

CAPACITY PLANNING IN A VIRTUAL CALL CENTER USING AUCTIONS TO PERFORM ROSTERING

Matthew F. Keblis
College of Business
University of Dallas
Irving, TX 75062
Phone: (972) 265-5719
Fax: (972) 265-5750
Email: mkeblis@udallas.edu

ABSTRACT

In this paper we consider a capacity planning problem facing the management of a virtual call center. More specifically, we study the strategic or long-term problem of determining the appropriate number of call center agents to hire and train in order to enable the hour-to-hour call center coverage desired. This is a somewhat classic problem, but it is made unique in this instance by the fact that the underlying tactical planning process, an integral part of the long-term planning process, makes use of auctions to perform rostering. We explore the use of hire up-to policies, similar to order up-to policies from inventory theory, for addressing the strategic planning problem. Simulation is used to identify optimal hire up-to levels.

Call centers continue to grow in importance. In many industries they are increasingly becoming the primary interface between a firm and its customers. The traditional call center is a large room filled with open-air cubicles each containing a person and equipment. Each person, herein referred to as an agent, typically sits in front of a computer wearing a telephone headset. The telephone headset allows the agent to communicate with a customer of the organization and the computer provides the agent with that information needed to effect a successful customer interaction.

Although call centers of the sort described above remain the norm, an increasing number of organizations have ceased to operate physical call centers, opting instead to allow all their agents to work remotely. These organizations have thus become operators of virtual call centers. This mode of operation is made more attractive with refinement of Voice over Internet Protocol (VoIP), a communications technology that enables voice calls to be inexpensively transmitted over packet switched networks, e.g. the Internet. According to the industry trade press, adoption of the virtual call center model allows an organization to lower infrastructure costs (building and equipment expenditures), gain access to more than just a local labor market, eliminate the commute of agents and hence much of their tardiness as well as absenteeism, lower employee turnover rate, and better match labor supply with labor demand.

The problem of matching labor supply with labor demand in a virtual call center, or capacity planning for short, is examined in this paper. This problem exists at essentially two levels. At the lowest level it is about call center management grappling with having the right number of agents manning the phones hour-to-hour, if not minute-to-minute; this is a tactical capacity planning problem. At the highest level, it is about that same management determining the appropriate number of agents to hire and train in order to enable the hour-to-hour coverage desired; this is a strategic capacity planning problem. In this paper we focus on the strategic planning problem. However as will become clear below, the tactical planning approach employed by a firm is an integral aspect of a firm's strategic planning, hence we shall also detail the underlying tactical planning problem.

The traditional approach to capacity planning in a call center has four steps: forecasting, staffing, scheduling, and rostering or hiring, where the last step depends on whether the planning is tactical or strategic in nature (see Gans et al. (2003) for a more detailed coverage of capacity planning in a call center). The goal of the forecasting step is to generate an estimate of the number of voice calls that will be received in each period (typically an hour or some fraction of one) during some block of time in the future, where the block of time in the future is at least several days and not uncommonly several weeks out. After forecasts have been generated for the future block of time, staffing is performed for each period within the block with the goal of determining the number of agents required to attain the objective of management. In call centers where the primary mission is to resolve customer troubles or complaints, service level objectives are typically the norm, e.g. X percent of calls answered in Y or less seconds. On the other hand in call centers that are revenue generating, i.e. where each caller is a potential product purchaser, economic objectives are regularly used to arrive at staffing decisions. Once staffing requirements are determined, the task of scheduling when during the future time block that agents will work is undertaken. This step amounts to deciding the number of agents to assign to each of a set of work schedules, where each work schedule consists of a collection of shifts, each shift indicating the periods of a day that an agent will be required to work. Once forecasting, staffing, and scheduling are completed, the planning process concludes with either rostering which amounts to assigning each individual agent to a specific work schedule, i.e. deciding who is going to work when, or hiring which amounts to the creation of a plan that indicates the number of agents to hire, fire, and train over some time horizon. Note that while rostering in almost all call centers is done on a seniority basis, in some industries like health care where hospitals find it difficult to find nurses to work unpopular shifts, auctions are employed, i.e. nurses who might be willing to work indicate through a bid how much they would need to be paid to get them to come in and if a nurse's bid is deemed acceptable, then they get to work the shift in question at the wage they require (see Li, 2005).

