Judgement Accuracy Under Congestion In Service Systems

Francis de Véricourt Peng Sun

July 10, 2009

Abstract

When serving a customer, service providers typically need to weigh the value of taking time to make an accurate judgement against the cost of delaying the provision of service to others in the system. Our paper presents the first analysis of how to dynamically manage this judgement accuracy/congestion tradeoff. To that end, we study an elementary system where a service provider faces a random stream of customers. The service consists of a basic judgement task. The key feature of our approach is an explicit representation of the agent's decision process as an elicitation sequence of binary probabilistic cues. When the objective is to maximize average profit to the firm, we show that the maximum number of elicited cues required to make the decision should decrease as the number of customers in the system increases. This structure yields several counter-intuitive results: i) Increasing cue validity may actually decrease overall judgement accuracy; ii) The level of congestion may increase with increasing cue validity; iii) Increasing the cue elicitation rate can increase the level of congestion. Further, as an alternative to optimal policies, we propose several simple cognitive heuristics adapted to congestion, which we construe as a form of endogenous time pressure. Our study suggests that judgements based only on the most relevant piece of information perform reasonably well. When basing judgements on all available information, simple fixed threshold rules appear to be very robust. These findings are consistent with prior results for decision making under exogenous time pressure.

Keywords: Service Operations, Queueing Theory, Dynamic Programming, Decision Making, Information Search, Bayes Rule.

1 Introduction

The accuracy of employees' judgement and decisions has an obvious impact on the value offered to customers in service delivery systems. Even small and rare inaccuracies can inflict significant aggregate damage on organization and its customers, given the high volume of activity observed in many service delivery systems. A dramatic example is the number of people who die annually in the United States from medical errors, a number larger than the number of people who die in motor vehicle accidents (Institute of Medicine 2001). In general, accumulating additional information is likely to improve decision and judgement accuracy. Accumulating and processing information, however, takes time and therefore may increase customer waiting times. The service provider is then required to weigh the benefit of collecting additional information against the cost of delaying the provision of service to others in the system.

Tradeoffs such as this are most prevalent in diagnostic services where service providers focus on forming a judgement, but do not perform themselves any subsequent treatments that may be indicated by the diagnosis. Triage nurse systems provide a typical example. Triage is a dynamic decision-making process that assesses patient condition for medical care (Gerdtz and Bucknall 2001). The nurse elicits different cues based on the patient's history and physical assessment to help form an accurate judgement. On the other hand, long triage processes can result in adverse patient outcomes (Travers 1999). Another example of accuracy/congestion tradeoff in a diagnostic service setting 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. Cues to inform this judgement can be intangible and this assessment task requires specialized technicians. On the other hand, the maintenance industry is subject to intensive time-based competition. A similar judgement 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 make accurate judgements under congestion also frequently occurs in some non-diagnostic systems, such as support centers and help desks (de Véricourt and Zhou 2005).

In this paper, we gain insight into the problem of dynamically balancing judgement accuracy against delays to reach a decision. To that end we study an elementary system where a service provider faces a random stream of customers, each of whom is one of two types. The service consists in determining the customer type. Accuracy is then defined as the probability that the customer type is correctly diagnosed. The key feature of our approach is to explicitly represent the judgement process of the service provider. In our framework, judgements about the unknown state of the world are based on binary probabilistic cues. The decision process consists of eliciting these cues one by one. Given a cue value, the normative approach assumes the service provider updates her prior probability regarding a customer's type according to Bayes' rule. The choice of stopping the cue search and reaching a decision depends on the level of congestion in the system as well as the agent's belief (i.e., her subjective probability distribution). Cues are imperfect and false negatives may occur while false positives are negligible. This assumption greatly simplifies the analysis while retaining the Bayesian updating mechanism. Our model therefore captures the most fundamental elements of making a judgement under time pressure in the form of congestion. A similar approach can be used when false negatives are negligible. Our main managerial insights also hold in general for more intricate cases.

We first characterize the optimal decision rules that maximize the long run average value to the firm, which includes rewards for identifying customer type and waiting costs. We show that the maximum number of elicited cues required to make a decision should decrease with the number of customers present in the system. Such decision policies yield counter-intuitive insights into improving the service capabilities. In particular, the firm may seek to use better cues, which correctly indicate a customer's type with higher probabilities. Surprisingly, we show that due to congestion effects, using better cues can decrease accuracy. Better cues should also lead to faster services as fewer cues might be needed to make the diagnosis. Nonetheless, we show that improving validity of cues can increase congestion. Increasing the cue elicitation rate has a similar impact. These results are illustrated on an efficient frontier which is derived from our model. Thus far, the study assumes that the decision maker relies on all available information and updates her belief using a normative Bayes' rule, which is cognitively demanding. Psychologists have suggested that cognitively simpler heuristics can perform well (Gigerenzer and Goldstein 1996). In line with this approach, we propose and study several cognitively simple decision rules adapted to environments with congestion.

