Markovian assumption clearly helps to reduce the state space of the problem, and it should be reasonable when the rate of movement through skill levels is fairly homogeneous across individual employees. In some environments, however, the progress through the skill levels is not homogeneous. In these cases, knowing how many periods an employee has spent at a particular level may significantly affect the conditional probability that a person will proceed to the next level in the following period. Thus, the Markovian assumption breaks down.

Because individual employees typically learn and turn over independently of each other, a natural distribution to use for modeling learning and turnover in each period is the multinomial distribution. It implies that a person’s length of stay at any level is geometrically distributed with different parameters for different levels. Note that this “mixed geometric” distribution is a discrete version of the “mixed exponential” tenure length distribution which has been widely used in the human resources literature. For more information on mixed exponential and other Markov models of employee tenure, see Bartholomew et al. (1991).

When employees turn over but do not learn—so that the number of employees that turn over is binomially

distributed—we can use well-known results to demonstrate that the value function is convex in the number of employees (see Karlin 1968). No such results exist for more complex multinomial distributions, however, and we approximate them with analytically more tractable stochastic proportions. More specifically, we assume that $\tilde{l}_{i,t}$ and $\tilde{q}_{i,t}$ are independently distributed random variables with support on $[0, 1]$ that represent the “stochastic proportion” of people who learn and turn over ($\tilde{l}_{m,t} = 0$). Therefore, in period t ,

$$\tilde{Q}_{i,t}(\cdot) = \begin{cases} \tilde{q}_{i,t} \cdot y_t & \text{if } i = 1 \\ \tilde{q}_{i,t} \cdot n_{i,t} & \text{if } i > 1 \end{cases} \quad \text{and}$$

$$\tilde{L}_{i,t}(\cdot) = \begin{cases} \tilde{l}_{1,t} \cdot (1 - \tilde{q}_{1,t}) \cdot y_t & \text{if } i = 1 \\ \tilde{l}_{i,t} \cdot (1 - \tilde{q}_{i,t}) \cdot n_{i,t} & \text{if } i > 1 \end{cases}.$$

Given the numbers of employees in the system, $(n_{1,t}, \dots, n_{m,t})$, the number of Type-1 employees hired, x_t , and the fractions of employees turning over and learning at time t , $\tilde{q}_{i,t}$ and $\tilde{l}_{i,t}$, the numbers of employees at $t+1$ are straightforward to calculate. The following equations represent the system evolution:

$$y_t = n_{1,t} + x_t; \quad x_t \geq 0 \quad (2)$$

$$n_{1,t+1} = (1 - \tilde{l}_{1,t})(1 - \tilde{q}_{1,t})y_t; \quad (3)$$

$$n_{2,t+1} = (1 - \tilde{l}_{2,t})(1 - \tilde{q}_{2,t})n_{2,t} + \tilde{l}_{1,t}(1 - \tilde{q}_{1,t})y_t; \quad (4)$$

$$n_{i,t+1} = (1 - \tilde{l}_{i,t})(1 - \tilde{q}_{i,t})n_{i,t} + \tilde{l}_{i-1,t}(1 - \tilde{q}_{i-1,t})n_{i-1,t}, \quad 2 < i \leq m. \quad (5)$$

Note that the means of the stochastic proportions can be set to equal the multinomial probabilities that an individual turns over or learns. Stochastic proportions do not naturally capture second-order effects, however. That is, when the distribution is multinomial, the relative dispersion of the fraction of employees quitting or advancing changes with the number of employees, but it does not change with stochastic proportions.

To the extent that in steady state the number of employees of one type remains large and relatively stable, these second-order differences should not induce undue bias. For example, we have tested problems with multinomially distributed transition probabilities that are analogues of the numerical examples presented in §5. In all cases, the cost difference observed between the multinomial and stochastic proportion models was less than 0.5%.

REMARK 1. Typically, the rates at which learning and turnover occur have particular functional forms. The log-linear “learning curve” has been widely used in manufacturing (see Yelle 1979), and there is empirical evidence that it also exists in service operations. An example with directory assistance operators can be found in Gustafson (1982). Turnover typically decreases with job tenure: Employees whose “job fit” is poor tend to leave after a short time;

employees with longer tenure are self-selected to have better job fit (for example, see Jovanovic 1979). In our experience working with telephone call centers, we have found that the turnover rate can decline by more than 50% after an initial training and adjustment period. Structurally, however, our model and its results do not depend on the rates at which learning and turnover take place. The “typical” learning curve and turnover behaviors described above affect only the problem data used in the model.

3.4. Operating Cost Function

Now we provide a more detailed example of the operating cost function $O_t(y_t, n_{2,t}, \dots, n_{m,t}; \vec{D}_t)$. We describe how it can be modeled to address uncertainty in the service requirement forecast, as well as to capture the essence of the assignment and workforce scheduling problems.

Suppose there are \mathcal{J} job types and that within any time period t there are in total \mathcal{S} subintervals, indexed by $s = 1, 2, \dots, \mathcal{S}$. The forecast $\vec{D}_t = \{D_t^{sj}\}_{s=1, \dots, \mathcal{S}}^{j=1, \dots, \mathcal{J}}$ specifies the service requirement for each job type in every subinterval. This may be a point forecast describing expected demand, or a distributional forecast that accounts for uncertainty.

For this time period there also exist \mathcal{W} feasible work schedules for employees, indexed by $w = 1, 2, \dots, \mathcal{W}$. These work schedules may or may not include overtime, but all of them satisfy workplace rules and regulations regarding breaks, overtime, and so on. Let x_{iw} denote the number of type- i employees on work schedule w , and let $C_{iw}(x_{iw})$ denote the cost of placing the x_{iw} type- i employees on schedule w .

When the service requirements cannot be met with available regular time and overtime, outsourcing will be used to satisfy the residual service requirements. We let z_s denote the amount of work to be outsourced in subinterval s and $OS_s(z_s)$ the outsourcing cost during s .

The solution to the following mathematical program defines the operating cost in period t :

$$\begin{aligned} & O_t(y_t, n_{2,t}, \dots, n_{m,t}; \vec{D}_t) \\ &= \min_{x_{iw}, z_{sj}} \sum_{i=1}^m \sum_{w=1}^{\mathcal{W}} C_{iw}(x_{iw}) + \sum_{s=1}^{\mathcal{S}} \sum_{j=1}^{\mathcal{J}} OS_{sj}(z_{sj}) \\ & \text{s.t. } f(x_{11}, \dots, x_{m\mathcal{W}}, z_{11}, \dots, z_{\mathcal{S}\mathcal{J}}; \vec{D}_t) \geq 0 \quad (6) \\ & \sum_{w=1}^{\mathcal{W}} x_{1w} \leq y_t \\ & \sum_{w=1}^{\mathcal{W}} x_{iw} \leq n_{i,t} \quad \forall 2 \leq i \leq m \\ & x_{iw}, z_{sj} \geq 0 \quad \forall i, w, s, j. \end{aligned}$$

Inequality (7) is a service-level constraint that reflects the result of the low-level work assignment aspect of the problem.

A simple example of (7) is that of set-covering formulation, in which the average processing rate assigned to each interval s must be enough to cover the service requirement:

$$\sum_{w: I(w,s)=1} \mu_i x_{iw} + z_s \geq D_i^s \quad s = 1, \dots, \mathcal{S}. \quad (7)$$

Here, service requirements D_i^s are for average processing rates, and the indicator $I(w, s) = 1$ if schedule w assigns employees to work in interval s , and it equals 0 otherwise. Berman et al. (1997) use a similar LP formulation to solve the workforce and workflow scheduling problem at high-volume service operations such as USPS Mail Processing Centers.

In general, however, the evaluation of Equation (7) is not transparent. For example, it may say “the choice of $(x_{11}, \dots, x_{m^{\mathcal{W}}}, z_{11}, \dots, z_{\mathcal{S}\mathcal{J}})$ should be such that the probability of meeting a certain service standard is more than a certain percentage, $a\%$.”

Nevertheless, by imposing a mild restriction on (7) we can guarantee the convexity of the function $O_t(\cdot)$. For a proof of the following proposition, please see Appendix A.

PROPOSITION 1. *If (i) the set $\{(x_{11}, \dots, x_{m^{\mathcal{W}}}, z_{11}, \dots, z_{\mathcal{S}\mathcal{J}}): f(x_{11}, \dots, x_{m^{\mathcal{W}}}, z_{11}, \dots, z_{\mathcal{S}\mathcal{J}}; \vec{D}_t) \geq 0\}$ is convex; and (ii) the cost functions $C_{iw}(\cdot)$ and $OS_{sj}(\cdot)$ are convex, then $O_t(\cdot)$ is jointly convex in $(y_t, n_{2,t}, \dots, n_{m,t})$.*

Conditions (i) and (ii) appear to be reasonable. For example, for condition (i), if two solutions to the math program provide acceptable service more than $a\%$ of the time, it is reasonable to assume that some convex combination of them will also do so. This assumption is true, for example, when service requirements, \vec{D}_t , are specified only as requirements for total processing capacity, as in (7). Because the marginal cost of each extra unit of outsourced work is typically increasing, it is also reasonable to assume $OS_{sj}(\cdot)$ to be convex. $C_{iw}(\cdot)$ is usually linear.