The rest of this paper is organized as follows. In the next section we detail the underlying tactical planning process after which we describe the strategic planning problem and our solution to it. In the third section of the paper we report the findings of a simulation study. We end the paper with a discussion of the managerial insights that we have obtained by conducting the research described herein.

PROBLEM SETTING AND SOLUTION APPROACH

In the problem setting of this paper we consider only revenue generating inbound call centers where the objective is to maximize expected net profit which we define as revenue less wages and hiring/training expenses. We begin with the tactical planning process. To fix ideas, assume that management is trying to put together a capacity plan for some hour a week ahead, hence the implication is that each shift is an hour in duration. Each such shift is put up for bid, i.e. auctioned. One possibility is to conduct a single auction for each shift immediately before its start. Such timing would require a workforce with tremendous flexibility in the sense that each agent would need to be able to commence working on very short notice, and while such a workforce could be assembled, it would necessarily exclude people who need to make arrangements before actually working.

That the exclusion of one group in an auction setting may be undesirable is illustrated by an example found in Vulcano et al. (2002). In their paper Vulcano et al. discussed the problem of auctioning seats on an airline flight in order to maximize revenue given the existence of two customer classes: business and leisure. They observed that if only a single auction is conducted for the flight, then many potential bidders/buyers will be eliminated as such an auction will be either too early for business travelers who often do not know of their need to travel much in advance, or too late for leisure travelers who need to coordinate their vacation travel with other arrangements, like making resort reservations, scheduling

time off from work, and finding child care. Vulcano et al. demonstrated that in order to maximize revenue in such settings, it is optimal to conduct auctions at multiple points in time.

In our virtual call center world the workforce can be divided into two groups: those agents who need to make arrangements prior to actually working, e.g. find childcare, and those who do not. Given this division and the research findings of Vulcano et al., conducting a pair of auctions for each shift is beneficial. The first auction of the pair is held sufficiently in advance to accommodate the needs of those having to make arrangements, while the other auction is conducted at a point in time much closer to the beginning of the actual shift. At the close of the first auction, management selects some number of winning bidders and in doing so commits to the employment of specific agents during the forthcoming period. Later, at the time of the second auction, management assesses whether there is a need for additional manpower beyond first auction commitments, and if so, management augments the commitments already made by engaging some number of second auction participants. In this setting management faces a pair of decision problems for each shift. We next examine each of these decision problems in detail, beginning with the one corresponding to the first auction.

So what information does management have at the time of a first auction? When bidding closes management knows the number of first auction bids received n_1 , as well as the value of each individual bid, the collection of which we denote as the vector B^{n_1} . Management also has some expectation as to the future call volume, the number of second auction bids that might be received, and the possible value of each second auction bid. In each auction conducted, bids are received only from people hired by management to work as agents. We assume that if an agent at the time of hire is someone who needs to make advance arrangements, then henceforth that agent participates only in first auctions. On the other hand an agent without the arrangements constraint at time of hire is assumed to participate in only second auctions throughout his career. The workforce is therefore divided into two pools of agents: the constrained (pool 1) and the unconstrained (pool 2), where the size of each is denoted by a_c and a_u respectively.

At the time management needs to decide the number of first auction winners, there is some information uncertainty. We now specify precisely the expectations of management regarding this uncertain information. Letting p_u be the probability that an agent in the unconstrained pool submits a bid for a shift, we model the total number of second auction bids received for a shift as a random variable n_2 drawn from a binomial distribution with parameters a_u and p_u . Hence the probability of observing k second auction bids for a shift is given by:

$$Pr\{n_2 = k\} = \binom{a_u}{k} p_u^k (1 - p_u)^{a_u - k}, \quad k = 0, \dots, a_u. \quad (1)$$

