

Diagnostic Accuracy Under Congestion

Saed Alizamir

Fuqua School of Business, Duke University, Durham, North Carolina 27708, saed.alizamir@duke.edu

Francis de Véricourt

INSEAD, 77305 Fontainebleau, France, francis.devericourt@insead.edu

Peng Sun

Fuqua School of Business, Duke University, Durham, North Carolina 27708, peng.sun@duke.edu

In diagnostic services, agents typically need to weigh the benefit of running an additional test and improving the accuracy of diagnosis against the cost of delaying the provision of services to others. Our paper analyzes how to dynamically manage this accuracy/congestion trade-off. To that end, we study an elementary congested system facing an arriving stream of customers. The diagnostic process consists of a search problem in which the service provider conducts a sequence of imperfect tests to determine the customer's type. We find that the agent should continue to perform the diagnosis as long as her current belief that the customer is of a given type falls into an interval that depends on the congestion level as well as the number of performed tests thus far. This search interval should shrink as congestion intensifies and as the number of performed tests increases if additional conditions hold. Our study reveals that, contrary to diagnostic services without congestion, the base rate (i.e., the prior probability of the customer type) has an effect on the agent's search strategy. In particular, the optimal search interval shrinks when customer types are more ambiguous a priori, i.e., as the base rate approaches the value at which the agent is indifferent between types. Finally, because of congestion effects, the agent should sometimes diagnose the customer as being of a given type, even if test results indicate otherwise. All these insights disappear in the absence of congestion.

Key words: service operations; queueing theory; dynamic programming; decision making; information search; Bayes' rule

History: Received December 22, 2009; accepted January 11, 2012, by Martin Lariviere, operations management. Published online in *Articles in Advance* September 21, 2012.

1. Introduction

Diagnostic services focus on determining customer needs but do not themselves perform any subsequent treatments that may be indicated by the diagnosis. Accumulating information and running additional tests on a customer is likely to improve the diagnosis, the accuracy of which obviously affects the value of the service offered. Accumulating and processing information, however, takes time and therefore increases congestion in the system. The service provider thus needs to weigh the benefit of running an additional test against the cost of delaying the provision of services to others.

Triage nursing systems provide a typical example of such trade-offs. The nurse elicits different pieces of information to assess the severity of the patient's symptoms (Gertz and Bucknall 2001). On the other hand, long triage processes can result in adverse patient outcomes (Travers 1999). Another example of an accuracy/congestion trade-off occurs at MTU Aero Engines, which is Germany's leading provider of engine maintenance. One key decision the company needs to make is whether to keep or replace expensive parts of an engine. The diagnosis is performed by a dedicated team of workers, who may have many

parts awaiting inspection. This can yield high costs of delay in an industry that is subject to intensive time-based competition. A similar task confronts those who carry out remanufacturing processes, which typically require determining whether returned parts are obsolete or not (Guide and Wassenhove 2001). Finally, the need to accurately determine a customer type under congestion also frequently occurs in some nondiagnostic systems, such as support centers and help desks (de Véricourt and Zhou 2005).

In this paper, we attempt to gain insight into this problem of dynamically balancing accuracy against delays in the process of rendering a diagnosis. To that end, we study an elementary diagnostic system facing a random stream of customers. The service consists in identifying each customer to be one of two types, $\bar{\tau}$ or $\underline{\tau}$. Accuracy is defined as the probability that the customer type is correctly diagnosed. We represent the diagnostic process as a sequential test problem in which the agent performs imperfect tests one by one. The test result may be either "positive" or "negative." We assume that all customers start with a base rate probability p_0 of being type $\bar{\tau}$, and each test result updates the subjective probability of customer type according to Bayes' rule.

Running a test takes time; service demands may therefore accumulate. More tests generally produce more accurate diagnosis, but also induce longer waiting times. To limit delays in the system, the agent may stop the diagnostic process at any time and move to the next customer. We formulate this problem of balancing diagnostic accuracy against delays as a partially observed Markov decision process (POMDP) and characterize the structure of the optimal decision rule that maximizes the long-run average value to the service provider, which includes rewards for correctly identifying customer type as well as costs associated with misidentifications and delays. In particular, we show that the service provider should perform additional tests as long as her current probability of the customer being of type $\bar{\tau}$ falls into an interval that depends on the congestion level as well as the number of performed tests thus far. We show that the optimal search interval shrinks as congestion intensifies. Further, we show that when more informative tests are run first, the optimal search interval also shrinks with the number of performed tests.

Our analysis reveals general insights into the impact of congestion on managing diagnostic systems. First, the base rate of a customer type has an effect on the agent's search strategy. In particular, the optimal search interval shrinks when customer types are more ambiguous a priori, i.e., as p_0 approaches the value at which the agent is indifferent between types. By contrast, the base rate can be ignored altogether when designing diagnostic services with no congestion, i.e., for dynamic search problems without accumulation of tasks.

Further, we demonstrate that the agent sometimes should stop the diagnostic process and identify the customer type against the result of a test. This means, for example, that a customer may be identified as type $\bar{\tau}$ even if test results are all positive. Again, effects such as these never occur for diagnostic services in the absence of congestion.

Finally, we apply these findings to the special case of diagnosis where a negative test perfectly reveals type $\bar{\tau}$ and terminates the process. This corresponds to settings where tests are treated as one sided and false negatives assumed to be negligible. An example is the protocol designed by Breiman et al. (1984), which classifies heart attack patients according to risks;¹ see HP Renew Program (2009) for examples

¹ The protocol tests in the following order: whether the blood pressure is above a prespecified threshold, whether the patient is older than 62.5 years, or whether sinus tachycardia is present. The diagnosis process stops and classifies the patient as low risk as soon as a test result is negative. The process moves to next test otherwise (see Goldstein and Gigerenzer 1999). In our framework, type $\bar{\tau}$ customers correspond to the low risk patients and the diagnosis process falls into the one-sided test case.

in a remanufacturing setting.² More generally, stress testing of products is commonly used in many industries (such as maintenance, production and information technology industries) and typically consists of sequentially checking whether a part conforms to different tolerance levels. A part fails the stress test as soon as it fails to satisfy one of these specifications. In our framework, type $\bar{\tau}$ customers correspond to non-conforming parts, and the tests with one-sided errors correspond to checking the different tolerance levels.

When the diagnostic process is one sided, managing the system does not require updating the service provider's belief probability with Bayes' rule, and our results can be expressed in terms of the number of performed tests and the number of customers in the system. This also yields a representation, which is akin to formulations found in the queueing literature. In this setting, the optimal decision rule may be characterized by a threshold on the number of customers allowed in the system. When tests are performed in the order of their accuracy, the optimal policy can be represented as two monotone thresholds in the number of tests performed.

Research in the field of operations management has addressed the problem of balancing congestion against the value offered to the customer. Hopp et al. (2007) proposed a queueing model with Poisson demand and deterministic service time, in which the value offered to the customer is an increasing concave function of the service time. The objective is to balance the congestion related costs against the generated value. The authors showed that, under the optimal policy, the service time should decrease with the level of congestion, or, equivalently, that the maximum number of customers allowed in the system should decrease with service time. Bouns (2003) studied the same model, except that the service time has an Erlang distribution. The agent can then dynamically adjust the number of stages of the distribution to maximize profit. The optimal policy possesses a similar structure in that the maximum number of customers allowed in the system should decrease with the current stage of the service time. A related problem is the speed-congestion trade-off studied by George and Harrison (2001), where the agent continuously adjust the service rate to minimize congestion-related and service rate costs³ (for earlier work on this problem, see also Crabill 1972, Crabill et al. 1977, Stidham and Weber 1993). The optimal policy is shown to increase the maximum

² In particular, an incoming product needs to be replaced if it fails one of four (binary) cosmetic tests.

³ Service rate costs can be interpreted in our context as customer disutility; the faster the service, the lower the generated value. With this interpretation, minimizing service rate cost is equivalent to maximizing value.

number of customers allowed in the system with the service rate.

Our diagnostic system retains the basic elements of most previous models, with the main departure in its representation of the service process as a sequential testing problem. This change, however, yields different results and insights. First, diagnostic services require tracking the belief probability in addition to the current congestion level and service time. For example, we are able to study the impact of the customer type base rate, which has no equivalent counterpart in the previous models. Nonetheless, managing the one-sided case does not require updating the belief probabilities with Bayes' rule and is more akin to the value/congestion trade-off models. However, whereas a decreasing threshold on congestion levels describes optimal rules in the previous queuing literature, our optimal structure requires two thresholds on the number of performed tests, or, equivalently, one unimodal threshold on the level of congestion (under some additional conditions on test accuracy). In essence, our optimal rule for diagnostic services retains a more general form of monotonicity property, with a search interval that shrinks as congestion intensifies.

Recent research on customer behavior in congested systems has also focused on issues related to the value/congestion trade-off. In particular, Wang et al. (2010) studied patient behaviors in a call center of triage nurses, where the service corresponds to performing diagnoses (see also Anand et al. 2011 for the study of speed-quality trade-offs with strategic customers). The service corresponds to a continuous search problem that does not dynamically depend on congestion levels, as it does in our study. Thus, their model cannot account for our findings. On the other hand, their model allows exploring the impact of different system parameters on demand, an issue we do not address.

Finally, when there is ample service capacity such that arriving customers always find an available agent (i.e., when the number of servers is infinite), the corresponding diagnostic system reduces to a single diagnostic task problem with no congestion. The optimal decision rules for these systems without congestion are well established. Specifically, the one-sided search problem is treated in Bertsekas (2007a) and the symmetrical two-sided test problem was first solved by Edwards (1965). More generally, these two systems are special cases of sequential hypothesis testing problems (see, for instance, Wald 1947, DeGroot 1970), which have not been studied under congestion, to the best of our knowledge. It is worth noting that all our insights that include the impact of the base rate on the optimal rule, or a type identification that goes against all test results, disappear when there is no congestion.

We present the model in the next section. The optimal decision rules are characterized in §3. In §4 we study the special case of one-sided tests. We present the effect of the base rate in §5. Section 6 concludes the paper.

2. Models of Diagnostic Services

Consider a service provider serving customers arriving according to a Poisson process with rate λ . The server performs tests one by one to identify a customer's type, which can be either $\bar{\tau}$ or $\underline{\tau}$. Each test takes an exponentially distributed time with rate μ . We denote $\rho \equiv \lambda/\mu$ to be the test utilization rate, that is, the average number of arriving customers while a test is performed. The number of available tests may be infinite. The system is preemptive in the sense that a test can be stopped at any time.