3.5. Training and Hiring Lead Times

In organizations in which employees need to acquire a certain level of service proficiency before they are put “online,” training is an important component of the hiring process. For example, in some call centers initial training for a CSR may take more than two months to complete. In this case, hiring lead time (of which training is a major component) becomes very important in capacity planning for future periods.

Because our model already includes multiple employee types, it is straightforward to incorporate this lead time. If the hiring/training lead time is λ periods, then we add λ employee types (thus $m > \lambda$) and use employee types 1 to λ to denote the trainees. For organizations in which training mainly takes place on the job, $\mu_i > 0, \forall i = 1, \dots, \lambda$ denotes the service output these trainees may generate. For other organizations in which employees must receive a minimum amount of initial training before working, $\mu_i = 0, \forall i = 1, \dots, \lambda$. Note that W_i for $1 \leq i \leq \lambda$ are the trainee wage and benefit rates. If the hiring process itself takes some time, then the first few W_i s may be 0.

3.6. Relationship to Inventory Models

In this system, the “inventory” is the population of employees, and “demand” corresponds to turnover. Similarly, multiple levels of skills correspond to multiple classes of inventory, and learning is the analogue of transfer of inventory from one class to another. There are two natural ways to view this system as an inventory system.

First, the model may be viewed as a single-location model with multiple types of inventories. Here, turnover corresponds to “spoilage” or some other special type of inventory-level-dependent demand, and learning corresponds to the change of type among different types of inventories.

Second, the model may be viewed as a multiechelon inventory system. Here, different levels on the learning curve correspond to different system echelons. In turn, learning is the analogue of transfer between echelons, and turnover represents the echelons’ inventory-level-dependent demands.

Note, however, that the dynamics of the model differ from those in traditional inventory models. In inventory models, demand depletes inventory and changes the state of the system. In this model, service requirements drive the operating cost function, but they do not change the state of the system.

4. STRUCTURAL RESULTS CONCERNING OPTIMAL POLICIES

In this section, we use MDP value iteration to characterize structural properties of optimal hiring policies. Our first result, which holds in great generality, shows that optimal policies are of the “hire-up-to” type. For special cases in which the state space can be collapsed into one dimension, we show that computationally tractable myopic policies are optimal. In these cases, we analyze systems in which service requirements are stationary or stochastically increasing, as well as those in which service requirements are periodic.

4.1. Optimality of Hire-Up-To Policies

We define a hire-up-to policy as follows:

DEFINITION 1. A policy is of the state-dependent hire-up-to type if, for any $(n_{1,t}, n_{2,t}, \dots, n_{m,t})$, there exists $y_t^*(n_{2,t}, \dots, n_{m,t})$ such that the optimal hiring number $x_t^*(n_{1,t}, \dots, n_{m,t})$ is

$$x_t^*(n_{1,t}, \dots, n_{m,t}) = \begin{cases} y_t^*(n_{2,t}, \dots, n_{m,t}) - n_{1,t} & \text{if } n_{1,t} < y_t^*(n_{2,t}, \dots, n_{m,t}) \\ 0 & \text{otherwise.} \end{cases}$$

Given our problem structure, we can prove the optimality of hire-up-to policies for finite-horizon problems with no additional assumptions. For the case of infinite-horizon problems, however, we need to make two technical assumptions about the state space and cost. In particular, we define

$H_t(y_t, n_{2,t}, \dots, n_{m,t}) \stackrel{\text{def}}{=} (h + W_1)y_t + O_t(y_t, n_{2,t}, \dots, n_{m,t})$, and we assume that for any planning horizon T there exist large constants, M and K , such that:

ASSUMPTION 1. $y_t^* + \sum_{i=2}^m n_{i,t} \leq M \quad \forall i, t$.

ASSUMPTION 2. $H_t(y_t, n_{2,t}, \dots, n_{m,t}) + \sum_{i=2}^m W_i n_{i,t} \leq K$, $\forall t$ and $\forall (y_t, n_{2,t}, \dots, n_{m,t})$ where $y_t + \sum_{i=2}^m n_{i,t} \leq M$.

Assumption 1 states that at any time the total number of employees after hiring does not exceed M . Because learning and turnover will not increase the total number of employees in the organization, this ensures that if we start with fewer than M people in total, then we will never exceed that number. As a result, the state space can be reduced to a bounded and closed, and therefore compact, subset of \mathbf{R}^m .

Assumption 2 states that whenever there is a finite number of people in the organization, the total one-period cost is bounded. This is straightforward for hiring and wage costs, and it should also hold for the operating costs. While we have not explicitly defined $O_t(\cdot)$, we know that as long as there exist finite bounds on all service requirements and cost parameters, the total operating cost in any time period is bounded.

Now we are ready to state our main result:

THEOREM 1. *Suppose either (a) $T < \infty$ or (b) $T = \infty$ and Assumptions 1 and 2 hold. Then a policy of the hire-up-to type is optimal.*

SKETCH OF PROOF. The proof of part (i) follows classic proofs of convexity results in the inventory literature (for example, see Heyman and Sobel 1982). It differs from that for more traditional inventory models, however, because of the dynamics in which “inventory” changes type. In particular, we use the following lemmas to show that convexity of the MDP value function is preserved in the presence of learning and turnover and that convexity is propagated through the minimization that is central to MDP recursion. They follow from Theorems 5.7 and 5.3 in Rockafellar (1970).

LEMMA 1. *If $g(y_1, \dots, y_n)$ is jointly convex in (y_1, \dots, y_n) and $f_i(x_1, \dots, x_m), i = 1, \dots, n$, are linear functions in (x_1, \dots, x_m) , then $h(x_1, x_2, \dots, x_m) = g(f_1(x_1, \dots, x_m), \dots, f_n(x_1, \dots, x_m))$ is also jointly convex in (x_1, \dots, x_m) .*

LEMMA 2. *Let*

$$f(n_1, n_2, \dots, n_m) = \inf_{y \geq n_1} \{h(y, n_2, \dots, n_m)\}.$$

If $h(y, n_2, \dots, n_m)$ is jointly convex in (y, n_2, \dots, n_m) , then $f(n_1, n_2, \dots, n_m)$ is jointly convex in (n_1, n_2, \dots, n_m) .

For a more detailed statement and proof of Theorem 1, please see Appendix B.

The fact that hire-up-to policies are optimal means that the current number of entry-level employees need not be taken into account when deciding on the appropriate hire-up-to number. Thus, the dimensionality of the search space (of policies) may be reduced by one. This policy-space reduction is clearly advantageous when solving the MDP.

Even in cases in which the MDP is not solved explicitly, the results of Theorem 1 reduce the search for optimal policies to those within the class of the hire-up-to type. In this way, they should help to speed the performance of other policy-based search methods such as simulation-based optimization techniques, as well as the LP heuristic developed in §5.

At the same time, the optimal hiring policy may still be difficult to implement. The high dimensionality of the state space makes the computational task required formidable, particularly in organizations with many employees. In certain cases, however, we may be able to exploit additional structure inherent in the problem to develop computationally efficient heuristics that perform well.

In the following sections, we examine two cases in which we can collapse the state space into one dimension so that myopic policies are optimal. These myopic hiring policies are computationally much more efficient than the solution of the general MDP.

4.2. Two Cases with One-Dimensional State Spaces

In some organizations very little initial training is used, and learning on the job is so fast or so little that we can model all the employees as one type. The prototypical example of this type of operation is a “fast-food” restaurant. Even though learning is not a factor in these organizations, turnover remains an important problem. With only one employee type and no hiring lead time, $m = 1$ and the state space becomes one-dimensional.

Even when there is a positive hiring lead time, if turnover during the hiring lead time occurs at a deterministic rate, then we can collapse the state space to one dimension. With a hiring lead time of λ and no learning, we make the following formal assumption:

ASSUMPTION 3. *A system with a hiring lead time of λ , no learning, and deterministic lead-time turnover has $m = \lambda + 1$ skill levels, as well as the following additional parameters: (i) only type $(\lambda + 1)$ has productive capacity: $\mu_{\lambda+1} > 0$, and $\mu_i = 0$, for $1 \leq i \leq m - 1$; (ii) turnover rates in the hiring pipeline are deterministic: $\tilde{q}_{i,t} = q_{i,t}$ with probability one for constants $q_{i,t}$, $i = 1, \dots, \lambda$; (iii) all employees in the hiring pipeline advance to the next level with probability one: $\tilde{l}_{i,t} = 1$ with probability one; and (iv) first period wage costs W_1 include hiring costs, so that $h = 0$.*

Deterministic turnover rates during the hiring lead time are reasonable when there is little turnover, as well as in cases of planned attrition. Note also that Type-1 employees advance to Type-2 with probability 1, so no employee will incur hiring costs more than once.

In this case we can define a “staffing position” that follows in spirit the “inventory position” concept in the inventory literature (e.g., see Arrow et al. 1958). Here, however, turnover of employees in the hiring/training pipeline makes the weights of employees “on order” less than one. For details, see Appendix C.

4.3. Stationary and Stochastically Increasing Service Requirements

When the sequence of service requirements is stationary or stochastically increasing, we can show that myopic policies are optimal. Formally, we have

DEFINITION 2. A policy is myopic if, in each period t , an action $x_t^G(n_{1,t}, \dots, n_{m,t})$ is taken where $x_t^G(n_{1,t}, \dots, n_{m,t})$ is the minimizer of some one-period cost function $G_t(n_{1,t} + x_t, \dots, n_{m,t}; \vec{D}_t)$.