To make any statement about the possible values taken on by a collection of n_2 bids received, we must be more definitive about the type of auction conducted. In this paper, each auction is presumed to be of the uniform-price variety. This means that when management chooses j out of n_2 second auction bidders for employment, each of the j winners receives a wage equal to the bid of the $j + 1^{th}$ smallest bid. In a uniform-price action in which each bidder demands only a single unit, as is our case since an agent cannot work a single shift twice, it is a weakly dominant strategy for each auction participant to bid his true valuation of the item being auctioned (see Krishna, 2002). Modeling the valuation of each agent as a uniform random variable on the range $[\underline{\alpha}, \bar{\alpha}]$ with $\bar{\alpha} > \underline{\alpha}$, and defining $B_{j+1}^{n_2}$ as the $j + 1^{th}$ smallest bid of n_2 second auction bidders, it can be shown that:

$$E[B_{j+1}^{n_2}] = (\bar{\alpha} - \underline{\alpha}) \frac{j + 1}{n_2 + 1} + \underline{\alpha}. \quad (2)$$

The expectations of management regarding second auction bidding are completely specified by (1) and (2). With respect to call volume, we make the assumption that the number of calls received each period t has a Poisson distribution with a mean λ_t . We assume λ_t is a random variable. Let $d_t = \lambda_t - E(\lambda_t)$ be the deviation of d_t from its long term mean. Then $E(d_t) = 0$. We assume that the fluctuation of $\{d_t\}$ around zero is described by an autoregressive process. Specifically, letting $-1 < \rho < 1$, we define

$$d_{t+1} = \rho d_t + \sigma \sqrt{1 - \rho^2} \epsilon_t \quad (3)$$

where $\{\epsilon_t\}$ are independent and identically distributed standard normal random variables with d_t independent of ϵ_t . If $\{d_t\}$ is stationary (same mean and variance) then it follows that $E(d_t) = 0$ for all t . Also,

$$Var(d_{t+1}) = \rho^2 Var(d_t) + \sigma^2(1 - \rho^2) \quad (4)$$

so that $Var(d_t) = \sigma^2$, all t . We use this in what follows. To find the correlation structure: multiply (3) by d_t and take expected values: $E(d_{t+1}d_t) = \rho Var(d_t) + \sigma \sqrt{1 - \rho^2} E(d_t \epsilon_t)$. The last term is equal to zero by independence so $E(d_{t+1}d_t) = \rho \sigma^2$. From this it follows that the correlation between d_t and d_{t+1} is ρ . Also we can show the model implies $E(d_{t+k}|d_t) = \rho^k d_t$ which converges to 0 as k increases. This says our forecast for λ_t in the indefinite future is the long-run average no matter what we just observed. If d_{t-1} has been observed then the conditional distribution of d_t is $N(\rho d_{t-1}, \sigma^2(1 - \rho^2))$.

These expectations of call volume and second auction bidding, along with n_1 and B^{n_1} , the expected revenue per call (R), the probability a call is successfully handled (q), and the assumption that a call not immediately answered is a call lost is the information management has available when determining the optimal number of first auction winners. In the sequel we refer to this as decision problem one (DP1). This is a problem that lends itself to formulation as a two-stage stochastic program with recourse. Before presenting such a formulation we first detail the other lowest level periodic decision problem facing management, which we will refer to as decision problem two (DP2).

DP2 arises in conjunction with each second auction. As already indicated, at the time of each second auction management needs to resolve whether any of its participants should be selected to augment the corresponding first auction winners. At the time this decision is made, we assume that management knows the number of second auction bids that will be received, i.e. the number of second auction participants, but not the value of any second auction bids nor the call volume for the shift. With respect to bid values, the expectations of management continue to be given by (2) above, as we assume actual bid values are not revealed to management until after management has announced the number of second auction bids that will be accepted, while in the case of call volume a more accurate forecast is now obtainable.

We now return to DP1 and its formulation as a stochastic program. For a comprehensive treatment of stochastic programming see Birge and Louveaux (1997). We formulate DP1 as the following two-stage stochastic program with recourse:

$$\max f(x) = \sum_{i=0}^{n_1} -i B_{i+1}^{n_1} x_i + Q(x) \quad (5)$$

s.t.