A test result is either positive (signaling type $\bar{\tau}$) or negative (signaling type $\underline{\tau}$). Tests, however, are not perfect and may produce false outcomes. In particular, conditional upon the customer being of type $\bar{\tau}$, the $(k+1)$ st test returns a positive result with probability α_k . Similarly, the $(k+1)$ st test returns a negative result with probability β_k given a type $\underline{\tau}$ customer. After receiving a test result, the agent updates her belief probability on the customer's type.

We denote p_k to represent the probability of a customer being type $\bar{\tau}$ after completing the first k tests. Probability p_0 is the agent's initial belief about a customer being type $\bar{\tau}$, or, the base rate of type $\bar{\tau}$ customers in the population. Probability p_k evolves according to the Bayes' rule. If the next test result is positive, the posterior probability becomes

$$p_{k+1} = \pi_k^+(p_k) \equiv \frac{\alpha_k p_k}{\alpha_k p_k + (1 - \beta_k)(1 - p_k)}. \quad (1)$$

If the result is negative, on the other hand, the posterior probability becomes

$$p_{k+1} = \pi_k^-(p_k) \equiv \frac{(1 - \alpha_k)p_k}{(1 - \alpha_k)p_k + \beta_k(1 - p_k)}. \quad (2)$$

At any time, the agent needs to decide whether to run a new test or terminate the diagnosis process and proceed to the next customer. Specifically, the agent can take one of the following three actions: stop testing and identify the customer as type $\bar{\tau}$, stop testing and identify the customer as type $\underline{\tau}$, or continue testing. Correctly identifying a type $\bar{\tau}$ customer brings value \bar{v} to the system. This reward may also include the benefit of subsequent services that $\bar{\tau}$ customers may receive. Examples include the value of repairing a part in maintenance services or treating a patient, etc. On the other hand, missing and releasing a type $\bar{\tau}$ customer as $\underline{\tau}$ incurs a misidentification cost \bar{c} . This corresponds to the disutility of not providing required healthcare or the expected cost of potential failures

when the part is not repaired. Similarly, we denote by \underline{v} and \underline{c} the reward and misidentification cost associated with a type $\underline{\tau}$ customer, respectively. If the agent identifies the customer as type $\bar{\tau}$ given her current probability p , the expected reward is equal to $\bar{r}(p) \equiv p\bar{v} - (1-p)\underline{c}$. The corresponding expected reward for identifying the customer as type $\underline{\tau}$ is $\underline{r}(p) \equiv (1-p)\underline{v} - p\bar{c}$.

We can then define θ as the unique value of p which satisfies $\bar{r}(p) = \underline{r}(p)$. That is,

$$\theta = \frac{\underline{v} + \underline{c}}{\bar{v} + \bar{c} + \underline{v} + \underline{c}}.$$

Critical fraction θ highlights the relative value of correctly identifying type $\bar{\tau}$ customers. In particular, θ decreases with $(\bar{v} + \bar{c})/(\underline{v} + \underline{c})$. From the definition of θ , we see that if $p \geq \theta$ (respectively, $p < \theta$), then $\bar{r}(p) \geq \underline{r}(p)$ (respectively, $\bar{r}(p) < \underline{r}(p)$) and the service provider would diagnose the customer as type $\bar{\tau}$ (respectively, $\underline{\tau}$), were she to stop the search process. We also consider a waiting cost $c_w(x)$ per unit of time that is incurred when x customers are present in the system. We impose no restrictions on $c_w(x)$ other than it being strictly increasing in $x \geq 0$ with $c_w(0) = 0$ and unbounded (i.e., $\lim_{x \rightarrow \infty} c_w(x) = +\infty$).

A control policy determines the agent's action at any point in time. The performance of a policy is then measured as the long-run average profit. For control policy u , we define the corresponding long-run average profit g^u as

$$g^u = \liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\bar{v} \bar{N}^u(T) + \underline{v} \underline{N}^u(T) - \bar{c} \bar{M}^u(T) - \underline{c} \underline{M}^u(T) - \int_0^T c_w(X^u(t)) dt \right], \quad (3)$$

where $\bar{N}^u(t)$ (respectively, $\underline{N}^u(t)$) is the random cumulative number of correctly identified type $\bar{\tau}$ (respectively, $\underline{\tau}$) customers up to time t ; similarly, $\bar{M}^u(t)$ (respectively, $\underline{M}^u(t)$) is the random cumulative number of misidentified type $\bar{\tau}$ (respectively, $\underline{\tau}$) customers up to time t ; and $X^u(t)$ is a random process representing the number of customers in the system at time t . Later we will show that the optimal policy that maximizes the long-run average profit is stationary, and we define g^* to be the optimal long-run average profit.

Equation (3) captures an accuracy/congestion trade-off. Indeed, given control policy u under which the system is stable, the expected total number of type $\bar{\tau}$ (respectively, $\underline{\tau}$) customers going through the service is $p_0 \lambda t$ (respectively, $(1-p_0) \lambda t$) as t goes to infinity. Denote accuracy $\bar{\delta}^u$ as the probability that the service correctly identifies a type $\bar{\tau}$ customer. Similarly, accuracy $\underline{\delta}^u$ is the corresponding conditional probability for type $\underline{\tau}$. It follows that

$$\bar{\delta}^u = \liminf_{T \rightarrow \infty} \frac{\mathbb{E}[\bar{N}^u(T)]}{p_0 \lambda T} \quad \text{and} \quad \underline{\delta}^u = \liminf_{T \rightarrow \infty} \frac{\mathbb{E}[\underline{N}^u(T)]}{(1-p_0) \lambda T}.$$

The long-run average profit is then equal to

$$g^u = \bar{a} \bar{\delta}^u + \underline{a} \underline{\delta}^u - \liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T c_w(X^u(t)) dt \right] + b, \quad (4)$$

where $\bar{a} \equiv \lambda p_0(\bar{v} + \bar{c})$, $\underline{a} \equiv \lambda(1-p_0)(\underline{v} + \underline{c})$ and $b \equiv \lambda p_0(\underline{c} - \bar{c}) - \lambda \underline{c}$, which holds from the conservation of flow of $\bar{\tau}$ and $\underline{\tau}$ customers. In particular, the impact of policy u on the system performance changes with \bar{v} , \underline{v} , \bar{c} , and \underline{c} only through the total rewards $\bar{v} + \bar{c}$ and $\underline{v} + \underline{c}$.

3. Optimal Rules for Diagnostic Services with Congestion

The problem of finding the optimal decision rule corresponds to a POMDP. The state space is represented as (x, k, p) , in which x is the number of customers in the system, k is the number of completed tests, and p is the probability of the customer in service being type $\bar{\tau}$. In fact, probability p is a sufficient statistic for the test results obtained thus far. Although in theory decisions can be made at any time, the exponential interarrival and test time assumption allows us to only consider decisions made when an arrival occurs or a test is completed without loss of optimality. The optimality equation is, therefore,

$$\begin{aligned} g + J(x, k, p) &= \max \left\{ -c_w(x) + \lambda J(x+1, k, p) \right. \\ &\quad + \mu(\alpha_k p + (1-\beta_k)(1-p)) J(x, k+1, \pi_k^+(p)) \\ &\quad + \mu((1-\alpha_k)p + \beta_k(1-p)) J(x, k+1, \pi_k^-(p)), \\ &\quad \left. g + \bar{r}(p) + J(x-1, 0, p_0), \right. \\ &\quad \left. g + \underline{r}(p) + J(x-1, 0, p_0) \right\}, \\ &\quad \text{for } x \geq 1 \text{ and } k \geq 0 \text{ and } 0 \leq p \leq 1, \quad (5) \\ g + J(0, 0, p_0) &= \lambda J(1, 0, p_0) + \mu J(0, 0, p_0), \end{aligned}$$

where $\bar{r}(p)$ and $\underline{r}(p)$ are the expected profits of assigning type $\bar{\tau}$ and $\underline{\tau}$, respectively, as defined in §2. Recall from the definition of θ that if $p \geq \theta$, then $\bar{r}(p) \geq \underline{r}(p)$. Therefore, according to optimality equation (5), the optimal diagnosis depends on whether p is larger or smaller than θ .

The following theorem, which presents our first theoretical result, describes the optimal policy structure as two thresholds in p .

THEOREM 1. *For diagnostic systems with congestion, the optimal rule can be characterized by two thresholds $\underline{p}(x, k)$ and $\bar{p}(x, k)$ on probability p for any given x and k , where $\underline{p}(x, k) \leq \theta \leq \bar{p}(x, k)$. Performing an additional test is optimal when probability p satisfies $\underline{p}(x, k) < p < \bar{p}(x, k)$. Otherwise, it is optimal to stop testing and identify the customer as type $\underline{\tau}$ when $p \leq \underline{p}(x, k)$, or as type $\bar{\tau}$*

when $p \geq \bar{p}(x, k)$. Furthermore, threshold $\underline{p}(x, k)$ is nondecreasing in x , and $\bar{p}(x, k)$ is nonincreasing in x ; and there exists \bar{x} such that $\underline{p}(x, k) = \bar{p}(x, k) = \theta$ for all $x \geq \bar{x}$.


The main idea of the proof is based on showing that the value function $J(x, k, p)$ is convex in p through value iteration of the corresponding discounted case⁴ along with subtle sample path arguments. The complete proof is presented in Appendix A.

Theorem 1 states that the service provider conducts tests as long as her probability about the customer type belongs to an (x, k) -dependent interval, the length of which decreases as congestion (i.e., the number of customers in the system, x) intensifies. Figure 1(a) illustrates the result for a diagnostic service with five tests. The figure depicts thresholds $\underline{p}(x, 3)$ and $\bar{p}(x, 3)$ for the optimal policy when $k = 3$, that is after three tests have been performed. The dots on the grids represent states of the system in which performing the next test is optimal. The upward (respectively, downward) triangles correspond to the states in which stopping the search and identifying the customer as type $\bar{\tau}$ (respectively, $\underline{\tau}$) is optimal.