Thus, myopic policies optimize one-period static problems for each period, rather than the MDP's dynamic, multiperiod problem. When $m = 1$ we can drop the subscript i and keep the meaning of variables and parameters such as n, x, y , and W . Then we let $\tilde{r}_t = 1 - \tilde{q}_t$ and define

$$G_t(y_t; \vec{D}_t) \stackrel{\text{def}}{=} \mathbf{E}_{\{\tilde{r}_t\}} \{((1 - \alpha\tilde{r}_t)h + W)y_t + O_t(y_t; \vec{D}_t)\} \\ = ((1 - \alpha\mathbf{E}\{\tilde{r}_t\})h + W)y_t + O_t(y_t; \vec{D}_t), \quad (8)$$

to be the one-period cost function $G(\cdot)$ that is the basis of a myopic policy.

With a one-dimensional state space and a stationary demand distribution $\{\vec{D}_t\}$, it is not difficult to show that repeatedly optimizing $G(\cdot)$ minimizes (1) as well. Given the following additional assumption, we can also show that myopic policies are optimal when demands are stochastically increasing.

Recall that a random variable X is said to be stochastically less than or equal to another random variable Y , $X \leq Y$, if their corresponding cumulative distribution functions F_X and F_Y satisfy $F_X(x) \geq F_Y(x)$ for all x . Then we have

ASSUMPTION 4. For two forecast distributions, if $\vec{D}_1 \leq \vec{D}_2$, then $dO_t(y_t; \vec{D}_1)/dy_t \geq dO_t(y_t; \vec{D}_2)/dy_t$ for all t and y_t .

Assumption 4 states that the marginal value of an extra employee in reducing the operating cost is higher when the service requirement is stochastically higher.

Using an argument that parallels that in Veinott (1965), we can prove the optimality of using (8) to determine hire-up-to numbers. A detailed statement and proof of the following theorem may be found in Appendix C:

THEOREM 2. Assume that all system parameters except service requirements are stationary. Suppose further that (i) either (a) $m = 1$ or (b) $m > 1$ and Assumption 3 holds; (ii) service requirements are either (a) stationary or (b) stochastically increasing and that Assumption 4 holds; and (iii) either (a) $T = \infty$ or (b) $T < \infty$ and the end-of-horizon cost equals $-\alpha hn_{T+1}$.

Then a myopic policy is optimal. Furthermore, let $y_t^G = \arg \min_{y_t} G_t(y_t)$. (Choose the smallest y_t^G if multiple minimizers exist.) Then when service requirements are stationary $x_{t+1}^* = \tilde{q}_t y_t^G$, and when service requirements are stochastically increasing $x_{t+1}^* = (y_{t+1}^G - y_t^G) + \tilde{q}_t y_t^G$.

The implication here is that if service requirements are stationary, we are simply hiring to replace those who have just left. If the service requirements are increasing, we hire to replace the turnover, as well as to expand in order to meet increasing demand. Therefore, not only is the optimal policy myopic and easy to compute, but the optimal hiring action in steady state is also easy to understand and to exercise.

4.4. Periodic Service Requirements

In many service organizations, customer demand for service is highly seasonal. For example, retail stores and catalog vendors experience holiday-season spikes in demand. Similarly, for call centers in retail financial services, tax season means high call volumes. In this section, we will derive structural properties of the optimal hiring policy when the service requirements are periodic.

We assume that service requirements are periodic with a fixed cycle length. Similarly, we assume that other model parameters, such as costs and turnover rates, are also either stationary or periodic, though not necessarily of the same cycle length. This allows us to treat the whole process as periodic, with a cycle length that is the least common multiplier of the lengths of all the cycles embedded in the system. We will denote this system cycle length by k .

We only consider the infinite planning horizon case. Note that the optimality of hire-up-to policy is shown in Theorem 1. As the optimal decision depends only on the current state and cost data and service requirements that are k -periodic, the optimal hire-up-to levels are also k -periodic. We denote them by $y_0^*, y_1^*, \dots, y_{k-1}^*$.

The following theorem parallels that of Karlin (1960b). Its proof can be found in Appendix D:

THEOREM 3. Suppose Assumptions 1 and 2 hold and that either (a) $m = 1$ or (b) $m > 1$ and Assumption 3 holds. If the problem data are k -periodic, then (i) there always exist periods $\underline{t} \pmod k$ in which a myopic policy is optimal; (ii) the optimal policy is "almost" myopic: The optimal hire-up-to levels can be found through solution of a k -period MDP with \underline{t} as the last period and $-\alpha hn_{\underline{t}+1}$ as the end-of-horizon cost function.

REMARK 2. Two notes: First, Theorem 3 only establishes the existence of \underline{t} , and a procedure still has to be developed to find \underline{t} . In the inventory literature, both Karlin (1960b) and Zipkin (1989) give k -stage algorithms to find \underline{t} and calculate the optimal order-up-to levels. Second, further results from inventory theory can be modified to demonstrate the existence of a more general, one-sided "smoothing" effect. The smoothing effect shows the following: When the service requirement decreases, so does the hire-up-to number; but when the service requirement increases, the hire-up-to number may go down instead. Proofs of smoothing effects for inventory systems can be found in Karlin (1960a) and Zipkin (1989). Analogous proofs, developed for the particular dynamics of this system, can be found in Gans and Zhou (2000).

5. NUMERICAL ANALYSIS

In this section, we compare the performance of the optimal policy with that of an LP-based heuristic that models learning and turnover behavior using deterministic rates. Our test set consists of 45 instances of an example problem with two types of employees and stationary, deterministic service requirements.

We find that in the majority of these examples, the LP-based heuristic performs very well, within 1% of optimality. However, when (1) there exists a lag between the time employees are hired and the time they become productive and (2) flexible capacity—in the form of overtime and outsourcing—is limited, then the LP heuristic's performance deteriorates when compared with that of the optimal policy.

In these instances, randomness in turnover behavior sometimes drives the system to fall short of the LP heuristic's target capacity level, and expensive flexible capacity must be used to make up for the shortfall. The optimal policy, however, explicitly accounts for these random events and provides a buffer of extra staff whose capacity can be used.

5.1. Example Problem

The problem used for the numerical analysis is set in the context of a telephone call center:

State Space. The state space is discrete. There are two types of employee (new and experienced), and element i of the two-dimensional state vector represents the current number of employees of type- i .

Learning. If a Type-1 (new) employee does not turn over, then s/he becomes Type-2 (experienced) with probability 1. That is, $P\{\tilde{l}_{1,t} = 1\} = 1$ for all t .

Turnover. The stochastic proportion $\tilde{q}_{i,t}$ is modeled using a discretized version of a beta distribution. We choose the beta distribution for the stochastic proportions because, by varying the mean and variance, we can use it to approximate the proportions associated with binomial distributions well.

Service Requirements. \vec{D}_t is one dimensional, representing the total number of calls to be handled each period. That is, in each period, some total number of calls must be served, and subintervals are not specified. We denote this scalar (rather than vector) as D_t .

Regular-Time Employee Capacity and Costs. Each type- i employee is paid period wages and benefits of W_i and can process up to μ_i calls per period without incurring overtime charges. Employee wages are capacity neutral: $W_1/\mu_1 = W_2/\mu_2 \equiv \omega$. In period t , total regular time wages are $W_t^{tot} \stackrel{\text{def}}{=} W_1 y_t + W_2 n_{2,t}$, and total regular-time capacity is $\mu_t^{tot} \stackrel{\text{def}}{=} \mu_1 y_t + \mu_2 n_{2,t}$.

Employee Overtime Capacity and Costs. Each type- i employee may also process up to $\alpha \times \mu_i$ calls by working overtime, and overtime wages are paid at $\beta \times W_i$ and are prorated for the number of calls handled versus μ_i .

Because base wages are capacity neutral, we only need consider overtime capacity and costs in the aggregate. Thus, in period t total overtime capacity of $\alpha \times \mu_t^{tot}$ calls is available at a cost of $\beta \times \omega$ per call.

Outsourcing Capacity and Costs. An essentially unlimited number of calls can be handled using outsourcing, and an outsourcing fee of os is paid per call.

Operating Cost Function. In each period, operating costs are determined by comparing total capacity with the total service requirements, as follows:

$$O_t = \begin{cases} 0, & \text{if } D_t \leq \mu_t^{tot} \\ (D_t - \mu_t^{tot}) \times \beta\omega, & \text{if } \mu_t^{tot} < D_t \\ \alpha\mu_t^{tot} \times \beta\omega \\ + (D_t - (1 + \alpha)\mu_t^{tot}) \times os, & \text{if } (1 + \alpha)\mu_t^{tot} < D_t. \end{cases} \quad (9)$$

Note that the operating cost function assumes that an unlimited number of calls can be outsourced at a constant marginal cost. In many service environments, however, the ability to outsource may be limited or nonexistent. For example, because of confidentiality concerns, many retail banks do not outsource any of their calls. When outsourcing is limited, then the "outsourcing cost" is more accurately described as a proxy for the degradation in quality of service experienced by customers when system capacity is strained. Most likely, this cost is sharply increasing and convex.