$$\sum_{i=0}^{n_1} x_i = 1 \quad (6)$$

$$x_i = 0 \text{ or } 1 \quad (7)$$

where n_1 is the number of first auction bids, $B_i^{n_1}$ is the i^{th} smallest of the collection of first auction

bids B^{n_1} (define $B_{n_1+1}^{n_1} \equiv B_{n_1}^{n_1}$), each x_i is a first-stage binary decision variable that takes on the value 1 (one) when it is optimal to hire the i lowest bidders of the first auction each at a wage of $B_{i+1}^{n_1}$, and

$$Q(x) = E_{\xi} Q(x, \xi(\omega)) \quad (8)$$

where

$$Q(x, \xi(\omega)) = \max \sum_{j=0}^{n_2(\omega)} \{-j B_{j+1}^{n_2}(\omega) + \sum_{i=0}^{n_1} \lambda(\omega) [1 - \mathcal{L}(i+j, \lambda(\omega)/\mu)] R q x_i\} y_j \quad (9)$$

s.t.

$$\sum_{j=0}^{n_2(\omega)} y_j = 1 \quad (10)$$

$$y_j = 0 \text{ or } 1 \quad (11)$$

The symbol ω denotes a random event or outcome. For each random outcome ω , $n_2(\omega)$ is a realization of the random variable n_2 , the number of second auction bids, $B^{n_2}(\omega)$ is a realization of the expectation of the random vector B^{n_2} , the collection of second auction bids, with $B_j^{n_2}(\omega)$ the expected value of the j^{th} smallest bid (define $B_{n_2+1}^{n_2}(\omega) \equiv B_{n_2}^{n_2}(\omega)$), and $\lambda(\omega)$ is a realization of the random variable λ the expected call volume for a shift. Piecing together the foregoing stochastic elements, we obtain a vector $\xi^T(\omega) = (n_2(\omega), B^{n_2}(\omega)^T, \lambda(\omega))$ with $N = 1 + n_2 + 1$ components. $Q(x, \xi(\omega))$ is thus the value of the second stage for a given realization of the random vector ξ , and the value of the recourse function $Q(x)$ is obtained by taking the expectation of $Q(x, \xi(\omega))$ with respect to the probability distribution of ξ . In the above formulation, for each realization of ξ , each y_j is a second-stage binary decision variable that takes on the value 1 (one) when it is optimal to hire the j lowest bidders of the second auction each at an expected wage of $B_{j+1}^{n_2}(\omega)$. The parameters R and q retain their previous meanings, while the parameter μ indicates the average number of calls per shift that an agent is able to process. Finally, $\mathcal{L}(i+j, \lambda(\omega)/\mu)$ denotes the Erlang loss probability, defined as:

$$\mathcal{L}(i+j, \lambda(\omega)/\mu) = \frac{(\lambda(\omega)/\mu)^{i+j} / (i+j)!}{\sum_{k=0}^{i+j} (\lambda(\omega)/\mu)^k / k!}, \quad (12)$$

which gives the expected fraction of calls that are not handled and hence lost when the number of agents working a shift is $i+j$ and the offered load is $\lambda(\omega)/\mu$.

This stochastic programming representation captures the sequence of events that embody the tactical capacity planning problem faced by management. First-stage decisions x are taken at the close of first auction bidding in the presence of uncertainty about future realizations of ξ . In the second stage, the actual number of second auction participants, one element of ξ , becomes known, while for another element, call volume, an improved estimate becomes available, after which a recourse action y , i.e. indicating the number of second auction bidders that will be selected to work, is taken. Most importantly, first stage (auction) decisions are chosen by taking their future effects into account, with these effects measured by the recourse function $Q(x)$, which computes the value of taking decision x .

Having specified the formulation of DP1 as a two-stage stochastic program, we now describe our method of arriving at its solution. Our approach recognizes that attempting to solve the above stochastic program as a nonlinear program is sure to involve numerical difficulties. Fortunately, for a given first-stage decision, the recourse function $Q(x)$ is computable, hence we can determine the optimal number of first auction winners by following this two step procedure:

Step 1: Let $C_1 = \{0, \dots, n_1\}$ where n_1 is the number of first auction bidders. Also let d_i denote the

first-stage decision where $x_i = 1$ and the remaining first-stage decision variables equal zero. Compute the recourse function $Q(d_i)$ for each $i \in C_1$.