The intuition behind Theorem 1 is as follows. The expected profit of stopping the diagnosis process (i.e., $\max\{\underline{r}(p), \bar{r}(p)\}$), is the lowest when $p = \theta$. In this case, the service provider is actually indifferent between identifying the customer as type $\underline{\tau}$ and $\bar{\tau}$. Therefore, the value of information from an additional test is also the highest when $p = \theta$. As a result, the service provider is willing to bear higher congestion costs for additional tests when probability p is close to θ . On the other hand, when p is away from θ , the value of running additional tests is low and the agent aborts the search at lower levels of congestion. Thus, the optimal policy takes the form of an interval around θ , which shrinks as the queue length increases. Note also that when p_0 is sufficiently far away from θ , or the waiting cost high enough, directly identifying customers without performing any test, i.e., a degenerate policy, is optimal. This happens when $p_0 \leq \underline{p}(1, 0)$ or $p_0 \geq \bar{p}(1, 0)$. In this case, no diagnostic service is required and every customer is identified as type $\bar{\tau}$ if $p_0 \geq \theta$ or $\underline{\tau}$ otherwise.

One immediate consequence of Theorem 1 is that the service provider may stop the search and identify the customer against a test result. This means, for instance, that the service provider may diagnose the customer as type $\bar{\tau}$, even though the latest test indicated $\underline{\tau}$. To see this, consider Figures 1(a) and 1(b), where Figure 1(b) depicts the optimal policy after an additional test has been performed (i.e., for

Figure 1 Optimal Policy for a Diagnostic System with Congestion; $p_0 = 0.52$, $\alpha_1 = 0.68$, $\alpha_2 = 0.55$, $\alpha_3 = 0.6$, $\alpha_4 = 0.65$, $\alpha_5 = 0.75$, $\beta_1 = 0.71$, $\beta_2 = 0.65$, $\beta_3 = 0.52$, $\beta_4 = 0.6$, $\beta_5 = 0.8$, $\rho = 0.4$, $\bar{v} = 180$, $\underline{v} = 135$, $\bar{c} = \underline{c} = 0$, and $c_w(x) = x$



Notes. The model parameters are chosen to demonstrate that thresholds may be nonmonotone in k in general. Further discussion on this is presented at the end of §3.

$k = 4$). In state $(x = 1, k = 3, p = 0.75)$, Figure 1(a) indicates that the agent should run an additional test. Assume then that this additional test result is negative so that the value of probability p decreases from 0.75 to $p = 0.64$. According to Figure 1(b), performing

⁴ Such a claim is far from obvious from the Bellman's equation (5). Furthermore, because of the maximization operator, the value function is not differentiable. Therefore, the proof requires a careful study of the subgradient.

the next test is still optimal in the new system state ($x = 1$, $k = 4$, $p = 0.64$). If, however, two more customers arrive while the next test is still running, the level of congestion reaches $x = 3$ and stopping the search to identify the customer as type $\bar{\tau}$ becomes optimal. This is despite the negative result of the last performed test, which indicated type τ . In other words, the service provider has incurred a cost for performing the previous test but then ignored its result and stopped the search.

In fact, an agent who makes a diagnosis in the presence of congestion can end up identifying a customer as type $\bar{\tau}$, even though all performed test results have indicated type τ . This is because congestion costs might become too high because of the accumulation of new tasks during the diagnostic process. Hence, this observation highlights a difference between diagnostic services with and without congestion. Without congestion, the agent never makes a diagnosis against the previous test result. (See Appendix C, Proposition 4.)

In short, Theorem 1 provides monotonicity properties of the optimal policy in the number of customers. On the other hand, it does not claim anything about the effect of the number of performed tests. In the special case of systems with an infinite number of identical tests such that $\alpha_k = \alpha$ and $\beta_k = \beta$ for all k , the optimal thresholds defined in Theorem 1 do not depend on k (see Equations (1), (2), and (5)). This corresponds, for instance, to situations where the agent can rerun the same test many times independently. Beside these specific systems, however, the number of tests has an effect on the search interval. In fact, thresholds $\bar{p}(x, k)$ and $p(x, k)$ may not be monotone in the number of tests, k , in general. Nonetheless, consider the following sequence of tests:

DEFINITION 1. Let probabilities (α_1, β_1) and (α_2, β_2) characterize tests 1 and 2, respectively. We say that test 1 is *more informative* than test 2 if there exist ξ_1 and ξ_2 , with $0 \leq \xi_1, \xi_2 \leq 1$, such that $\alpha_2 = \xi_1\alpha_1 + \xi_2(1 - \alpha_1)$ and $1 - \beta_2 = \xi_2\beta_1 + \xi_1(1 - \beta_1)$. A sequence of K tests is then *well ordered* if test k is more informative than test $(k + 1)$ for all $k < K$.

When tests are well ordered according to Definition 1, the next result states that the thresholds are also monotone in k .

PROPOSITION 1. *When tests are well ordered, optimal thresholds $\bar{p}(x, k)$ and $p(x, k)$ are nonincreasing and nondecreasing in k , respectively.*

Intuitively, the conditions of Definition 1 state that test 2 outcomes are noisy signals of test 1 outcomes. In particular, conditioning on a positive test 1 result, test 2 gives a positive signal with probability ξ_1 ; similarly, given a negative test 1 result, test 2 gives

a negative signal with probability $(1 - \xi_2)$. These conditions are similar to the conditions in Blackwell's theorem (Blackwell 1953), which establishes a connection between noisy information structures and stochastic dominance, albeit in a setting very different from ours.

When the the service provider can choose the order in which tests are performed, a well-ordered sequence of tests always maximizes the system profit (the proof is omitted). However, Definition 1 describes a partial order and well-ordered sequences may not always exist (this is, for instance, the case for $\alpha_1 = 0.7$, $\beta_1 = 0.8$, $\alpha_2 = 0.75$ and $\beta_2 = 0.75$). The agent may also need to perform less informative tests first because of factors that our model does not directly capture.⁵ In any case, when tests are not well ordered, the search interval can actually expand in k . This means that for fixed values of x and p , the agent stops the process after k but continues the search after $k + 1$ tests. Figure 1 depicts such a case. According to Figure 1(a), the agent stops the process when $k = 3$, for $x = 9$ and $p = 0.37$. However, according to Figure 1(b) the agent continues the search when $k = 4$, for $x = 9$ and $p = 0.37$. This is because test 4 is less informative than tests 3 and 5 in our example. It is not worth performing an additional less informative test after test 3. However, if test 4 has been performed, continuing the search with test 5 becomes valuable again.

4. The One-Sided Case

In many settings, diagnostic processes are treated as one sided in the sense that false negatives are negligible. In our framework, this means that $\alpha_k = 1$ for all k and the process stops as soon as a test is negative, in which case the customer is identified as type τ without possible error. Managing one-sided systems does not require tracking belief probability p anymore because this information is fully captured by the number of tests performed (and therefore positive test results obtained) thus far. The number of performed tests is then directly related to the time spent serving the customer. Formulating the accuracy/congestion trade-off with state (x, k) , therefore, allows contrasting our findings with the value/congestion trade-off studied in the queueing literature.

More specifically, for one-sided systems, test number k uniquely determines belief probability p_k , which is defined as the probability of the customer being type $\bar{\tau}$ after k positive test results. The state space thus reduces to (x, k) and Bayes' rules (1) and (2) simplify to

$$p_{k+1} = \frac{p_k}{1 - \beta_k + \beta_k p_k}, \quad (6)$$

⁵ For example, a biopsy might be the definitive test for a medical condition but starting with a simple blood test might be economically desirable.

such that probability p_k is increasing in k . We also define k_θ as the smallest k for which $p_k \geq \theta$. Hence, for all $k \geq k_\theta$, we have $\bar{r}(p_k) \geq \underline{r}(p_k)$, and diagnosing the customer as $\bar{\tau}$ is optimal should the diagnosis process stop.

Theorem 1 implies that the optimal policy of the one-sided case is characterized by a threshold $\bar{x}(k)$ for any given k (and therefore the corresponding p_k). That is, continuing the search is optimal when x is less than $\bar{x}(k)$ given that k tests have been positive thus far; when $x \geq \bar{x}(k)$, on the other hand, it is optimal to stop the search and identify the customer as type $\bar{\tau}$ if $k \geq k_\theta$, or as type $\underline{\tau}$ otherwise. However, in general, threshold $\bar{x}(k)$ is not monotone or unimodal in the number of tests. Therefore, the optimal policy cannot always be characterized as a search interval (or two thresholds) in k for a given x . Nonetheless, when the tests are well ordered (Definition 1), optimal threshold $\bar{x}(k)$ can be shown to be unimodal, which implies that the optimal policy can be described as a search interval in k . For the one-sided test case, well-ordered tests are equivalent to nonincreasing β_k s. The next proposition formally makes this point:

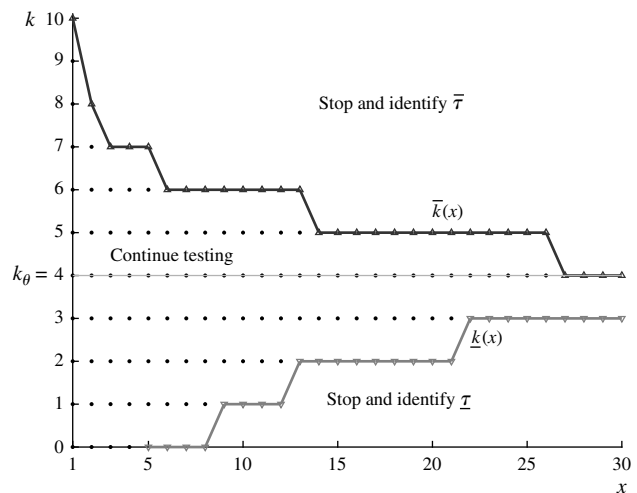
PROPOSITION 2. *When tests are one sided, if β_k is non-increasing in k , then the optimal rule can be characterized by two queue length dependent thresholds $\underline{k}(x)$ and $\bar{k}(x)$ such that $\underline{k}(x) < k_\theta \leq \bar{k}(x)$. Performing an additional test is optimal when $\underline{k}(x) < k < \bar{k}(x)$. Otherwise, it is optimal to stop the search and identify the customer as type $\underline{\tau}$ when $k \leq \underline{k}(x)$, or as type $\bar{\tau}$ when $k \geq \bar{k}(x)$. Furthermore, $\underline{k}(x)$ and $\bar{k}(x)$ are nondecreasing and nonincreasing in x , respectively.*

We should stress that Proposition 2 is not an immediate consequence of Theorem 1 and Proposition 1. Indeed, Proposition 2 requires that when the optimal decision in state (x, k) stops the process and identifies the customer as type $\underline{\tau}$, the same decision must also be optimal in state $(x, k - 1)$. This statement, however, is not directly supported by Theorem 1 or Proposition 1. The complete proof is in Appendix B.