Our representation of the decision process is an elementary version of typical cognitive processes studied by psychologists. The approach is related to Brunswik's lens model (Hammond and Stewart 2001) in which inference about unknown states of the world are based on probabilistic cues. However, eliciting and comparing cues may be cognitively demanding and time consuming. Therefore, the fast and frugal heuristics program (Goldstein and Gigerenzer 1999) promotes the study of simple cognitive rules for making decisions when time and cognitive resources are limited (a related approach is the accuracy-effort framework of Payne, Bettman, and Johnson 1993). Fast and frugal heuristics typically rely on a fraction of available information and combines cue values using simple rules to reach a decision.

Psychologists, however, generally have not studied the types of cognitive environments that fundamentally characterize service delivery systems. In particular, the above research considers isolated decision tasks, in which time pressure usually takes the form of deadlines (Payne et al. 1993, Rieskamp and Hoffrage 1999) or opportunity costs (Payne et al. 1996) and is exogenously determined. This cannot account for congestions and delays which are crucial aspects of service systems. As in our model, the time spent on a given task affects the system congestion and hence the level of time pressure imposed on subsequent tasks.

In other words, this paper is concerned with decision rules adapted to environments where time pressure is *endogenized*. In fact, as we show in our numerical study, decision rules that ignore congestion phenomena sometimes perform poorly in these environments. One contribution of this paper is to propose fast and frugal heuristics that are adaptive under this form of endogenized time pressure. We study two simple heuristics, each of which relies on a separate piece of information: the current congestion level or the number of elicited cues, respectively. Our study suggests that judgements based only on the most relevant piece of information perform reasonably well. When the decision is based on all available information, we show that simple combination rules based on fixed-thresholds perform extremely well.

Our model can be thought of as a search problem with endogenous costs. In the classical search problem (Bertsekas 2007a, Vol. I, Chapter 5.4.4), the decision maker searches for a treasure of a given value. Each search incurs a fixed cost and generates a false negative with a given probability. The problem is to decide when to stop in order to maximize total expected value. Typically, the cost of each search is exogenously given and a single task is considered. Psychologists have also studied actual decision processes in similar frameworks (Wallsten et al. 2005; Bearden and Connolly 2007). In our case, however, the cost is determined by the congestion resulting from the decision time. In particular, we consider a stream of search tasks such that the time spent on a given task affects the value of subsequent searches. Search costs are therefore endogenously determined.

Alternatively, our model corresponds to a system with discretionary task completion where employees determine how much time to allocate to a job. Hopp et al. (2007) analyzes such systems and represents quality (or profit) as a concave increasing function of service time. George and Harrison (2001) also study a queueing system where the service rate is dynamically changed so as to minimize waiting and capacity costs. In this stream of research, congestion effects are the main source of uncertainty. By contrast, we model quality as judgement accuracy. This requires considering additional sources of uncertainty related to the decision process and the associated dynamics. We are, for instance, able to derive an accuracy/congestion efficient frontier, which cannot be obtained from the previous models. This also means that exiting proof techniques and results cannot be directly extended to our case. Further, our model allows defining heuristics akin to simple cognitive rules studied by psychologists. For instance we consider the first impression policy which relies on a single elicited cue. (This rule corresponds to a special case of the so-called Take-The-Best heuristic proposed by Gigerenzer and Goldstein 1996, or the Single Variable rule introduced by Hogarth and Karelaia 2007 for decision problems without congestion.) Such cognitive rules cannot be easily defined in the frameworks presented by Hopp et al. (2007) or George and Harrison (2001).

Anand et al. (2008) also study the impact of a similar quality/congestion trade-off on customer behavior in an equilibrium framework. Quality is modelled as a decreasing function in service capacity that does not capture the judgement updating mechanisms we are interested in. To our knowledge, Wang, Debo, Scheller-Wolf, and Smith (2008) is the only work concerned with judgement accuracy/congestion tradeoffs. Wang et al. (2008) study a call center of triage nurses. The decision process, however, is not dynamic in their case, and the emphasis is on customer behavior. They find system equilibriums such that the effect of nurse skills on congestion is not monotonic. This finding is analogous with some of the effects on congestion described in the present paper. The underlying reasons are, however, quite different. Their results are driven by customer behaviors, which do not play any role in our model.

We present the model in the next section. The optimal decision rule is characterized in Section 3. We establish our counter-intuitive managerial insights in Section 4 and illustrate then with an accuracy/congestion frontier in Section 5. We study cognitively simple decision rules in Section 6. We discuss other managerial implications of our findings and conclude in Section 7.