Step 2: Let i^* denote the optimal number of first auction winners. Determine i^* by evaluating:

$$i^* = \operatorname{argmax}_{i \in C_1} -iB_{1,i+1} + Q(d_i) \quad (13)$$

Our approach to solving DP2 is more straightforward. To determine the number of second auction winners, which we denote as j^* , we solve the following deterministic optimization program:

$$j^* = \min \operatorname{argmax}_{j \in C_2} -jB_{j+1}^{n_2} + \lambda[1 - \mathcal{L}(i^* + j, \lambda/\mu)](Rq) \quad (14)$$

where n_2 is the number of second auction participants, $C_2 = \{0, \dots, n_2\}$, $B_j^{n_2}$ is the expected value of the j^{th} smallest bid of the second auction given n_2 bidders (define $B_{n_2+1}^{n_2} \equiv B_{n_2}^{n_2}$), λ is the expected call volume, $\mathcal{L}(i^* + j, \lambda/\mu)$ the Erlang loss probability, R the expected revenue per call, q the probability a call is successfully handled, μ the average number of calls per shift an agent is able to process, and i^* the already determined number of first auction winners.

The above completes our treatment of the tactical planning process. We now turn our attention to the strategic planning problem. The strategic planning problem amounts to determining the number of people to hire and train up to work as agents in order to ensure that lower level staffing requirements can be met. This is a problem faced by every call center management and is therefore in some sense classical, but the problem as considered here is different because of its auction underpinnings. In our setting the problem amounts to sizing the pools of constrained (pool 1) and unconstrained (pool 2) agents. In what follows we will refer to the sum of agents in pools 1 and 2 as the hire up-to level. Management must decide the hire up-to level and the workforce split of agents across the two pools. In the following we will limit our attention to five workforce splits: all agents hired into pool 1 (denoted 1/0), all agents hired into pool 2 (denoted 0/1), agents hired into both pools equally (denoted 50/50), 80% of all agents hired into pool 1 (denoted 80/20), and 20% of all agents hired into pool 1 (denoted 20/80). A hire up-to policy is therefore given by a workforce split and a hire up-to level. For a given workforce split the fundamental problem is to determine the hire up-to level that maximizes expected net profit (revenue less wages and hiring/training expenses). Operationally speaking, adoption of a hire up-to policy requires management to periodically compare the total number of trained agents (a number always being drawn down through attrition) and the total number of agents in training with the target hire up-to level, and if the number trained plus in training is less than the target, then the difference is hired to begin agent training. We discuss next how we model agent attrition.

We model agent attrition as a function of bidding success (or lack thereof) among other things (see Whitt (2006) for a similar model). We treat the career of each agent as a nonhomogeneous discrete-time Markov chain $\{X_n\}$ with states $\{0, 1\}$, where n denotes the length-of-service (LOS), i.e. the number of days since hire, and the states indicate respectively whether an agent is an active or inactive bidder, the latter synonymous with an agent that has quit the job. We define the transition probabilities of the chain as follows: $p_{00}(s, n) = g(s) - h(n)$, $p_{01}(s, n) = 1 - g(s) + h(n)$, $p_{10}(s, n) = 0$, and $p_{11}(s, n) = 1$, where $0 \leq s \leq 1$ measures the cumulative bidding success of an agent since hire, $g(s)$ is a function of that bidding success mapping into a subrange of $(0, 1]$, and $h(n) = \min(n/\Gamma, g(s))$ is a function of LOS through the value of n , with Γ a suitably defined positive constant. The value of $p_{00}(s, n)$ gives the probability that an active bidder in period n will remain one in the following period, where that probability increases the more successful an agent is as a bidder as measured by the term $g(s)$, but where that probability decreases the longer an agent is on the job as measured by the term $h(n)$; the value of $p_{01}(s, n)$ is simply $1 - p_{00}(s, n)$. Meanwhile $p_{11}(s, n)$, as defined, indi-

Scenario	Call Success (q)	Call Volume Standard Deviation (σ)
1	0.80	25
2	0.80	75
3	0.80	125
4	0.60	25
5	0.60	75
6	0.60	125
7	0.40	25
8	0.40	75
9	0.40	125

Table 1: Scenarios of simulation study

cates that state 1 is absorbing, that is once an agent quits, he never participates in an auction again; consistent with that reality the value of $p_{10}(s, n)$ is $1 - p_{11}(s, n)$.