Figure 2 depicts an example of the optimal policy for a one-sided system. The service provider performs up to 10 tests when a single customer is present in the system (i.e., $x = 1$ with $\bar{k}(1) = 10$ and $\underline{k}(1) < 0$). The maximum number of tests, $\bar{k}(\cdot)$, decreases with x to reach $k_\theta = 4$. At the same time, lower threshold $\underline{k}(\cdot)$ increases with congestion and eventually reaches $k_\theta - 1$, such that no test should be performed when $x \geq 26$.

More generally, Proposition 2 implies that the overall maximum number of customers allowed in the system is achieved when $k = k_\theta - 1$ or $k = k_\theta$. Note further that, because k always increases for a customer in service, in steady state, threshold $\underline{k}(\cdot)$ can only be reached with a customer arrival, that is, threshold

Figure 2 Optimal Policy for a One-Sided Diagnostic System with Congestion; $p_0 = 0.3$, $\beta_k = 0.5$ for all $k \geq 0$, $\rho = 0.1$, $\bar{v} = 100$, $\underline{v} = 500$, $\bar{c} = \underline{c} = 0$, and $c_w(x) = x$



$\underline{k}(\cdot)$ is always crossed from the left-hand side, and never from above. As a result, an alternative way of presenting the optimal structure is through the threshold $\bar{x}(k)$. That is, the agent should first let the maximum level of congestion allowed in the system increase with the number of performed tests. Only when enough tests have been run should the maximum number of customers in the system decrease. This structure highlights the distinction between our result and the commonly seen monotone threshold results in the existing queueing control literature.

The structure of the policy also retains the observation made in §3 that the service provider may make a diagnosis against the test results. For one-sided tests, she may interrupt the search and identify the customer as type $\underline{\tau}$ even though all test results thus far are positive. This never occurs in systems without congestion. With congestion, this can only occur in the one-sided case when the congestion level increases over the threshold during the elicitation of the first few tests (i.e., as long as $k < k_\theta$). On the other hand, if type $\bar{\tau}$ base rate is high enough (i.e., $p_0 \geq \theta$), the service provider never makes a diagnosis against the test results, as stated by the following proposition, which directly follows from Proposition 2 and the definition of k_θ :

PROPOSITION 3. *Assume that β_k is nonincreasing in k . The optimal policy for one-sided systems is fully characterized by nonincreasing thresholds $\bar{k}(x)$ if and only if $p_0 \geq \theta$. In this case, the agent performs an additional test if $k < \bar{k}(x)$ and identifies the customer as type $\bar{\tau}$ otherwise.*

In other words, when the base rate is large enough, the management of one-side diagnostic services under congestion is consistent with insights from systems without congestion (the service providers never make a diagnosis against the test results; see Appendix C)

and from the existing queueing literature (the optimal policy is characterized by one monotone threshold; see Hopp et al. 2007). When $p_0 < \theta$, however, managing diagnostics systems with congestion becomes significantly different.

5. Effect of Base Rate p_0

In this section, we explore further the impact of base rate p_0 . This allows deriving additional insights into the effect of congestion on the management of diagnostic services. In particular, the insights of this section disappear in systems with no congestion, i.e., when there is ample service capacity such that arriving customers always find an available server.

More precisely, we are interested in exploring how the optimal search intervals react to changes in p_0 . In a system without congestion, belief probability p is a sufficient statistic, and thus past test results, as well as p 's initial value p_0 , do not affect the optimal decision (see Appendix C, Proposition 5 for a formal proof). In diagnostic services with congestion, however, base rate p_0 reflects the service provider's belief about the types of customers waiting in queue. This influences the decision on whether to continue the current diagnostic process, or stop and start a new search on the next customer in line. Thus, the optimal policy of a diagnostic service with congestion should change with base rate p_0 . (This can also be seen from optimality equation (5).)


Our numerical study reveals that the optimal search interval shrinks as p_0 moves closer to θ . Figure 3 depicts the optimal thresholds for a system with identical and symmetrical tests where $\alpha_k = \beta_k = 0.6$ for all k , for two values of p_0 . Note that because tests are identical, the optimal thresholds do not depend on k (see §3). The solid thresholds correspond to the

optimal rule when the base rate is equal to $p_0^a = 0.6$. The dashed thresholds represent the optimal thresholds when the base rate equals $p_0^b = 0.5$. For this example $\theta = 0.44$ so that p_0^b is closer to θ than p_0^a (i.e., $\theta < p_0^b < p_0^a$).⁶ As it is evident from Figure 3, the dashed thresholds are within the solid thresholds. In other words, the search intervals shrink as p_0 approaches θ .

To investigate the intuition behind this phenomenon, consider state, for example, $(x = 20, p = 0.4)$ for both systems in Figure 3. The service provider faces the trade-off between running an additional test on the current customer, or stopping the search and starting diagnosing the next available customer. The service provider's belief about the current customer is the same ($p = 0.4$) in both systems. As a result, the expected value of performing an additional test is the same in both systems. However, her belief about the next customer type is equal to 0.6 and 0.5 in systems (a) and (b), respectively. Thus, the values of starting diagnosing the next customer in line differ in the two systems. As base rate p_0 gets closer to θ from p_0^a to p_0^b , the next customer type becomes more ambiguous. As a result, the marginal benefit of starting the search on a new customer increases. This explains why at state $(x = 20, p = 0.4)$, continuing the process for the current customer is optimal when base rate is $p_0^a = 0.6$, and starting the process on the next customer in line becomes optimal when $p_0^b = 0.5$.

More generally, at a given state (x, p) , when the decision of stopping the search and starting the next diagnose is optimal for a given value of p_0 , the decision remains optimal for values of p_0 that are closer to θ . Consequently, the optimal search interval shrinks. This is further illustrated by Figure 4, in which the congestion level is fixed at $x = 2$, and p_0 takes different values.⁷ The vertical axis corresponds to p_0 , and the horizontal one represents belief probability p . For a given p_0 , each dot indicates states $(x = 2, p)$ in which the service provider continues the search. The left- and right-hand side triangles denote the states where the service provider stops the search and diagnoses the customer as type τ and $\bar{\tau}$, respectively. As p_0 increases from 0 to θ the optimal search interval shrinks. But as p_0 continues to increase and moves away from θ , the interval expands again.⁸

Figure 3 Optimal Policy for $p_0^a = 0.6$ (Solid Thresholds) and $p_0^b = 0.5$ (Dashed Thresholds); $\alpha_k = \beta_k = 0.6$, for all k , $\rho = 0.2$, $\bar{v} = 700$, $\underline{v} = 550$, $\bar{c} = \underline{c} = 0$, and $c_w(x) = x$



⁶ In general, the posterior probabilities and hence the dots in Figure 3 should not be the same for different base rates. They coincide in our example because the tests are identical and symmetrical, and since we choose $p_0^b = 0.5 = \pi_k^-(0.6) = \pi_k^-(p_0^a)$.

⁷ The values of p_0 correspond to the different posterior probabilities determined by Bayes' rules (1–2), starting from $p_0 = 0.5$.

⁸ Figure 4 does not imply that more tests are conducted to each customer with p_0 farther away from θ . The accuracy/congestion trade-off may cause more or fewer tests conducted on a customer in steady state as p_0 gets closer to θ .

Figure 4 Search Interval for $x = 2$ as a Function of p_0 ; $\alpha_k = \beta_k = 0.6$, $\rho = 0.2$, $\bar{v} = 700$, $\underline{v} = 550$, $\bar{c} = \underline{c} = 0$, and $c_w(x) = x$




Table 1 Model Parameters for the Numerical Study

Parameters	Uniform distribution
ρ	[0.1, 0.6]
$\underline{c} + \underline{v}$	[10, 700]
α_k, β_k	[0.5, 0.8] for $k = 1, \dots, 5$
θ	[0.1, 0.9]
p_0^a	[0.1, 0.9]
p_0^b	$[p_0^a, \theta]$ or $[\theta, p_0^a]$

This observation is not limited to systems with identical and symmetrical tests. We have run a large scale numerical study with model parameters randomly generated from uniform distributions according to intervals listed in Table 1. In total we generated 10^6 cases. For all tested cases, we find that given any two systems (a) and (b), which differ only in their base rates where $p_0^a > p_0^b > \theta$ or $p_0^a < p_0^b < \theta$, there exist thresholds $\bar{p}^a(x, k)$, $\underline{p}^a(x, k)$, $\bar{p}^b(x, k)$, and $\underline{p}^b(x, k)$ such that $\underline{p}^a(x, k) < \underline{p}^b(x, k) < \theta < \bar{p}^b(x, k) < \bar{p}^a(x, k)$ for all state (x, k) .

6. Conclusion

This paper is the first to study how to dynamically perform diagnosis under time pressure in the form of congestion. We formulate this problem as a POMDP and demonstrate that the service provider should perform additional tests as long as her subjective

probability belongs to a given interval, the length of which decreases as congestion intensifies. This structure reveals several important aspects of managing diagnostic services under congestion, which significantly differ from more established search problems with no congestion.

First, diagnostic processes where tasks do not accumulate can be designed without knowledge of base rate p_0 . This is not true anymore for congested diagnostic services. This means, for instance, that a medical diagnostic process needs to account for the population of patients it serves and track changes in the base rate of the searched-for type. In particular, our numerical analysis suggests that the search interval should shrink as p_0 approaches θ .

Second, because of congestion effects, the agent should sometimes stop the diagnostic process and make a diagnosis against the latest test result. In the one-sided case, this means that the agent can identify the customer as a given type even though all test results have indicated otherwise. Decisions like these never occur in diagnostic systems with no congestion.

Finally, we find that in the one-sided case, the agent should first let the maximum level of congestion allowed in the system increase with the number of performed tests. Only when enough tests have been run should the maximum number of customers in the system decrease.