5.2. Data for the Problem Instances

The following data concerning system dynamics are the same across all instances that we test. Each time period corresponds to a three-month quarter. Service requirements are 250,000 calls per quarter. Average turnover percentages vary by employee type. Type-1 (new) employees have an average turnover rate of 15% per quarter (47.8% annually), and Type-2 (fast) employees have a mean turnover percentage of 10% (34.4% annually). These mean rates were chosen to be consistent with typical call center data.

The examples' common cost parameters are also constructed to be consistent with data from actual call center operations. A typical employee, with a processing rate of 10,000 calls per quarter is paid a base wage of \$4,500 per quarter (\$18,000 per year), plus about 22.2% benefits, which brings the base compensation up to \$5,500 per quarter (\$22,000 per year), or \$0.55 per call. Employees with other capacities are paid in proportion to the rate at which they can handle calls. Overtime is paid at a 50% premium of the base wage (without benefits), or \$0.68 per call ($1.5 \times 4,500/10,000$). This is equivalent to $\beta = 1.23$ ($1.5/1.222$). There is a \$1,000 fixed cost for each employee hired.

We vary three sets of data across the problem instances we test.

First, we vary the outsourcing cost per call. Values include \$1, \$10, \$20, \$50, and \$100. Again, one may think of high outsourcing costs as dual prices, or proxies, for the

degradation of service that comes about when the system is caught without adequate capacity to meet demand.

Second, we vary α , the maximum amount of overtime available to the organization as a percentage of the total regular time available, from 10% to 30%. While typical work rules nominally allow the overtime limit to be 50% of regular time, in practice the percentage available tends to be much less. In particular, daily or weekly peaks in demand for service often make the use of overtime, which is available mainly at off-peak times, an ineffective means of adding marginal capacity. Because our examples do not model subintervals within each quarterly period, they do not directly reflect these scheduling constraints. To capture this effect, we therefore impose tighter limits on the theoretical maximum overtime available.

Finally, we vary the percentage increase in capacity that comes with learning. In one set of instances, experienced employees can process work at a rate that is 40% above that of new employees. In a second set, experienced employees are 80% faster than new employees. This range is consistent with capacity increases reported to us by managers of call center operations, as well as those reported by Gustafson (1982). In a last set of instances, Type-1 employees are trainees and process *no* work, while Type-2 employees have been trained and process at “full” capacity.

Actual processing rates for Type-1 and Type-2 employees were chosen so that, given the equilibrium population of employees (in all cases about 90% experienced and about 10% new in steady state) the average processing rate equals roughly 10,000 per employee. In cases in which capacity increases 40% with learning, the Type-1 and Type-2 processing rates are, respectively, 7,400 and 10,360 calls per quarter. In cases with 80% capacity increases, the Type-1 and Type-2 rates are 5,800 and 10,440 calls per quarter. In the cases with lead times, Type-1 employees have processing rates of zero, while Type-2s process 10,000 calls per quarter.

Differences in learning and overtime percentages have also prompted us to (slightly) vary the variances of the turnover fractions across the examples. More specifically, we have chosen the variances so that, given the number

of Type-1 and Type-2 employees in equilibrium (under the optimal policy), they roughly match the aggregate variances (across all states) that would have been generated by binomial distributions with the same means. Because the optimal numbers of employees vary across the cases, the standard deviations of the turnover percentages vary slightly as well. For a summary of turnover variances, see Appendix E.

5.3. The LP Heuristic

The heuristic substitutes mean rates for the random variables $\tilde{l}_{i,t}$ and $\tilde{q}_{i,t}$ and proceeds to solve (1). Given the piecewise linear form of $O(\cdot)$, the heuristic effectively solves an LP. More generally, the heuristic would solve a convex optimization problem. The time horizon used for (1) in our numerical tests is 2,500 periods, much longer than the number of periods required for the solution to reach a steady state.

The heuristic assumes that the state space is continuous. Given $n_{2,0}$ Type-2 employees, it solves an LP to find the hire-up-to number y_0 that minimizes the total cost as defined above. A separate LP is to be solved for each value of $n_{2,0}$. Note, however, that Theorem 1 guarantees the optimality of hire-up-to policies—even when $\tilde{l}_{i,t}$ and $\tilde{q}_{i,t}$ are degenerate. This significantly reduces the number of LPs that must be solved, since an LP need *not* be solved for every $(n_{1,0}, n_{2,0})$ pair.

5.4. Numerical Results

We report results of numerical tests that were made using an average-cost model. Because average costs are independent of beginning staffing levels, they are more straightforward to interpret than discounted cost results, which depend on the starting state. Similarly, the LP heuristic is run without discounting: $\alpha = 1$. For a justification of the use of average costs, please see Appendix F.

Table 1 displays the average cost per period incurred by the optimal policy, as well as the percentage increase over optimal obtained by the LP heuristic. The table shows that the LP policy performs consistently within 0.4% to 0.8% when the percentage speedup from slow to fast is 40% or

Table 1. Average cost per period in 45 example problems.

	$os = \$1$		$os = \$10$		$os = \$20$		$os = \$50$		$os = \$100$	
	OPT ¹	LP ²	OPT ¹	LP ²	OPT ¹	LP ²	OPT ¹	LP ²	OPT ¹	LP ²
30% OT 40% speedup	140.6	0.4%	140.6	0.4%	140.6	0.4%	140.6	0.4%	140.6	0.4%
20% OT 40% speedup	140.6	0.4%	140.6	0.4%	140.6	0.4%	140.6	0.4%	140.6	0.4%
10% OT 40% speedup	140.6	0.4%	140.6	0.4%	140.6	0.4%	140.6	0.4%	140.6	0.4%
30% OT 80% speedup	141.0	0.8%	141.0	0.8%	141.0	0.8%	141.0	0.8%	141.0	0.8%
20% OT 80% speedup	141.0	0.8%	141.0	0.8%	141.0	0.8%	141.0	0.8%	141.0	0.8%
10% OT 80% speedup	141.0	0.8%	141.0	0.8%	141.0	0.8%	141.0	0.8%	141.0	0.9%
30% OT lead time	154.6	2.4%	157.4	1.1%	159.4	0.5%	163.3	0.0%	167.5	0.6%
20% OT lead time	154.9	2.2%	161.9	0.0%	165.7	0.0%	172.1	3.1%	177.7	10.8%
10% OT lead time	156.1	1.7%	170.6	1.7%	176.5	7.6%	185.0	29.2%	191.9	67.4%

¹Average cost of optimal policy in \$000's.

²Percent increase of average cost using LP policy over optimal.

Table 2. Average cost per period by category in problems with lead times.

		<i>os</i> = \$1			<i>os</i> = \$20			<i>os</i> = \$100		
		OPT ¹	LP ¹	% diff. ²	OPT ¹	LP ¹	% diff. ²	OPT ¹	LP ¹	% diff. ²
30% OT	hiring	14.6	15.7	7.8%	15.1	15.7	3.7%	16.3	15.7	-3.6%
	wage	128.2	138.0	7.7%	133.2	138.0	3.6%	143.0	138.0	-3.5%
	operating	11.9	4.5	-61.8%	11.0	6.5	-40.9%	8.1	14.8	82.3%
	total	154.6	158.2	2.4%	159.4	160.2	0.5%	167.5	168.5	0.6%
20% OT	hiring	14.6	15.7	7.8%	15.8	15.7	-0.3%	17.2	15.7	-8.7%
	wage	128.2	138.0	7.7%	138.5	138.0	-0.4%	150.9	138.0	-8.5%
	operating	12.2	4.6	-62.0%	11.4	12.0	5.7%	9.6	43.3	349.8%
	total	154.9	158.3	2.2%	165.7	165.8	0.0%	177.7	197.0	10.8%
10% OT	hiring	14.6	15.7	7.8%	16.9	15.7	-7.0%	18.6	15.7	-15.6%
	wage	128.2	138.0	7.7%	148.3	138.0	-6.9%	163.1	138.0	-15.4%
	operating	13.4	5.0	-62.4%	11.4	36.2	217.7%	10.2	167.4	1543.3%
	total	156.1	158.7	1.7%	176.5	189.9	7.6%	191.9	321.1	67.4%

¹Average cost in \$000's.

²Percent increase of LP policy over optimal.

80%. In the instance with a hiring lead time, however, performance deteriorates with decreases in available overtime and with increases in outsourcing costs.

Table 2 displays the average cost per period by category for a subset of the instances with hiring lead times. From these results one can see that the LP heuristic's hiring policy does not change with the availability of overtime or with outsourcing costs: Average hiring costs remain a constant \$15,700 per period, and average wage costs \$138,000 per period. As available overtime decreases and as unit outsourcing costs increase, the LP heuristic's operating costs grow. In contrast, the optimal policy increases its hire-up-to levels as overtime becomes less freely available and as the cost of outsourcing grows.

The results of Table 2 provide insight into conditions under which the LP policy performs well. Whenever relatively inexpensive, flexible capacity is ample enough to cover for the event of high turnover, the LP policy does well. When flexible capacity is limited, however, a buffer of excess "regular-time" capacity is needed to provide for the event of excess turnover.