Incorporating the above model of agent attrition into the strategic planning problem greatly increases the strategic planning problem’s complexity and necessitates the use of simulation to determine the optimal hire up-to level for a given workforce split. We conjecture that the net profit function is concave, hence our approach to determining the optimal hire up-to level for a given workforce split is to simulate the operation of a virtual call center for a hire up-to level and obtain a measure of average hourly net profit. We then perform additional simulations for successively larger hire up-to levels until average hourly net profit falls, which given our conjecture that the net profit function is concave, signals that larger hire up-to levels are not optimal.

A SIMULATION STUDY

In this section we report the results of a study we conducted to learn how the strategic planning process may impact the profitability of a virtual call center that uses a tactical planning approach employing auctions. In our study we adopted an hourly perspective and we set the long-run average demand to 500 calls per hour. We also fixed ρ , the correlation, choosing the value of 0.8. Such a value is consistent with the observation in Brown et al. (2002) that call center arrival rates tend to “vary smoothly and not too rapidly throughout the day”. As to σ , the standard deviation of demand, we used values of 25, 75, and 125, while we assumed that R , the expected revenue per call, equals \$20. We adopted the convention that auctions for members of pool 1 were concluded one week in advance of the period to be worked, while for members of pool 2 the conclusion of an auction coincided with the beginning of the period up for bid. We fixed p_c (p_u), the probability that an agent in pool 1 (2) submits a bid in any auction, at the value of 0.167 (0.167). With respect to the valuation of an agent, as indicated previously we assume it is uniformly distributed on a range $[\underline{\alpha}, \bar{\alpha}]$; in our study we set $\underline{\alpha} = \$10$, and $\bar{\alpha} = \$20$. The parameter μ , the service rate of each agent, was fixed at 6, while for q , the probability a call is successfully handled, we used values of 0.40, 0.60, and 0.80. Lastly, we set agent hiring and training expenses at \$1,000, and let an agent immediately begin taking calls upon hire, the implication being that a new employee is drawn from a pool of qualified people. Our decision to vary q and σ while holding all other parameters constant led to an investigation of nine different scenarios (see Table 1). For each of these scenarios we simulated the operation of a virtual call center for the five workforce splits mentioned in the previous section and measured how average hourly net profit varies with hire up-to level. We next discuss in detail three of the scenarios.

We begin with Scenario 5 because it represents what might be called the middling scenario of our study (see Figure 1). Figure 1 shows that for all the workforce splits considered the optimal hire up-to level lies in the vicinity of 600-700. It can be seen that the net profit function is rather flat in the neighborhood of the profit maximizing hire up-to level. This is good news because it means that when it comes to establishing a hire up-to level an operator only has to get into the vicinity of what is optimal in order to achieve a high level of financial performance since there is no sharp drop off in profitability as one moves away from the optimal hire up-to level.

We next consider Scenario 1 as it represents one extreme of our study (see Figure 2). Figure 2 shows that the financial performance achievable under the different workforce splits is indistinguishable. Under such circumstances it does not matter much which workforce split is adopted, only that about 600-700 people are hired and trained up to work as agents. That is if management decides to go with a hire up-to level of 600, then the same level of profitability will be achieved regardless of whether all 600 agents are members of pool 1, or all 600 agents are members of pool 2, or the agents are split across the pools in some fashion, e.g. with 480 in pool 1 and 120 in pool 2. Note that under each workforce split the net profit curves of this scenario are flatter than those of Scenario 5, thus as uncertainties diminish we see that the task of selecting a hire up-to level that will deliver a high level of financial performance becomes even easier for an operator. In this case that translates into a level of profitability in Scenario 1 that is about 57% higher than that found in Scenario 5.

We end this section with consideration of Scenario 9 which corresponds to the other extreme of our study (see Figure 3). We see from Figure 3 that the curve corresponding to a 0/1 workforce split largely defines an upper envelope, the exception being some of the largest hire up-to levels investigated, hence a 0/1 workforce split is always preferred to all others in the most meaningful range of hire up-to levels. At the same time, the curve corresponding to a 1/0 workforce split now defines a lower envelope, hence the adoption of any workforce split other than 1/0 will lead to greater profitability. At the hire up-to level of 600, we find that the workforce split (0/1) holds a slightly greater than 5% advantage over the workforce split (1/0), while the workforce splits 20/80 and 50/50 are within 1% and 2% respectively of what can be achieved under 0/1. We also see that the net profit curves of this scenario are more peaked than those of Scenarios 1 and 5 and the corresponding decline in profitability that comes with moving from the observed peak at 600 to either of the adjacent hire up-to levels investigated is approximately 3% in the case of the 0/1 workforce split.