From a more technical point of view, some of our assumptions can actually be relaxed. In particular, Proposition 2 for the one-sided case holds under more general conditions. Specifically, two monotone thresholds are also optimal when $\mu_k \beta_k (1 - p_k)$ is nonincreasing in k , with test-dependent rate μ_k . This, however, requires a very different proof than the one offered in this paper, which does not directly extend to the two-sided asymmetrical case. Another point worth mentioning is what happens when the test elicitation times are not exponentially distributed. In general, this requires expanding the state space to include the time elapsed since the last event. However, if decisions are only made when a new customer arrives or a test is completed, the system can be cast into a discrete time format similar to our model (see Bertsekas 2007b, Chap. 5). In this case, we believe that most of our insights continue to hold.

Other natural extensions of our system should further help explore many other questions related to forming judgements under congestion. For instance, the diagnostic process could include nonhomogeneous costs of performing tests. Multiserver systems are also relevant as they may shed light on staffing rules for diagnostic services. In addition, agents may be in charge of both making the diagnosis and taking follow-up actions, which would raise the design problem of striking the right balance between providing the diagnostic and subsequent services. These extensions, however, may not always yield tractable models. A fruitful direction may consist in exploring simple heuristics that perform well.

Finally, our approach constitutes a very promising framework for understanding how individuals make actual decisions when tasks can accumulate. In fact, a similar diagnostic process without congestion has been proposed to represent how individuals decide (Busemeyer and Rapoport 1988). Psychologists have long recognized the importance of time pressure in human decision making. However, the cognitive environments psychologists consider are not typical of service organizations. In particular, situations where time pressure takes the form of accumulation of tasks have systematically been ignored. Our model naturally lends itself to experimental studies. Our results also offer a normative benchmark against which performances can be compared.

Appendices

Appendices A and B contain the proofs to our theoretical results. We first focus on the corresponding infinite horizon total discounted profit model in Appendix A and show various properties of the value function that are required for the structural results. Then, in Appendix B, we extend the results to the long-run average case, which completes the proofs. Appendix C provides more details on the no congestion case.

Appendix A. Analysis of the Total Discounted Profit Model

In this appendix, we consider a different version of the model in §3 in which the objective is to maximize the expected total discounted profit of the infinite horizon Markov decision process. The proofs focus on establishing value function properties that enable the structures of the optimal policy. We let γ represent the discount rate and assume, without loss of generality, that $\lambda + \mu + \gamma = 1$. Note that all results derived in this appendix are independent of this simplification and hold in general. Later, in Appendix B, we show that results from this section can be extended to the corresponding long-run average profit model by letting γ go to zero.

First, we can derive the optimality equation for the discounted case as $J = \Gamma J$, in which operator Γ is defined as

$$\begin{aligned} (\Gamma J)(x, k, p) &= \max \{ -c_w(x) + \lambda J(x+1, k, p) \\ &\quad + \mu(\alpha_k p + (1 - \beta_k)(1 - p))J(x, k+1, \pi_k^+(p)) \\ &\quad + \mu((1 - \alpha_k)p + \beta_k(1 - p))J(x, k+1, \pi_k^-(p)), \\ &\quad \bar{r}(p) + J(x-1, 0, p_0), \underline{r}(p) + J(x-1, 0, p_0) \}, \\ &\quad \text{for } x \geq 1 \text{ and } k \geq 0 \text{ and } 0 \leq p \leq 1, \quad (\text{A1}) \\ (\Gamma J)(0, 0, p_0) &= \lambda J(1, 0, p_0) + \mu J(0, 0, p_0). \end{aligned}$$

The two linear functions $\bar{r}(p)$ and $\underline{r}(p)$ represent the expected profit of announcing the customer as type $\bar{\tau}$ and $\underline{\tau}$, respectively, when the subjective probability is p ; $\bar{r}(p)$ is increasing in p , while $\underline{r}(p)$ decreases with p .

Before deriving the properties of the optimal value function, we need to impose the following assumption which is without loss of generality:

ASSUMPTION 1. Parameters \bar{v} , \underline{v} , \bar{c} and \underline{c} are such that, $r_0 := \max\{\bar{r}(p_0), \underline{r}(p_0)\} = 0$.

Recall from Equation (4) that changing parameters \bar{v} , \underline{v} , \bar{c} , and \underline{c} does not affect the optimal policy as long as the total rewards $\bar{v} + \bar{c}$ and $\underline{v} + \underline{c}$ remain unchanged. Assumption 1 is without loss of generality because, for any given \bar{v} , \underline{v} , \bar{c} , and \underline{c} , we can redistribute $\bar{v} + \bar{c}$ between \bar{v} and \bar{c} , and $\underline{v} + \underline{c}$ between \underline{v} and \underline{c} , so that the assumption is satisfied.

Let J_γ^* denote the solution to the Bellman's equation (A1). That is, J_γ^* represents the optimal value function for the discounted model when the discount rate is γ . Then, $J_\gamma^*(x, 0, p_0) \geq J_\gamma^*(x-1, 0, p_0) + r_0$ from the optimality equation, which reduces to $J_\gamma^*(x, 0, p_0) \geq J_\gamma^*(x-1, 0, p_0)$ using Assumption 1. From this, we deduce that $J_\gamma^*(x, 0, p_0) \geq 0$ and $J_\gamma^*(x, 0, p_0) \nearrow x$, which will be repeatedly used later in the proofs.

Next, we prove the following properties of the optimal value function J_γ^* :

C1. For any fixed x and k , $J_\gamma^*(x, k, p)$ is convex in p for $0 \leq p \leq 1$.

C2. For any fixed k and p , $J_\gamma^*(x, k, p) - J_\gamma^*(x-1, 0, p_0) \searrow x$.

C3. If the sequence of tests is well ordered, for any fixed x and p , $J_\gamma^*(x, k, p) \searrow k$.

To prove the above properties, we first prove the convergence of the value iteration algorithm and existence of the optimal value function.

LEMMA 1. The optimal value function, J_γ^* , satisfies the Bellman's equation (A1) such that $\Gamma J_\gamma^* = J_\gamma^*$, and can be obtained by the value iteration algorithm starting from any arbitrary function J_0 , i.e.:

$$\lim_{n \rightarrow \infty} \Gamma^{(n)} J_0 = J_\gamma^*.$$

PROOF. Because we have a maximization problem and the instantaneous costs $-c_w(x)$, $r(p)$, and $\bar{r}(p)$ are bounded from above, the negativity assumption holds and the result follows (Bertsekas 2007b, Vol. II, Proposition 3.1.6). \square

Having the existence of the optimal value function established, we next show its convexity in parameter p .

LEMMA 2. Operator Γ propagates property C1. Hence, the optimal value function, $J_\gamma^*(x, k, p)$, is convex in p for $0 \leq p \leq 1$.

PROOF. Consider the Bellman's equation (A1). First note that the maximum of a finite collection of convex functions is still convex. Moreover, the second and the third term in the maximization are linear and therefore convex in p . It remains to show the convexity of the first term in p . For this, we prove that $(\alpha_k p + (1 - \beta_k)(1 - p))J(x, k + 1, \pi_k^+(p))$ is convex in p . The proof for convexity of $((1 - \alpha_k)p + \beta_k(1 - p))J(x, k + 1, \pi_k^-(p))$ is exactly similar.

For notational simplicity, we use $J(p)$, α , β , and $\pi(p)$ to represent $J(x, k + 1, p)$, α_k , β_k , and $\pi_k^+(p)$, respectively. Assume $J(p)$ is convex in p , with a subgradient $\mathcal{F}(p) \in \partial J(p)$. That is,

$$J(p + \epsilon) \geq J(p) + \mathcal{F}(p)\epsilon.$$

Let

$$h(p) = (\alpha p + (1 - \beta)(1 - p))J(\pi(p))$$

in which

$$\pi(p) = \frac{\alpha p}{\alpha p + (1 - \beta)(1 - p)}.$$

Define

$$p_\epsilon = \frac{\alpha(p + \epsilon)}{\alpha(p + \epsilon) + (1 - \beta)(1 - p - \epsilon)}.$$

Note that for ϵ sufficiently small (i.e., $\epsilon < 1 - p$), we have $0 < p_\epsilon < 1$. Next,

$$\begin{aligned} h(p + \epsilon) - h(p) &= (\alpha(p + \epsilon) + (1 - \beta)(1 - p - \epsilon))J(p_\epsilon) \\ &\quad - (\alpha p + (1 - \beta)(1 - p))J(\pi(p)) \\ &\geq (\alpha(p + \epsilon) + (1 - \beta)(1 - p - \epsilon)) \\ &\quad \cdot (J(\pi(p)) + \mathcal{F}(\pi(p))(p_\epsilon - \pi(p))) \\ &\quad - (\alpha p + (1 - \beta)(1 - p))J(\pi(p)) \\ &= \epsilon \left((\alpha + \beta - 1)J(\pi(p)) \right. \\ &\quad \left. + \frac{\alpha(1 - \beta)}{\alpha p + (1 - \beta)(1 - p)} \mathcal{F}(\pi(p)) \right). \end{aligned}$$

In the first inequality, we have used both the convexity of J at $\pi(p)$ and the fact that the denominator for p_ϵ is positive. Therefore, $h(p)$ is convex in p with the following quantity as a subgradient:

$$(\alpha + \beta - 1)J(\pi(p)) + \frac{\alpha(1 - \beta)}{\alpha p + (1 - \beta)(1 - p)} \mathcal{F}(\pi(p)).$$

Thus, convexity of J in p implies convexity of ΓJ . Then, the convexity of J_γ^* in p directly follows from Lemma 1. \square

LEMMA 3. When $p = 1$, announcing $\bar{\tau}$ is the optimal decision so that $J_\gamma^*(x, k, 1) = \bar{v} + J_\gamma^*(x - 1, 0, p_0)$. Similarly, when $p = 0$, announcing $\underline{\tau}$ is the optimal decision so that $J_\gamma^*(x, k, 0) = \underline{v} + J_\gamma^*(x - 1, 0, p_0)$.

PROOF. Suppose the system is at state $(x, k, 1)$. If the search on the current customer continues for a (random) period of length $T > 0$, the server's belief about the customer's type remains equal to one during T because $\pi_k^+(1) = \pi_k^-(1) = 1$. After T , the customer is released and profit $e^{-\gamma T} \bar{v}$ is obtained. Alternatively, if the server releases the current customer right away and remains idle during T , the higher profit of \bar{v} is gained and less waiting cost is incurred during T . (In both cases, the system transits into state $(x_T + x - 1, 0, p_0)$ after T , where x_T is the number of arrivals during T .) Therefore, the optimal decision at state $(x, k, 1)$ is to stop and identify the customer as type $\bar{\tau}$.