REMARK 3. The LP heuristic tracks differences in employee capacities, but it does not explicitly model randomness in learning and turnover behavior. A complementary heuristic approach explicitly models randomness in turnover but does not track differences among employee capacities. Such a "headcount" policy considers only the numbers of employees on hand when making hiring decisions. In Gans and Zhou (2000), we evaluate this headcount policy. The LP heuristic reported here consistently outperforms this headcount policy.

Finally, it is important to note the limitations of the numerical examples presented here. First, they are constructed with stationary demands for service, rather than with more complex periodic patterns. Second, the deterministic operating cost function does not account for errors in demand forecasting. We imagine that these factors would both strain the effectiveness of LP-based heuristics and make the computation of MDP-based policies more difficult still.

6. CONCLUSION

In this paper, we have developed an MDP approach to the employee staffing problem. Our model allows for direct representation of the stochastic nature of learning and turnover, as well as the inclusion of a very general class of single-period operating cost functions. Our analysis yields structural results that provide insight into system dynamics and control: Hire-up-to policies are optimal; when there is no learning and no hiring lead-time, myopic policies are optimal.

Our numerical results show that, when lead times and learning are not significant or flexible capacity—in the form of relatively inexpensive overtime or outsourcing—is available, then computationally efficient, LP-based, deterministic heuristics perform well. Furthermore, this efficiency is aided by the hire-up-to property identified in Theorem 1.

Conversely, the numerical results also demonstrate a consistent set of conditions under which the computational effort required by the MDP may be worthwhile. In particular, when (1) hiring lead times are significant and (2) the cost of buying incremental flexible capacity is high, then more sophisticated policies can provide a significant improvement in performance.

Thus, when there is little learning and when there is relatively inexpensive flexible capacity, optimal (myopic) or near optimal (LP) policies that are computationally efficient are likely to be available. When these conditions do not hold, however, the problem of finding effective, computationally efficient policies requires additional work.

Two future paths are worth investigating. First, we may consider developing more sophisticated heuristics that make fuller use of some of the structural properties suggested by the MDP analysis. One approach would be to extend the myopic policy of §4.3 to explicitly account for the randomness in capacity realized over one hiring lead time. Rather than solving the full MDP, a myopic heuristic would use stochastic information about learning and turnover to explicitly model the distribution of capacity one

lead time into the future, perhaps solving a stochastic program to determine the appropriate number to hire. Another approach would be to continue to use LP policies, adding a safety margin to the capacity target to be hit in each period. Alternatively, we may consider simulation-based optimization techniques that do not require a great deal of additional specialized structure to be effective. In either case, further work is required to validate and evaluate the effectiveness of the techniques.

Additional work must also be done to develop practical implementations of this general approach. As we noted in the introduction, large call centers use commercial software to solve the integer program associated with $O_t(\cdot)$. In principle, it should be feasible to extend these commercial systems to (1) solve the type of long-term staffing problem described in this paper, and (2) include in the model a one-period operating cost function that is based on an LP relaxation of the scheduling IP. To our knowledge, no such implementation yet exists.

APPENDIX A. PROOF OF PROPOSITION 1

The proof follows the argument for the convexity of an LP’s objective function value with respect to its right-hand side. Let \bar{n}^1 and \bar{n}^2 be the right-hand sides of two (otherwise identical) instances of $O(\cdot)$, and let \bar{x}^1 and \bar{x}^2 be the corresponding optimal solutions. Then condition (i) ensures the convexity of the feasible region, so that for any $\lambda \in [0, 1]$, $(\lambda\bar{x}^1 + (1 - \lambda)\bar{x}^2)$ is feasible for the right-hand side $(\lambda\bar{n}^1 + (1 - \lambda)\bar{n}^2)$. In addition, condition (ii) implies that for every $C_{iw}(\cdot)$ we have $\lambda C_{iw}(x_{iw}^1) + (1 - \lambda)C_{iw}(x_{iw}^2) \geq C_{iw}(\lambda x_{iw}^1 + (1 - \lambda)x_{iw}^2)$, and for every $OS_{sj}(\cdot)$ we have $\lambda OS_{sj}(z_{sj}^1) + (1 - \lambda)OS_{sj}(z_{sj}^2) \geq OS_{sj}(\lambda z_{sj}^1 + (1 - \lambda)z_{sj}^2)$. \square

APPENDIX B. PROOF OF THEOREM 1

We will not directly analyze the mathematical program (1) to demonstrate Theorem 1. Instead, we analyze the following equivalent MDP problem. (See Puterman 1994, §4.2 and §6.1 for this equivalence.)

Let $V_t(n_{1,t}, \dots, n_{m,t}|T)$ denote the total discounted future cost at the beginning of period t when the numbers of employees on hand are $(n_{1,t}, \dots, n_{m,t})$ and the planning horizon is T time periods. Then we have the following recursive expression:

$$\begin{aligned} &V_t(n_{1,t}, \dots, n_{m,t}|T) \\ &= \min_{x_t \geq 0} \left\{ hx_t + W_1(n_{1,t} + x_t) \right. \\ &\quad \left. + \sum_{i=2}^m W_i n_{i,t} + O_t(n_{1,t} + x_t, n_{2,t}, \dots, n_{m,t}) \right. \\ &\quad \left. + \alpha \mathbf{E}_{\{\tilde{q}_{1,t}, \dots, \tilde{l}_{1,t}, \dots\}} \{V_{t+1}(n_{1,t+1}, \dots, n_{m,t+1}|T)\} \right\} \\ &= \min_{y_t \geq n_{1,t}} \{J_t(y_t, n_{2,t}, \dots, n_{m,t}|T)\} \\ &\quad - hn_{1,t} + \sum_{i=2}^m W_i n_{i,t}, \end{aligned}$$

where

$$\begin{aligned} &J_t(y_t, n_{2,t}, \dots, n_{m,t}|T) \\ &= H_t(y_t, n_{2,t}, \dots, n_{m,t}) \\ &\quad + \alpha \mathbf{E}_{\{\tilde{q}_{1,t}, \dots, \tilde{l}_{1,t}, \dots\}} \{V_{t+1}(n_{1,t+1}, \dots, n_{m,t+1}|T)\}, \\ &H_t(y_t, n_{2,t}, \dots, n_{m,t}) = (h + W_1)y_t + O_t(y_t, n_{2,t}, \dots, n_{m,t}), \end{aligned}$$

and (3), (4), and (5) hold.

REMARK 4. Note that $O_t(y_t, \dots, n_{m,t})$, and therefore $J_t(y_t, \dots, n_{m,t}|T)$, is not defined for $y_t < 0$. Later in this appendix, we will prove the convexity of J_t , but for now, suppose convexity holds for J_t . When $\frac{\partial J_t}{\partial y_t}(0, \dots, n_{m,t}|T) \leq 0$, the minimum of $J_t(y_t, \dots, n_{m,t}|T)$ is achieved at $y_t^*(n_{2,t}, \dots, n_{m,t}|T) \geq 0$, and there is no need for this definition. However, when $\frac{\partial J_t}{\partial y_t}(0, \dots, n_{m,t}|T) > 0$, it is optimal to not hire: $y_t^*(n_{2,t}, \dots, n_{m,t}|T) = 0$. In this case, we extend the definition of J_t so that $J_t(y_t, \dots, n_{m,t}|T) = J_t(0, \dots, n_{m,t}|T)$ for $y_t < 0$. This extension makes the first-order condition, that

$$\lim_{y_t \rightarrow y_t^*(n_{2,t}, \dots, n_{m,t}|T)^-} \frac{\partial J_t}{\partial y_t}(y_t, n_{2,t}, \dots, n_{m,t}|T) = 0,$$

valid, and it preserves the convexity of J_t .

We now formally state and prove Theorem 1:

THEOREM 1.

(i) For $T < \infty$, $J_t(y_t, n_{2,t}, \dots, n_{m,t}|T)$ is jointly convex in $(y_t, n_{2,t}, \dots, n_{m,t})$ and $V_t(n_{1,t}, \dots, n_{m,t}|T)$ is jointly convex in $(n_{1,t}, \dots, n_{m,t})$ for all t .

(ii) For $T = \infty$, under Assumptions 1 and 2, $J_t(y_t, n_{2,t}, \dots, n_{m,t}|\infty)$ is jointly convex in $(y_t, n_{2,t}, \dots, n_{m,t})$ and $V_t(n_{1,t}, \dots, n_{m,t}|\infty)$ is jointly convex in $(n_{1,t}, \dots, n_{m,t})$ for all t .

Therefore, for both cases, in any time period a policy of the hire-up-to type is optimal.

PROOF OF PART (i). First, let T be finite and fixed. We use induction.

For $t = T + 1$, Formulation (1) implies $V_{T+1}(n_{1,T+1}, \dots, n_{m,T+1}|T) \equiv 0$. Then $V_{T+1}(\cdot|T)$ is (trivially) jointly convex in $(n_{1,T+1}, \dots, n_{m,T+1})$. Next, suppose $V_{t+1}(n_{1,t+1}, n_{2,t+1}, \dots, n_{m,t+1}|T)$ is jointly convex in $(n_{1,t+1}, n_{2,t+1}, \dots, n_{m,t+1})$. Then in two steps we prove that $J_t(y_t, \dots, n_{m,t}|T)$ and $V_t(n_{1,t}, \dots, n_{m,t}|T)$ are, in turn, jointly convex in $(y_t, \dots, n_{m,t})$ and $(n_{1,t}, \dots, n_{m,t})$ respectively.