CONCLUSIONS

In this paper we investigated the impact the strategic planning process may have on the profitability of a virtual call center that uses a tactical planning approach employing auctions. To gain the understanding desired we undertook a simulation study. We found that achieving a high level of profitability is most challenging when call volume variability is high and worker success in handling a call is low.

REFERENCES

References are available upon request from Matthew Keblis.

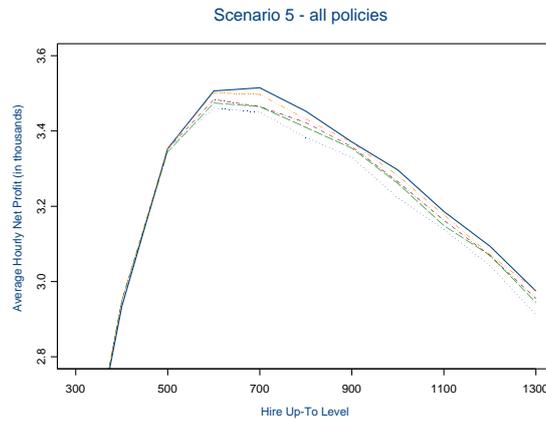


Figure 1: Average hourly net profit for each hire up-to level for the following workforce splits: all agents in pool 2 (solid line), all agents in pool 1 (dotted line), 50% of agents in pool 1 (dotted-dashed line), 80% of agents in pool 1 (dashed line), 20% of agents in pool 1 (three-dot-dashed line).

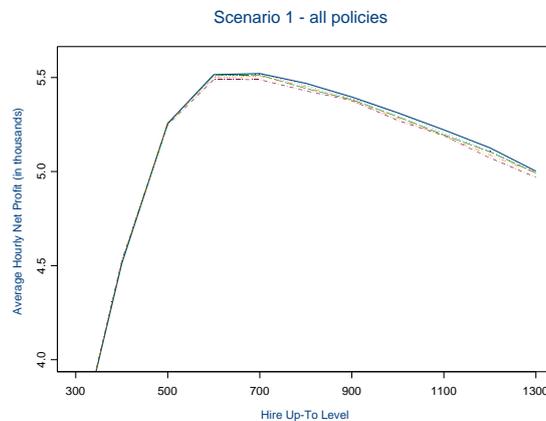


Figure 2: Average hourly net profit for each hire up-to level for the following workforce splits: all agents in pool 2 (solid line), all agents in pool 1 (dotted line), 50% of agents in pool 1 (dotted-dashed line), 80% of agents in pool 1 (dashed line), 20% of agents in pool 1 (three-dot-dashed line).

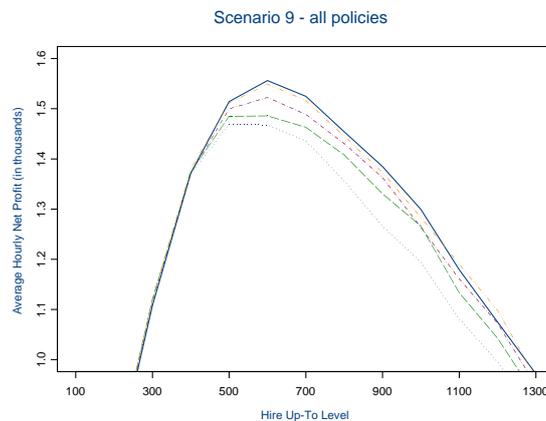


Figure 3: Average hourly net profit for each hire up-to level for the following workforce splits: all agents in pool 2 (solid line), all agents in pool 1 (dotted line), 50% of agents in pool 1 (dotted-dashed line), 80% of agents in pool 1 (dashed line), 20% of agents in pool 1 (three-dot-dashed line).