The proof for the second part is similar. \square

Lemmas 2 and 3 will be used in Appendix B to prove Theorem 1. Next, we show that the optimal value function holds property C2.

LEMMA 4. The optimal value function, J_γ^* , is such that $J_\gamma^*(x, k, p) - J_\gamma^*(x - 1, 0, p_0) \searrow x$.

PROOF. Let T be a random variable representing the time needed to go from state $(x + 1, k, p)$ to state $(x, 0, p_0)$, when the system operates under the optimal policy. First, $T < \infty$ because otherwise, the queue length must always be higher than x , which can not be true under the optimal policy.

Now suppose the system is at state (x, k, p) and consider a policy u which, instead of following the optimal policy for the current system, follows the optimal policy for an otherwise identical system with a queue length of $x + 1$. This continues until the queue length drops to $x - 1$ for the first time. From that point onward, policy u coincides with the optimal policy for the current system. Define $\hat{c} = \min_x \{c_w(x + 1) - c_w(x)\}$. It follows that,

$$\begin{aligned} J_\gamma^u(x, k, p) &\geq (J_\gamma^*(x + 1, k, p) - \mathbb{E}[e^{-\gamma T}]J_\gamma^*(x, 0, p_0)) \\ &\quad + \mathbb{E}[e^{-\gamma T}]J_\gamma^*(x - 1, 0, p_0) + \mathbb{E}\left[\int_0^T \hat{c}e^{-\gamma t} dt\right]. \end{aligned}$$

The right-hand side of the above inequality is the value function under policy u except that we have used a lower bound for the waiting cost of the one additional customer during T .

Note that $\mathbb{E}\left[\int_0^T \hat{c}e^{-\gamma t} dt\right] \geq 0$, and hence this term can be dropped from the right-hand side. Next, because $J_\gamma^u(x, k, p)$ provides a lower bound for $J_\gamma^*(x, k, p)$, we have

$$\begin{aligned} J_\gamma^*(x, k, p) &\geq (J_\gamma^*(x + 1, k, p) - \mathbb{E}[e^{-\gamma T}]J_\gamma^*(x, 0, p_0)) \\ &\quad + \mathbb{E}[e^{-\gamma T}]J_\gamma^*(x - 1, 0, p_0) \\ &\Rightarrow J_\gamma^*(x, k, p) - \mathbb{E}[e^{-\gamma T}]J_\gamma^*(x - 1, 0, p_0) \\ &\geq J_\gamma^*(x + 1, k, p) - \mathbb{E}[e^{-\gamma T}]J_\gamma^*(x, 0, p_0). \end{aligned}$$

Furthermore, $J_\gamma^*(x, 0, p_0) \nearrow x$ implies that

$$-\mathbb{E}[1 - e^{-\gamma T}]J_\gamma^*(x - 1, 0, p_0) \geq -\mathbb{E}[1 - e^{-\gamma T}]J_\gamma^*(x, 0, p_0).$$

Adding this inequality to the one obtained above gives us

$$J_\gamma^*(x, k, p) - J_\gamma^*(x - 1, 0, p_0) \geq J_\gamma^*(x + 1, k, p) - J_\gamma^*(x, 0, p_0). \quad \square$$

According to Lemma 4, the added value of the current task is nonincreasing in the number of accumulated tasks, when the agent operates based on the optimal policy.

Finally, assume that the sequence of tests is well ordered. The next lemma describes the behavior of the optimal value function in k .

LEMMA 5. *If the sequence of tests is well ordered, the optimal value function, J_γ^* , is such that $J_\gamma^*(x, k, p) \searrow k$.*

PROOF. Suppose the sequence of tests is well ordered, and let ξ_{1k} and ξ_{2k} be the corresponding coefficients for the ordering of the $(k+1)$ st and the $(k+2)$ nd tests. Assuming that the function J holds property C3, we show that ΓJ preserves this property.

First, after plugging in the formulas for $\pi_{k+1}^+(\cdot)$ and $\pi_{k+1}^-(\cdot)$ and then replacing α_{k+1} and β_{k+1} by their equivalent values in terms of α_k and β_k , we can use the convexity of J in p (see Lemma 2, which ensures the convexity of the value function in each step of the value iteration algorithm) to get the following inequalities:

$$\begin{aligned} & J(x, k+2, \pi_{k+1}^+(p)) \\ & \leq \frac{\xi_{1k}(\alpha_k p + (1-\beta_k)(1-p))}{(\xi_{1k}\alpha_k + \xi_{2k}(1-\alpha_k))p + (\xi_{2k}\beta_k + \xi_{1k}(1-\beta_k))(1-p)} \\ & \quad \cdot J(x, k+2, \pi_k^+(p)) \\ & \quad + \frac{\xi_{2k}((1-\alpha_k)p + \beta_k(1-p))}{(\xi_{1k}\alpha_k + \xi_{2k}(1-\alpha_k))p + (\xi_{2k}\beta_k + \xi_{1k}(1-\beta_k))(1-p)} \\ & \quad \cdot J(x, k+2, \pi_k^-(p)), \end{aligned}$$

$$\begin{aligned} & J(x, k+2, \pi_{k+1}^-(p)) \\ & \leq ((1-\xi_{1k})(\alpha_k p + (1-\beta_k)(1-p))) \\ & \quad \cdot [((1-\xi_{1k})\alpha_k + (1-\xi_{2k})(1-\alpha_k))p \\ & \quad + ((1-\xi_{2k})\beta_k + (1-\xi_{1k})(1-\beta_k))(1-p)]^{-1} \\ & \quad \cdot J(x, k+2, \pi_k^+(p)) + ((1-\xi_{2k})((1-\alpha_k)p + \beta_k(1-p))) \\ & \quad \cdot [((1-\xi_{1k})\alpha_k + (1-\xi_{2k})(1-\alpha_k))p \\ & \quad + ((1-\xi_{2k})\beta_k + (1-\xi_{1k})(1-\beta_k))(1-p)]^{-1} \\ & \quad \cdot J(x, k+2, \pi_k^-(p)). \end{aligned}$$

Multiplying the first inequality by $(\alpha_{k+1}p + (1-\beta_{k+1})(1-p))$ and the second inequality by $((1-\alpha_{k+1})p + \beta_{k+1}(1-p))$ and adding them up leads to (after some algebraic simplifications)

$$\begin{aligned} & (\alpha_{k+1}p + (1-\beta_{k+1})(1-p))J(x, k+2, \pi_{k+1}^+(p)) \\ & \quad + ((1-\alpha_{k+1})p + \beta_{k+1}(1-p))J(x, k+2, \pi_{k+1}^-(p)) \\ & \leq (\alpha_k p + (1-\beta_k)(1-p))J(x, k+2, \pi_k^+(p)) \\ & \quad + ((1-\alpha_k)p + \beta_k(1-p))J(x, k+2, \pi_k^-(p)). \end{aligned}$$

This inequality together with property C3 implies that

$$\begin{aligned} & (\alpha_{k+1}p + (1-\beta_{k+1})(1-p))J(x, k+2, \pi_{k+1}^+(p)) \\ & \quad + ((1-\alpha_{k+1})p + \beta_{k+1}(1-p))J(x, k+2, \pi_{k+1}^-(p)) \\ & \leq (\alpha_k p + (1-\beta_k)(1-p))J(x, k+1, \pi_k^+(p)) \\ & \quad + ((1-\alpha_k)p + \beta_k(1-p))J(x, k+1, \pi_k^-(p)). \end{aligned}$$

Multiplying both sides by μ and another use of property C3 results in

$$\begin{aligned} & \lambda J(x+1, k+1, p) \\ & \quad + \mu(\alpha_{k+1}p + (1-\beta_{k+1})(1-p))J(x, k+2, \pi_{k+1}^+(p)) \\ & \quad + \mu((1-\alpha_{k+1})p + \beta_{k+1}(1-p))J(x, k+2, \pi_{k+1}^-(p)) \\ & \leq \lambda J(x+1, k, p) \\ & \quad + \mu(\alpha_k p + (1-\beta_k)(1-p))J(x, k+1, \pi_k^+(p)) \\ & \quad + \mu((1-\alpha_k)p + \beta_k(1-p))J(x, k+1, \pi_k^-(p)) \\ & \Rightarrow \Gamma J(x, k+1, p) \leq \Gamma J(x, k, p). \end{aligned}$$

Because Operator Γ propagates property C3, by Lemma 1, the optimal value function holds this property as well. This completes the proof. \square

Therefore, the optimal value function holds properties C1–C3.

Finally, we present the following lemma that limits the maximum number of customers allowed in the system for the discounted model.

LEMMA 6. *For any given $k \geq 0$ and $0 \leq p \leq 1$, there exists a finite value $\bar{x}(k, p)$ so that it is optimal to stop at state (x, k, p) with $x \geq \bar{x}(k, p)$. Further, finite \bar{x} exists so that it is always optimal to stop at state (x, k, p) with $x \geq \bar{x}$ for all k and p .*

PROOF. Define $\Phi_\gamma^*(x, k, p) = J_\gamma^*(x, k, p) - J_\gamma^*(x-1, 0, p_0)$ and rewrite the Bellman's equation (A1) for the optimal value function in the following form that is obtained by subtracting $J_\gamma^*(x-1, 0, p_0)$ from both sides:

$$\begin{aligned} & \Phi_\gamma^*(x, k, p) \\ & = \max\{-c_w(x) + \lambda\Phi_\gamma^*(x+1, k, p) + \lambda\Phi_\gamma^*(x, 0, p_0) \\ & \quad + \mu(\alpha_k p + (1-\beta_k)(1-p))\Phi_\gamma^*(x, k+1, \pi_k^+(p)) \\ & \quad + \mu((1-\alpha_k)p + \beta_k(1-p))\Phi_\gamma^*(x, k+1, \pi_k^-(p)) \\ & \quad - \gamma J_\gamma^*(x-1, 0, p_0), \bar{r}(p), \underline{r}(p)\}. \end{aligned}$$

Note that $-c_w(x)$ is strictly decreasing in x by assumption, all terms involving Φ_γ^* are nonincreasing in x from property C2, and $-\gamma J_\gamma^*(x-1, 0, p_0)$ is nonincreasing in x because $J_\gamma^*(x-1, 0, p_0) \nearrow x$. Thus, the first term in the maximization is strictly decreasing in x , and the second term ($\bar{r}(p)$) and the third term ($\underline{r}(p)$) do not change with x . It follows that as x increases, there is a break even point, denoted by $\bar{x}(k, p)$, above which $\max\{\bar{r}(p), \underline{r}(p)\}$ dominates the first term and it is optimal to stop.