Step 1. We first prove that $J_t(y_t, n_{2,t}, \dots, n_{m,t}|T)$ is jointly convex in $(y_t, n_{2,t}, \dots, n_{m,t})$. Because of the assumption that $O_t(y_t, n_{2,t}, \dots, n_{m,t})$ is jointly convex in $(y_t, n_{2,t}, \dots, n_{m,t})$, $H_t(y_t, n_{2,t}, \dots, n_{m,t})$ is also jointly convex in $(y_t, n_{2,t}, \dots, n_{m,t})$. Because the integral of a convex function is again convex, we only need to prove that $V_{t+1}(n_{1,t+1}, n_{2,t+1}, \dots, n_{m,t+1}|T)$ is jointly convex in $(y_t, n_{2,t}, \dots, n_{m,t})$ for any realization of $(\tilde{q}_{1,t}, \dots, \tilde{q}_{m,t}, \tilde{l}_{1,t}, \dots, \tilde{l}_{m-1,t})$. This is true because of Lemma 1 and the fact—from (3), (4), and (5)—that for fixed $(\tilde{q}_{1,t}, \dots, \tilde{q}_{m,t}, \tilde{l}_{1,t}, \dots, \tilde{l}_{m-1,t})$, $n_{1,t+1}, n_{2,t+1}, \dots, n_{m,t+1}$ are all linear in $(y_t, n_{2,t}, \dots, n_{m,t})$.

Step 2. Now, given that $J_t(y_t, n_{2,t}, \dots, n_{m,t}|T)$ is jointly convex in $(y_t, n_{2,t}, \dots, n_{m,t})$, we prove that $V_t(n_{1,t}, \dots, n_{m,t}|T)$ is jointly convex in $(n_{1,t}, \dots, n_{m,t})$. This follows from a direct application of Lemma 2. Note that $-hn_{1,t} + \sum_{i=2}^m W_i n_{i,t}$ is jointly convex in $(n_{1,t}, \dots, n_{m,t})$.

Therefore, by repeating Steps 1 and 2, we know that $J_t(y_t, n_{2,t}, \dots, n_{m,t}|T)$ is jointly convex in $(y_t, n_{2,t}, \dots, n_{m,t})$ and $V_t(n_{1,t}, \dots, n_{m,t}|T)$ is jointly convex in $(n_{1,t}, \dots, n_{m,t})$ for all t .

As a result of $J_t(n_{1,t}, \dots, n_{m,t}|T)$ being jointly convex in $(n_{1,t}, \dots, n_{m,t})$ for all t we know that there exists $y_t^*(n_{2,t}, \dots, n_{m,t}|T)$ such that $y_t^*(n_{2,t}, \dots, n_{m,t}|T)$ minimizes $J_t(y_t, n_{2,t}, \dots, n_{m,t}|T)$ without the constraint $y_t \geq n_{1,t}$. When there are multiple minimizers to $J_t(y_t, n_{2,t}, \dots, n_{m,t}|T)$, we will take the smallest. Furthermore, given any $(n_{1,t}, n_{2,t}, \dots, n_{m,t})$, because $J_t(y_t, n_{2,t}, \dots, n_{m,t}|T)$ is convex, the optimal hiring policy with the constraint will be to hire up to $y_t^*(n_{2,t}, \dots, n_{m,t}|T)$ if $n_{1,t} < y_t^*(n_{2,t}, \dots, n_{m,t}|T)$, and not hire (and thus remain $n_{1,t}$) otherwise. This is exactly the hire-up-to policy of Definition 1. \square

To prove part (ii) of the theorem, we will need the following lemma.

LEMMA 3. $V_t(\cdot|T)$ is increasing in T .

PROOF OF LEMMA 3. Pick any two cost functions $V_{t+1}^1(\cdot)$ and $V_{t+1}^2(\cdot)$ such that $V_{t+1}^1(n_{1,t+1}, \dots, n_{m,t+1}|T) \leq V_{t+1}^2(n_{1,t+1}, \dots, n_{m,t+1}|T)$ for any $(n_{1,t+1}, \dots, n_{m,t+1})$. Moreover, let the minima of $V_t^1(\cdot)$ and $V_t^2(\cdot)$ be achieved at y_t^{*1} and y_t^{*2} , respectively. Then

$$\begin{aligned}
& V_t^1(n_{1,t}, \dots, n_{m,t}|T) \\
&= H_t(y_t^{*1}, n_{2,t}, \dots, n_{m,t}) \\
&\quad + \alpha \mathbf{E}_{\{\tilde{q}_{1,t}, \dots, \tilde{l}_{1,t}, \dots\}} \{V_{t+1}^1(n_{1,t+1}^1, \dots, n_{m,t+1}^1|T)\} \\
&\quad - hn_{1,t} + \sum_{i=2}^m W_i n_{i,t} \\
&\leq H_t(y_t^{*2}, n_{2,t}, \dots, n_{m,t}) \\
&\quad + \alpha \mathbf{E}_{\{\tilde{q}_{1,t}, \dots, \tilde{l}_{1,t}, \dots\}} \{V_{t+1}^1(n_{1,t+1}^2, \dots, n_{m,t+1}^2|T)\} \\
&\quad - hn_{1,t} + \sum_{i=2}^m W_i n_{i,t} \\
&\leq H_t(y_t^{*2}, n_{2,t}, \dots, n_{m,t}) \\
&\quad + \alpha \mathbf{E}_{\{\tilde{q}_{1,t}, \dots, \tilde{l}_{1,t}, \dots\}} \{V_{t+1}^2(n_{1,t+1}^2, \dots, n_{m,t+1}^2|T)\} \\
&\quad - hn_{1,t} + \sum_{i=2}^m W_i n_{i,t} \\
&= V_t^2(n_{1,t}, \dots, n_{m,t}|T).
\end{aligned}$$

Therefore, the MDP value iteration is an increasing mapping. When T is increased to $T+1$, the values of $V_t(\cdot|T)$ change as if the end-of-horizon cost function had been

increased from 0 to a positive function. Therefore, by applying the monotonicity of the mapping repeatedly, we find that the function $V_t(\cdot|T)$ also increases. \square

PROOF OF PART (ii). It is not difficult to see that, because total one-period cost is bounded by K (due to Assumption 2), $V_t(\cdot|T) \leq \frac{K}{1-\alpha}$ for any T and $t \leq T$.

Thus, $\{V_t(\cdot|T) : T = 0, 1, \dots\}$ is a monotone, bounded sequence. Because monotone bounded sequences always converge, if we let $T \rightarrow \infty$, then

$$\lim_{T \rightarrow \infty} V_t(n_{1,t}, \dots, n_{m,t}|T) = U_t(n_{1,t}, \dots, n_{m,t}) \quad (10)$$

for some finite function $U_t(n_{1,t}, \dots, n_{m,t})$. Moreover, from Theorem 10.8 in Rockafellar (1970), $V_t(n_{1,t}, \dots, n_{m,t}|T)$ converges to $U_t(n_{1,t}, \dots, n_{m,t})$ uniformly, and $U_t(n_{1,t}, \dots, n_{m,t})$ is jointly convex in its arguments. By applying Lebesgue's Dominated Convergence Theorem, we obtain

$$\begin{aligned}
& U_t(n_{1,t}, \dots, n_{m,t}) \\
&= \min_{y_t \geq n_{1,t}} \{H_t(y_t, n_{2,t}, \dots, n_{m,t}) \\
&\quad + \alpha \mathbf{E}_{\{\tilde{q}_{1,t}, \dots, \tilde{l}_{1,t}, \dots\}} \{U_{t+1}(n_{1,t+1}, \dots, n_{m,t+1})\}\} \\
&\quad - hn_{1,t} + \sum_{i=2}^m W_i n_{i,t}.
\end{aligned}$$

Hence, these $U_t(\cdot)$ s satisfy the MDP optimality equation. As a result, they are the infinite-horizon cost functions. That is,

$$\begin{aligned}
V_t(n_{1,t}, \dots, n_{m,t}|\infty) &= U_t(n_{1,t}, \dots, n_{m,t}) \\
&= \lim_{T \rightarrow \infty} V_t(n_{1,t}, \dots, n_{m,t}|T).
\end{aligned}$$

Therefore, the infinite-horizon cost functions $V_t(\cdot|\infty)$ are jointly convex, and the hire-up-to policy is, again, optimal in the infinite-horizon case. \square

APPENDIX C. PROOF OF THEOREM 2

Given Assumption 3, we define the ‘‘staffing position’’ as follows. Let $r_{i,t}^s \stackrel{\text{def}}{=} \prod_{\tau=0}^{s-1} (1 - q_{\min\{\lambda+1, t+\tau\}, t+\tau})$ be the fraction of type- i employees on hand in period t that remain employed after s periods. Then the staffing position at t before hiring is $\hat{n}_t \stackrel{\text{def}}{=} \sum_{i=2}^{\lambda+1} n_{i,t} r_{i,t}^\lambda$, and the net number hired is $\hat{x}_t \stackrel{\text{def}}{=} x_t r_{1,t}^\lambda$. Thus, \hat{n}_t and \hat{x}_t are the analogues of $n_{1,t+\lambda}$ and $x_{t+\lambda}$ in a system with no hiring lead time.

We can also aggregate wage costs over one lead time to define the effective hiring cost of one person. More specifically, the future cost of hiring enough people in period t to yield one person in period $t+\lambda$ is $\hat{h}_t \stackrel{\text{def}}{=} \sum_{s=0}^{\lambda-1} \frac{W_{1+t+s}}{\alpha^{\lambda-s} r_{1,t+s}^{\lambda-s}}$.

Here each W_{1+s} is “grossed up” to account for two factors. First $1/r_{1+s}^{\lambda-s}$ type-(1+s) people must be in the pipeline in period $t+s$ to yield one person in period $t+\lambda$. Second, every dollar paid in period $t+s$ is equivalent to $1/\alpha^{\lambda-s}$ dollars paid in period $t+\lambda$.

Then using \hat{n}_t , \hat{x}_t , and \hat{h}_t , we may solve the hiring problem as if $m=1$. Note, however, that the optimal hiring number \hat{x}_t^* corresponds to the net number of new hires remaining in period $t+\lambda$. When implementing the policy, the actual number of people hired in period t is

$$x_t^* = \hat{x}_t^*/r_{1,t}^\lambda. \quad (11)$$

With staffing position precisely defined, we are ready to formally state and prove the theorem:

THEOREM 2. *Suppose the condition listed in the previous statement of the theorem holds, and let $y_t^G = \arg \min_{y_t} G_t(y_t)$. (Choose the smallest y_t^G if multiple minimizers exist.) Then*

(i) *for both the infinite-horizon case and the finite-horizon case with end-of-horizon cost function $V_{T+1}(n_{T+1}) = -\alpha hn_{T+1}$, if the service requirements are stationary or stochastically increasing, $y_t^G \leq y_{t+1}^G$ for all t ;*

(ii) *if $y_t^G \leq y_{t+1}^G$ for all t , then the myopic policy of minimizing $G_t(y_t)$ in each period t is optimal. That is, the hire-up-to numbers in the optimal policy are $(y_0^*, y_1^*, \dots) = (y_0^G, y_1^G, \dots)$.*

Therefore, when service requirements are stationary or stochastically increasing, the myopic policy is optimal. In steady state, with stationary requirements, $x_{t+1}^ = \tilde{q}_t y_t^G$, and with increasing requirements, $x_{t+1}^* = (y_{t+1}^G - y_t^G) + \tilde{q}_t y_t^G$.*

PROOF. Given any hiring policy, π ,

$$\begin{aligned} & V(n_0 | \pi; \vec{D}_t) \\ &= \mathbf{E}_{\{\{\tilde{r}_t\}_{t=0}^\infty\}} \left\{ \sum_{t=0}^{\infty} \alpha^t [h(y_t - n_t) + W y_t + O_t(y_t; \vec{D}_t)] \right\} \\ &= \mathbf{E}_{\{\{\tilde{r}_t\}_{t=0}^\infty\}} \left\{ \sum_{t=0}^{\infty} \alpha^t [(h+W)y_t + O_t(y_t; \vec{D}_t)] \right\} \\ &\quad - \mathbf{E}_{\{\{\tilde{r}_t\}_{t=0}^\infty\}} \left\{ \sum_{t=1}^{\infty} \alpha^t h n_t \right\} - h n_0 \\ &\stackrel{(**)}{=} -h n_0 + \mathbf{E}_{\{\{\tilde{r}_t\}_{t=0}^\infty\}} \left\{ \sum_{t=0}^{\infty} \alpha^t [(h+W)y_t + O_t(y_t; \vec{D}_t)] \right\} \\ &\quad - \mathbf{E}_{\{\{\tilde{r}_t\}_{t=0}^\infty\}} \left\{ \sum_{t=0}^{\infty} \alpha^{t+1} h \tilde{r}_t y_t \right\} \\ &= -h n_0 + \mathbf{E}_{\{\{\tilde{r}_t\}_{t=0}^\infty\}} \left\{ \sum_{t=0}^{\infty} \alpha^t [(1 - \alpha \tilde{r}_t)h + W] y_t + O_t(y_t) \right\} \\ &= -h n_0 + \mathbf{E}_{\{\{\tilde{r}_t\}_{t=0}^\infty\}} \left\{ \sum_{t=0}^{\infty} \alpha^t G_t(y_t; \vec{D}_t) \right\}, \quad (12) \end{aligned}$$

where $y_t \geq n_t$ and $n_{t+1} = \tilde{r}_t y_t$.

Equality (**) holds because $n_{\tau+1}$ is independent of $\{\tilde{r}_t\}_{t=\tau+1}^\infty$ and y_τ is independent of $\{\tilde{r}_t\}_{t=\tau}^\infty$, which implies that for any τ , $\mathbf{E}_{\{\{\tilde{r}_t\}_{t=0}^\infty\}} \{n_{\tau+1}\} = \mathbf{E}_{\{\{\tilde{r}_t\}_{t=0}^\infty\}} \{n_{\tau+1}\} = r_\tau \mathbf{E}_{\{\{\tilde{r}_t\}_{t=0}^\infty\}} \{y_\tau\} = r_\tau \mathbf{E}_{\{\{\tilde{r}_t\}_{t=0}^\infty\}} \{y_\tau\}$. Note that the above transformation (12) also holds in the finite horizon case, if we assume an end-of-horizon cost function of $V_{T+1}(n_{T+1}) = -\alpha hn_{T+1}$.

Part (i). When service requirements are stationary, $G_t(y_t)$ is the same for all t , therefore $y_t^G = y_{t+1}^G$ for all t . When service requirements are stochastically increasing, $\vec{D}_t \leq \vec{D}_{t+1}$, then, by Assumption 4, $O'_t(y_t; \vec{D}_t) \geq O'_{t+1}(y_t; \vec{D}_{t+1})$ and $G'_t(y) \geq G'_{t+1}(y)$. First-order conditions state that $\lim_{y \rightarrow y_t^{G-}} G'_t(y) = \lim_{y \rightarrow y_{t+1}^{G-}} G'_{t+1}(y) = 0$. Hence, the convexity of both $G_t(\cdot)$ and $G_{t+1}(\cdot)$ implies $y_t^G \leq y_{t+1}^G$.

Part (ii). We note that even though the costs underlying $G_t(\cdot)$ in the staffing problem differ from the costs underlying one-period costs in inventory problems, the convexity and separability of the $G_t(\cdot)$ s, along with the fact that the sequence of solutions $\{y_t^G\}$ is increasing, are sufficient to ensure that classic arguments from inventory theory hold. For example, see Proposition 3–2 and Theorem 3–1 in Heyman and Sobel (1982).

As a result, in steady state, when service requirements are stationary, $y_t = y_t^* = y_t^G$, then $n_{t+1} = \tilde{r}_t y_t = \tilde{r}_t y_t^G$, and $x_{t+1}^* = y_{t+1}^* - n_{t+1} = (1 - \tilde{r}_t) y_t^* = \tilde{q}_t y_t^G$. When the service requirements are increasing, so are y_t^G . And $x_{t+1}^* = y_{t+1}^* - n_{t+1} = (y_{t+1}^G - y_t^G) + \tilde{q}_t y_t^G$. \square

Note that when $m > 1$ and Assumption 3 holds, the myopic policy uses \hat{n}_t as the system state before hiring, \hat{h}_t as the hiring cost, and $W_{\lambda+1}$ as the wage cost within (8) to determine “net” hiring numbers $\{\hat{x}_t^*\}$. It then calculates optimal hiring numbers $\{x_t^*\}$ from $\{\hat{x}_t^*\}$ using (11).

APPENDIX D. PROOF OF THEOREM 3

For part (i), we first prove that if $y_t^* \leq y_{t+1}^*$, then the myopic policy is optimal in period t , i.e., $y_t^* = y_t^G = \arg \min_{y_t} G_t(y_t)$. Then if we let y_t^* be the smallest hire-up-to number among $y_0^*, y_1^*, \dots, y_{k-1}^*$, the myopic policy is optimal in period t .

To prove this, we note that since y_t^G is the smallest minimizer of $G_t(y_t)$, it is also the smallest value such that its left derivative $\lim_{y_t \rightarrow y_t^{G-}} G'_t(y_t) = 0$. Similarly, y_t^* is the smallest value such that $\lim_{y_t \rightarrow y_t^{*-}} J'_t(y_t) = 0$.

Because $y_t^* \leq y_{t+1}^*$, when $y_t \leq y_t^*$, $J_t(y_t) = (h+W)y_t + O_t(y_t) + \alpha(J_{t+1}(y_{t+1}^*) - h r_t y_t) = G_t(y_t) + \alpha J_{t+1}(y_{t+1}^*)$. Therefore, when $y_t \leq y_t^*$, $G'_t(y_t) = J'_t(y_t)$. Hence, $\lim_{y_t \rightarrow y_t^{*-}} G'_t(y_t) = \lim_{y_t \rightarrow y_t^{*-}} J'_t(y_t) = 0$, and y_t^{*-} is the smallest such number. Therefore $y_t^G = y_t^*$.

Part (ii) then follows directly. \square

Again, when $m > 1$ the optimal net staffing numbers obtained in the zero lead time analogue, $\{\hat{x}_t\}$, must be transformed using (11) to obtain the optimal hiring numbers $\{x_t^*\}$.

APPENDIX E. DATA USED IN THE NUMERICAL ANALYSIS

Table 3. Turnover rate variances used in example problems.