Finally, we introduce $\bar{x} = \max_{(k,p)}\{\bar{x}(k, p)\}$. Then, it is never optimal to continue testing when $x \geq \bar{x}$, and all states (x, k, p) with $x \geq \bar{x}$ are transient states. \square

Note that all properties obtained in this appendix for the optimal value function of the discounted model hold for any value of the discount rate γ and regardless of the simplification $\lambda + \mu + \gamma = 1$.

Appendix B. Extension to the Long-Run Average Profit Model

In this appendix, we show that all the results derived for the total discounted profit model in Appendix A still hold

when we change the objective to maximizing the long-run average profit. Then, we use this fact to provide the proofs for the main results stated in the paper. To this end, we verify SEN conditions (Sennott 1999, Chap. 7) in our setting. If the three SEN conditions hold, the convergence of the total discounted profit model to the long-run average profit model is guaranteed, as we let γ go to zero. The SEN conditions can be verified as discussed below.

SEN1. Consider state $(0, 0, p_0)$. Suppose t_i indicates the arrival time of the i th customer in the future. If the agent can make the correct diagnosis for all upcoming tasks in zero time, the total discounted profit obtained is

$$\begin{aligned} E\left[\sum_{i=1}^{\infty}(p_0\bar{v} + (1-p_0)\underline{v})e^{-\gamma t_i}\right] &= (p_0\bar{v} + (1-p_0)\underline{v}) \sum_{i=1}^{\infty} E[e^{-\gamma t_i}] \\ &= (p_0\bar{v} + (1-p_0)\underline{v}) \sum_{i=1}^{\infty} \left(1 + \frac{\gamma}{\lambda}\right)^{-i} \\ &= (p_0\bar{v} + (1-p_0)\underline{v}) \frac{\lambda}{\gamma}. \end{aligned}$$

In the above derivation, we have used the fact that t_i has Erlang distribution with parameters λ and i . It follows that

$$\gamma J_{\gamma}^*(0, 0, p_0) \leq \gamma(p_0\bar{v} + (1-p_0)\underline{v}) \frac{\lambda}{\gamma} = \lambda(p_0\bar{v} + (1-p_0)\underline{v}) < \infty.$$

SEN2. Suppose the system is at state (x, k, p) . If the agent releases all existing customers without performing any test to reach state $(0, 0, p_0)$ and follows the optimal policy thereafter, the value of $\max\{\bar{r}(p), \underline{r}(p)\} + (x-1)r_0 + J_{\gamma}^*(0, 0, p_0)$ is generated, which gives us a lower bound for $J_{\gamma}^*(x, k, p)$. Thus, $J_{\gamma}^*(x, k, p) - J_{\gamma}^*(0, 0, p_0)$ has $\max\{\bar{r}(p), \underline{r}(p)\}$ as its lower bound which is further bounded from below by $\bar{r}(\theta)$.

SEN3. To verify this condition, we need to find a finite constant that uniformly bounds $J_{\gamma}^*(x, k, p) - J_{\gamma}^*(0, 0, p_0)$ from above, for all (x, k, p) and γ . Lemma 6 allows us to deduce that $J_{\gamma}^*(x, 0, p_0) = J_{\gamma}^*(\bar{x}, 0, p_0)$ for $x \geq \bar{x}$. On the other hand, $J_{\gamma}^*(x, 0, p_0) \nearrow x$ implies that $J_{\gamma}^*(x, 0, p_0) \leq J_{\gamma}^*(\bar{x}, 0, p_0)$ for $x < \bar{x}$. These two together give us

$$\begin{aligned} J_{\gamma}^*(x, k, p) - J_{\gamma}^*(0, 0, p_0) &\leq \max\{\bar{v}, \underline{v}\} + J_{\gamma}^*(x-1, 0, p_0) - J_{\gamma}^*(0, 0, p_0) \\ &\leq \max\{\bar{v}, \underline{v}\} + J_{\gamma}^*(\bar{x}, 0, p_0) - J_{\gamma}^*(0, 0, p_0) \\ &\leq (\bar{x}+1) \max\{\bar{v}, \underline{v}\} < \infty. \end{aligned}$$

The last inequality is obtained by assuming that when the system is at state $(\bar{x}, 0, p_0)$, all customers are correctly diagnosed without performing any test and the maximum possible value is collected from each task so that the system transits into state $(0, 0, p_0)$ in zero time.

Therefore, the SEN conditions are satisfied and all the results from Appendix A can be extended to the long-run average profit model of §3. In particular, the optimal value function associated with the model of §3, which we denote by J^* , holds properties C1–C3.

Now, we are in the position to present the proofs for the theoretical results stated in the paper.

PROOF OF THEOREM 1. For any fixed $x \geq 1$ and $k \geq 0$, denote the three terms in Bellman's equation (5) by $H(p)$,

$\bar{R}(p)$ and $\underline{R}(p)$, respectively. First, an argument similar to the proof of Lemma 2 together with property C1 shows that $H(p)$ is convex in p . Moreover, $\bar{R}(p)$ is linearly increasing and $\underline{R}(p)$ is linearly decreasing in p , and the two linear functions intersect at $p = \theta$.

There are two possible cases to consider:

(i) If $\max\{\bar{R}(p), \underline{R}(p)\} \geq H(p)$ for all $p \in [0, 1]$, then continuing search is always dominated by stopping. In this case, $\underline{p}(x, k) = \bar{p}(x, k) = \theta$.

(ii) Otherwise, we know from Lemma 3 that $\underline{R}(0) \geq H(0)$ and $\bar{R}(1) \geq H(1)$, and from the definition of $\bar{R}(\cdot)$ and $\underline{R}(\cdot)$ that $\underline{R}(0) > \bar{R}(0)$ and $\underline{R}(1) < \bar{R}(1)$. Then, the inequality $\max\{\bar{R}(p), \underline{R}(p)\} < H(p)$ for some p , together with convexity of $H(p)$ in p implies that $\bar{R}(p)$ and $\underline{R}(p)$ each must intersect with $H(p)$ exactly once in the interval $[0, 1]$. The intersection points are $\bar{p}(x, k)$ and $\underline{p}(x, k)$, respectively. In this case, $\underline{p}(x, k) < \theta < \bar{p}(x, k)$.

To show the monotonicity of thresholds, suppose it is optimal to stop testing and identify the customer as type $\bar{\tau}$ at state (x, k, p) , i.e., $J^*(x, k, p) = J^*(x-1, 0, p_0) + \bar{r}(p)$. From property C2 and because $J^*(x+1, k, p) \geq J^*(x, 0, p_0) + \bar{r}(p)$ (Bellman's equation (5)), we should have $J^*(x+1, k, p) = J^*(x, 0, p_0) + \bar{r}(p)$. This, in turn, implies that it is optimal to stop the search at state $(x+1, k, p)$ and announce $\bar{\tau}$. In other words, $p \geq \bar{p}(x, k)$ implies $p \geq \bar{p}(x+1, k)$, and hence $\bar{p}(x, k) \searrow x$.

Following the same logic, one can show that if it is optimal to stop testing and identify the customer as type $\underline{\tau}$ at state (x, k, p) , it is optimal to do so as well at state $(x+1, k, p)$, and hence $\underline{p}(x, k) \nearrow x$.

Finally, the existence of \bar{x} follows directly from Lemma 6. \square

PROOF OF PROPOSITION 1. From property C3, $J^*(x, k, p) \searrow k$. Suppose it is optimal to announce $\bar{\tau}$ at state (x, k, p) , i.e., $J^*(x, k, p) = J^*(x-1, 0, p_0) + \bar{r}(p)$. From the monotonicity of $J^*(x, k, p)$ in k and because $J^*(x, k+1, p) \geq J^*(x-1, 0, p_0) + \bar{r}(p)$ (Bellman's equation (5)), we should have $J^*(x, k+1, p) = J^*(x-1, 0, p_0) + \bar{r}(p)$. This, in turn, implies that it is optimal to stop the search at state $(x, k+1, p)$ and announce $\bar{\tau}$. In other words, $p \geq \bar{p}(x, k)$ implies $p \geq \bar{p}(x, k+1)$, and hence $\bar{p}(x, k) \searrow k$.

A similar argument shows that $p \leq \underline{p}(x, k)$ implies $p \leq \underline{p}(x, k+1)$, and hence $\underline{p}(x, k) \nearrow k$. \square

PROOF OF PROPOSITION 2. The proof is complete if we can show the following three results for the one-sided model:

(i) If it is optimal to stop and announce type $\bar{\tau}$ at state (x, k) , it is also optimal to stop and announce type $\bar{\tau}$ at state $(x, k+1)$.

(ii) If it is optimal to stop and announce type $\underline{\tau}$ at state (x, k) , it is also optimal to stop and announce type $\underline{\tau}$ at state $(x, k-1)$.

(iii) If it is optimal to stop and announce type $\bar{\tau}$ (respectively, $\underline{\tau}$) at state (x, k) , it is also optimal to stop and announce type $\bar{\tau}$ (respectively, $\underline{\tau}$) at state $(x+1, k)$.

Statements (i) and (ii) correspond to the existence of two thresholds for each value of x , and statement (iii) implies the monotonicity of these thresholds in x . To proceed, we use the results for the general two-sided model and note that state (x, k) in the one-sided model is equivalent to state (x, k, p_k) in the general case.

(i) Suppose it is optimal to stop and announce type $\bar{\tau}$ at state (x, k) (i.e., (x, k, p_k)). By Theorem 1, the same decision is optimal at state (x, k, p_{k+1}) because $p_{k+1} > p_k$. Then, Proposition 1 implies that announcing $\bar{\tau}$ is also optimal at state $(x, k+1, p_{k+1})$, which is equivalent to state $(x, k+1)$ in the one-sided case.