OT ¹	CI ²	os = \$1		os = \$10		os = \$20		os = \$50		os = \$100	
		Type-1	Type-2	Type-1	Type-2	Type-1	Type-2	Type-1	Type-2	Type-1	Type-2
30%	40%	0.0519	0.0041	0.0519	0.0041	0.0519	0.0041	0.0519	0.0041	0.0519	0.0041
20%	40%	0.0519	0.0041	0.0519	0.0041	0.0519	0.0041	0.0519	0.0041	0.0519	0.0041
10%	40%	0.0519	0.0041	0.0518	0.0041	0.0518	0.0041	0.0518	0.0041	0.0518	0.0041
30%	80%	0.0519	0.0041	0.0519	0.0041	0.0519	0.0041	0.0519	0.0041	0.0519	0.0041
20%	80%	0.0519	0.0041	0.0519	0.0041	0.0519	0.0041	0.0519	0.0041	0.0519	0.0041
10%	80%	0.0516	0.0041	0.0515	0.0041	0.0515	0.0041	0.0515	0.0041	0.0515	0.0041
30%	∞	0.0611	0.0039	0.0596	0.0037	0.0594	0.0037	0.0580	0.0036	0.0566	0.0035
20%	∞	0.0611	0.0039	0.0582	0.0036	0.0643	0.0036	0.0529	0.0034	0.0511	0.0033
10%	∞	0.0611	0.0039	0.0621	0.0035	0.0607	0.0034	0.0528	0.0032	0.0559	0.0031

¹OverTime available as a percentage of regular time.

²Percentage Capacity Increase as employee progresses from Type-1 to Type-2. "∞" is the lead-time case.

APPENDIX F. AVERAGE COST AS THE METHOD OF COMPARISON

When the planning horizon is finite, the average-cost problem is a special case of the discounted-cost problem (with $\alpha = 1$). When planning horizon is infinite, the optimal average-cost policy generates a constant average one-period cost (for all the states that are in the same chain of the MDP), rather than the discounted costs that may depend on the system's starting state. The state-independent nature of the average cost makes the comparison of alternative policies and MDP systems clearer.

To use average costs, we must first develop a set of properties for the optimal hiring policy under the average-cost criterion. The proposition below assumes that there exists an average-cost optimal policy and its corresponding value function.

PROPOSITION 2. Let $V_0^\alpha(\cdot)$ be the α -discounted cost function, and suppose the average-cost function, $\lim_{T \rightarrow \infty} \sum_{t=0}^{T-1} [hx_t + W_1 y_t + \sum_{i=2}^m W_i n_{i,t} + O_t(y_t, n_{2,t}, \dots, n_{m,t})] / T$ exists. Then

(i) $\lim_{\alpha \rightarrow 1} (1 - \alpha) V_0^\alpha(n_{1,0}, \dots, n_{m,0})$ exists, and

$$\lim_{T \rightarrow \infty} \sum_{t=0}^{T-1} \frac{1}{T} \left[hx_t + W_1 y_t + \sum_{i=2}^m W_i n_{i,t} + O_t(y_t, n_{2,t}, \dots, n_{m,t}) \right] = \lim_{\alpha \rightarrow 1} (1 - \alpha) V_0^\alpha(n_{1,0}, \dots, n_{m,0});$$

(ii) the average-cost function is convex and the optimal hiring policy is of the hire-up-to type.

PROOF. Part (i) of the proposition follows from Proposition 4–7 in Heyman and Sobel (1982). Part (ii) follows naturally from (i) since the convexity of discounted-cost functions has already been established in Theorem 1. \square

Part (i) of the proposition establishes the relationship between the discounted-cost and average-cost criteria: When α is close to 1, the average cost function divided

by $(1 - \alpha)$ can be used to approximate the discounted cost function. Therefore comparisons based on an average-cost criterion can be used to infer comparisons based on discounted costs as well.

Note that we do not prove the existence of an average-cost optimal policy for our system, so technically we should regard our numerical results with some caution. Nevertheless, in all of our numerical examples, MDP value iteration has converged without problems. Furthermore, we find the clarity of the numerical comparisons afforded by the average-cost model to outweigh this technical limitation.

ACKNOWLEDGMENTS

The authors thank Avishai Mandelbaum for the data displayed in Figure 1, and they are grateful to the referees and area editor for their many helpful comments. This research was supported by the Wharton Financial Institutions Center and by NSF Grant SBR-9733739.

REFERENCES

- Akşin, O. Zeynep. 1999. Accounting for human assets in a service sector firm. Working paper, INSEAD, Fontainebleau, France.
- Arrow, Kenneth J., Samuel Karlin, Herbert Scarf. 1958. *Studies in the Mathematical Theory of Inventory and Production*. Stanford University Press, Stanford, CA.
- Aviv, Yossi, Awi Federgruen. 1997. Stochastic inventory models with limited production capacity and periodically varying parameters. *Probab. Engrg. Inform. Sci.* **11**(1) 107–135.
- Bartholomew, D. J., A. F. Forbes, S. I. McClean. 1991. *Statistical Techniques for Manpower Planning*, 2nd ed. John Wiley and Sons, New York.
- Berman, Oded, Richard C. Larson, Edieal Pinker. 1997. Scheduling workforce and workflow in a high volume factory. *Management Sci.* **43**(2) 158–172.
- Bordoloi, Sanjeev, Hirofumi Matsuo. 2002. Human resource planning in knowledge-intensive operations: A model for learning with stochastic turnover. *Eur. J. Oper. Res.* **130**(1) 169–189.

- Ebert, Ronald J. 1976. Aggregate planning with learning curve productivity. *Management Sci.* **23**(2) 171–182.
- Gaimon, Cheryl, Gerald L. Thompson. 1984. A distributed parameter cohort personnel planning model that uses cross-sectional data. *Management Sci.* **30**(6) 750–764.
- Gans, Noah, Yong-Pin Zhou. 2000. Managing learning and turnover in employee staffing. Working paper 99-07-05, OPIM Department, The Wharton School, University of Pennsylvania, Philadelphia, PA.
- Gerchak, Yigal, Mahmut Parlar, S. Sankar Sengupta. 1990. On manpower planning in the presence of learning. *Engrg. Costs Production Econom.* **20** 295–303.
- Grinold, Richard C., Robert E. Stanford. 1974. Optimal control of a graded manpower system. *Management Sci.* **20**(8) 1201–1216.
- Gustafson, H. W. 1982. Force-loss cost analysis. W. H. Mobley, ed. *Employee Turnover: Causes, Consequences, and Control*. Addison-Wesley, Reading, MA.
- Heyman, Daniel P., Matthew J. Sobel. 1982. *Stochastic Models in Operations Research. Volume II: Stochastic Optimization*. McGraw-Hill Book Company, New York.
- Holt, C. C., F. Modigliani, J. Muth, H. A. Simon. 1960. *Planning Production, Inventory and Work Force*. Prentice-Hall, Englewood Cliffs, NJ.
- Iglehart, Donald L., Stratton C. Jaquette. 1969. Multi-class inventory models with demand a function of inventory level. *Naval Res. Logist. Quart.* **16** 495–502.
- Jovanovic, Boyan. 1979. Job matching and the theory of turnover. *J. Political Econom.* **87**(5) 972–990.
- Kapuściński, Roman, Sridhar Tayur. 1998. A capacitated production-inventory model with periodic demand. *Oper. Res.* **46**(6) 899–911.
- Karlin, Samuel. 1955. The structure of dynamic programming models. *Naval Res. Logist. Quart.* **2** 285–294.
- . 1960a. Dynamic inventory policy with varying stochastic demands. *Management Sci.* **6**(3) 231–258.
- . 1960b. Optimal policy for dynamic inventory process with stochastic demands subject to seasonal variations. *J. SIAM* **8**(4) 611–629.
- . 1968. *Total Positivity*, Volume 1. Stanford University Press, Stanford, CA.
- Khoshnevis, Behrokh, Philip M. Wolfe. 1986. A short-cycle product aggregate planning model incorporating improvement curve productivity. *Engrg. Costs Production Econom.* **10** 217–233.
- Mandelbaum, Avishai, Anat Sakov, Sergei Zeltyn. 2000. Empirical analysis of a call center. Working paper, Technion, Haifa, Israel.
- Orrbeck, M. G., D. R. Schuette, H. E. Thompson. 1968. The effect of worker productivity on production smoothing. *Management Sci.*, **14**(6) B332–B342.
- Pinker, Edieal, Robert Shumsky. 2000. The efficiency-quality tradeoff of crosstrained workers. *Manufacturing & Service Oper. Management* **2**(1) 32–48.
- Puterman, Martin L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Mathematical Statistics. John Wiley and Sons, New York.
- Rockafellar, R. Tyrrell. 1970. *Convex Analysis*. Princeton University Press, Princeton, NJ.
- Veinott, Arthur F., Jr. 1965. Optimal policy for a multi-product, dynamic, nonstationary inventory problem. *Management Sci.* **12**(3) 206–222.
- Yelle, Louis E. 1979. The learning curve: Historical review and comprehensive survey. *Decision Sci.* **10** 302–328.
- Zipkin, Paul. 1989. Critical number policies for inventory models with periodic data. *Management Sci.* **35**(1) 71–80.