(ii) Next, suppose it is optimal to stop and announce type $\underline{\tau}$ at state (x, k) (i.e., (x, k, p_k)). Because $\pi_{k-1}^+(p_{k-1}) = p_k$, this translates into announcing $\underline{\tau}$ being optimal at state $(x, k, \pi_{k-1}^+(p_{k-1}))$. Then, from Theorem 1, the same decision is also optimal at state $(x, k, \pi_{k-1}^-(p_{k-1}))$ because $\pi_{k-1}^-(p_{k-1}) = 0 < \pi_{k-1}^+(p_{k-1})$. Further, again by monotonicity of thresholds in x as stated in Theorem 1, announcing $\underline{\tau}$ is optimal at states $(x', k, \pi_{k-1}^+(p_{k-1}))$ and $(x', k, \pi_{k-1}^-(p_{k-1}))$ for all $x' > x$. In particular, consider $x' = \bar{x} - 1$. From Lemma 6, stop decision is optimal at state $(\bar{x}, k-1, p_{k-1})$. Moreover, announcing $\underline{\tau}$ dominates announcing $\bar{\tau}$ at this state because $p_{k-1} < \theta$ (this is implied by $p_k = \pi_{k-1}^+(p_{k-1}) < \theta$). Hence, announcing $\underline{\tau}$ is the optimal decision at all three possible states that can be reached from state $(\bar{x} - 1, k-1, p_{k-1})$. It follows that the same decision should be optimal at state $(\bar{x} - 1, k-1, p_{k-1})$ as well. Iterative use of this argument implies that announcing $\underline{\tau}$ is optimal at all states $(x', k-1, p_{k-1})$ with $x' \geq x$, which are equivalent to states $(x', k-1)$ in the one-sided case.

(iii) The proof for this result is exactly the same as the proof for the monotonicity of thresholds in x as stated in Theorem 1. \square

PROOF OF PROPOSITION 3. The condition $p_0 \geq \theta$ is equivalent to $k_\theta = 0$. The necessity directly follows from the structure of the optimal policy described in Proposition 2. To prove the sufficiency, consider state $(\bar{x}, 0)$. According to Lemma 6, the optimal policy should release the customer at this state without performing any test. Because the optimal policy is fully characterized by thresholds $\bar{k}(x)$, we should have $\bar{k}(\bar{x}) = 0$. That is, the optimal decision at this state is to announce the customer as type $\bar{\tau}$. This implies that announcing $\bar{\tau}$ dominates announcing $\underline{\tau}$ at this state, which further implies that $\bar{r}(p_0) \geq r(p_0)$, and this is equivalent to $p_0 \geq \theta$. \square

Appendix C. No Congestion Case

This appendix formalizes the model and analysis of the diagnostic system without congestion. Various propositions presented in this appendix are referred to in the main text as a comparison to the system with congestion studied in the paper. The no congestion case in our context corresponds to systems with ample service capacity, i.e., with an infinite number of agents.

Diagnostic problems without congestion effects are traditionally modeled as sequential hypothesis testing problems (see, for instance, Wald 1947, DeGroot 1970) or search problems (Bertsekas 2007a), where each additional test incurs an exogenous cost. In our context, it can be considered as a system with infinite number of servers, such that arriving customers always find an available server. The problem consists in balancing the diagnosis accuracy against the search cost, c_s , which corresponds to the expected cost of holding one customer during the elicitation time, i.e., $c_s = c_w(1)/\mu$.

The objective is to maximize the expected total profit. The corresponding optimality equation is

$$\begin{aligned} J_s(k, p) &= (\Gamma J_s)(k, p) \\ &= \max\{-c_s + (\alpha_k p + (1 - \beta_k)(1 - p))J_s(k+1, \pi_k^+(p)) \\ &\quad + ((1 - \alpha_k)p + \beta_k(1 - p)) \\ &\quad \cdot J_s(k+1, \pi_k^-(p)), \bar{r}(p), r(p)\}. \end{aligned} \quad (C1)$$

Note that when the number of servers is infinite, the optimal value of objective function g^u (Equation (3)) is equal to $\lambda J_s(0, p_0)$. In other words, with ample capacity, diagnostic tasks do not accumulate in the sense that the time spent diagnosing one customer does not have any influence on the waiting time of the other customers. The problem becomes then separable.

The following characterization of the optimal policy for the diagnostic system without congestion can be directly implied from Theorem 1 and Proposition 1.

COROLLARY 1. *The optimal policy for the diagnostic system without congestion may be characterized by two thresholds $\underline{p}(k)$ and $\bar{p}(k)$ such that if state (p, k) satisfies $\underline{p}(k) < p < \bar{p}(k)$, it is optimal to continue testing. Otherwise, it is optimal to stop and identify the customer as type $\underline{\tau}$ when $p \leq \underline{p}(k)$, or as type $\bar{\tau}$ when $p \geq \bar{p}(k)$. Moreover, if the sequence of tests is well ordered, $\underline{p}(k)$ is nondecreasing and $\bar{p}(k)$ is nonincreasing in k .*

C.1. Type Identification Never Against Test Result

In the single diagnostic task case, the expected profit of identifying a customer as type $\underline{\tau}$ (respectively, $\bar{\tau}$) decreases after observing a positive (respectively, negative) test result. Therefore, the profit would have been higher had the service provider identified the customer as type $\underline{\tau}$ (respectively, $\bar{\tau}$) before performing the test, rather than after receiving an opposite signal. The next proposition formally states this intuition.

PROPOSITION 4. *In the search problem without congestion, it is never optimal to stop the search process and diagnose the customer as type $\underline{\tau}$ after a positive signal or as type $\bar{\tau}$ after a negative signal.*

PROOF. On the contrary to the statement, suppose it is optimal to continue at state (k, p) , but optimal to stop and announce $\bar{\tau}$ after a negative signal, i.e., at state $(k+1, \pi_k^-(p))$. It follows that $J_s^*(k+1, \pi_k^-(p)) = \bar{r}(\pi_k^-(p))$. Further, by Corollary 1, announcing $\bar{\tau}$ is also optimal at state $(k+1, \pi_k^+(p))$ because $\pi_k^+(p) > \pi_k^-(p)$, and hence $J_s^*(k+1, \pi_k^+(p)) = \bar{r}(\pi_k^+(p))$. Plugging in these values in Equation (C1) and some algebraic simplifications gives us

$$J_s^*(k, p) = \max\{-c_s + \bar{r}(p), \bar{r}(p), r(p)\},$$

which contradicts the assumption of continuing the search being optimal at state (k, p) . Thus, announcing $\bar{\tau}$ after a negative signal can not be optimal. The proof for the other case (announcing $\underline{\tau}$ after a positive signal) is similar. \square

C.2. Effect of Base Rate p_0

In the system without congestion, the current subjective probability p is already a sufficient statistic, regardless of where it is from, in particular, p_0 .

PROPOSITION 5. *The specifications of the optimal policy, $p(k)$ and $\bar{p}(k)$, in the diagnostic system without congestion do not change with the base rate p_0 .*

PROOF. The proof directly follows from Equation (C1) because the parameter p_0 is not present in the equation. Therefore, two different systems with different base rates share the same optimal policy as long as all other parameters are the same for the two systems. \square

References

- Anand KS, Paç MF, Veeraraghavan S (2011) Quality–speed conundrum: Trade-offs in customer-intensive services. *Management Sci.* 57(1):40–56.
- Bertsekas DP (2007a) *Dynamic Programming and Optimal Control*, Vol. I. (Athena Scientific, Nashua, NH).
- Bertsekas DP (2007b) *Dynamic Programming and Optimal Control*, Vol. II. (Athena Scientific, Nashua, NH).
- Blackwell D (1953) Equivalent comparisons of experiments. *Ann. Math. Statist.* 24(2):265–272.
- Bouns G (2003) Queueing models with admission and termination control-monotonicity and threshold results. Ph.D. thesis, Technische Universiteit Eindhoven, Eindhoven, The Netherlands.
- Breiman L, Friedman J, Olshen R, Stone C (1984) *Classification and Regression Trees* (Chapman and Hall, New York).
- Bussemeyer J, Rapoport A (1988) Psychological models of deferred decision making. *J. Math. Psych.* 32(2):91–134.
- Crabill T (1972) Optimal control of a service facility with variable exponential service times and constant arrival rate. *Management Sci.* 18(9):560–566.
- Crabill T, Gross D, Magazine M (1977) A classified bibliography of research on optimal design and control of queues. *Oper. Res.* 25(2):219–232.
- DeGroot M (1970) *Optimal Statistical Decisions* (McGraw-Hill, New York).
- de Véricourt F, Zhou Y-P (2005) Managing response time in a call-routing problem with service failure. *Oper. Res.* 53(6):968–981.
- Edwards W (1965) Optimal strategies for seeking information: Models for statistics, choice reaction times, and human information processing. *J. Math. Psych.* 2(2):312–329.
- George J, Harrison J (2001) Dynamic control of a queue with adjustable service rate. *Oper. Res.* 49(5):720–731.
- Gerdtz M, Bucknall T (2001) Triage nurses' clinical decision making: An observational study of urgency assessment. *J. Adv. Nursing* 35(4):550–561.
- Goldstein D, Gigerenzer G (1999) The recognition heuristic, how ignorance makes us smart. Gigerenzer G, Todd P, eds. *Simple Heuristics That Make Us Smart* (ABC Research Group, Oxford University Press, New York), 37–58.
- Guide VJ, Wassenhove L (2001) Managing product returns for remanufacturing. *Production Oper. Management* 10(2):142–155.
- Hopp W, Irvani S, Yuen G (2007) Operations systems with discretionary task completion. *Management Sci.* 53(1):61–77.
- HP Renew Program (2009) HP Renew refurbishment checklist. Retrieved April 15, 2010, <http://h20427.www2.hp.com/program/hprenew/ap/en/Refurbishment.html>.
- Sennott LI (1999) *Stochastic Dynamic Programming and the Control of Queueing Systems* (Wiley-Interscience, New York).
- Stidham S, Weber R (1993) A survey of Markov decision models for control of networks of queues. *Queueing Systems* 13(1): 291–314.
- Travers D (1999) Triage: How long does it take? How long should it take? *J. Emergency Nursing* 25(3):238–240.
- Wald A (1947) *Sequential Analysis* (Wiley & Sons, New York).
- Wang X, Debo LG, Scheller-Wolf A, Smith SF (2010) Design and analysis of diagnostic service centers. *Management Sci.* 56(11):1873–1890.