2 A Model of Judgement Accuracy under Congestion

Our model captures the most fundamental elements of making a judgement under time pressure in the form of congestion, while still being tractable. In particular, we model the stochastic dynamics of both an agent's belief and a queueing process. Consider a service provider serving customers arriving according to a Poisson process with rate λ. The service consists in determining whether a customer is of type τ or not. To that end, the service provider elicits cues one by one. Each elicitation time is exponentially distributed with rate μ. The system is preemptive in the sense that the elicitation of a cue can be stopped at any time. We denote by $\rho \equiv \lambda/\mu$ the elicitation utilization rate, that is, the average number of customers arriving during an elicitation period. Cues are not perfect and false negatives may occur. More precisely, if a customer is of type τ, after k elicited cues fail to identify the type, the next cue will reveal it with probability β_k, which is independent of the previous cues. The cue elicitation process stops when a cue indicates type τ. Otherwise, the service provider updates her prior probability that the customer is of type τ and may decide to elicit an additional cue depending on the number of customers in the system and her subjective probability distribution.

Identifying a type τ customer brings value V to the system. This reward also represents the benefit of subsequent services that τ customers may receive. This corresponds to the value of repairing a part in maintenance services, of re-using a good in a remanufacturing system, etc. On the other hand, missing and releasing a type τ customer incurs a mismatch cost C. This corresponds to the disutility of not providing required health care or the expected cost of potential failures when the part is re-used. Important special cases includes $V = 0$ or $C = 0$, where the service focuses on generating value or avoiding mismatch costs, respectively. In some applications, such as in a health care setting, minimizing mismatch costs is arguably more relevant than generating profit, which is captured by our model with $V = 0$ and $C > 0$. The decision maker needs to balance these rewards against the cost of delay. We consider a waiting cost $c(x)$ per unit of time which is incurred when x customers are present in the system. We impose no restrictions on $c(x)$ other than it being non-decreasing in $x \geq 0$ with $c(0) = 0$ and unbounded (i.e., $\lim_{x\to\infty} c(x) = +\infty$).

We denote by p_k the probability of a customer being type τ after the k^{th} cue elicitation does not identify the type, where p_0 is given and represents the proportion of type τ customers in the population. In other words, probability p_k represents the current belief of the service provider that the customer is of type τ. Probability p_k evolves according to Bayes' rule,

$$p_{k+1} = f_k(p_k) \equiv \frac{p_k(1-\beta_k)}{p_k(1-\beta_k)+1-p_k} = \frac{1-\beta_k}{1-\beta_k p_k} p_k \ .$$

That is, if the customer has not been identified as type τ after the kth cue, then

$$p_k = f_{k-1} \circ f_{k-2} \circ \ldots \circ f_0(p_0) \ .$$

We assume that the cues are such that the probability of finding a type τ customer is non-increasing in the the number of elicited cues.

Assumption 1 *$p_k\beta_k$ is non-increasing in k.*

Assumption 1 essentially requires that β_k does not increase too fast in k. In particular, the assumption is satisfied when β_k is decreasing in k, that is, when cues with highest validity are elicited first. This sequence is actually optimal for the corresponding search problem without congestion. With congestion, if the firm can choose the sequence in which cues are elicited, choosing them in order of decreasing validity maximizes total profit (as shown in Appendix A), and satisfies Assumption 1. Further, decision making studies suggest that in some environments, individuals elicit cues in order of decreasing validity (Rieskamp and Hoffrage 1999). An important special case is when cues are identical with $\beta = \beta_k$ for all k, which corresponds to running the same test several times. Assumption 1 also allows for a limited number of cues by taking $\beta_k = 0$ when k is larger than a given κ.

At any time, the decision maker needs to decide whether to elicit a new cue or to terminate the decision process and move to the next customer. Such a control policy makes a dynamic tradeoff between accuracy and congestion. The performance of a decision rule is then measured as the the long-run average profit. For control policy u, we define the corresponding long run average profit g^u as,

$$g^u \equiv \liminf_{T\to\infty} \frac{1}{T}\mathrm{E}\left[VN^u(T) - CM^u(T) - \int_0^T c(X^u(t))dt\right], \tag{1}$$

where $N^u(t)$ is the random cumulative number of identified type τ customers up to t, $M^u(t)$ is the number of missed type τ customers up to t, and $X^u(t)$ the random process corresponding to the number of customers in the system at time t.

Equation (1) captures an accuracy/congestion tradeoff. Given control policy u, accuracy A^u is defined as the probability that the system identifies a type τ customer given that she is of type τ. It follows that,

$$A^u = \liminf_{T\to\infty} \frac{\mathrm{E}\left[N^u(T)\right]}{p_0\lambda T}.$$

Similarly, congestion is measured by L^u, the average number of customers in the system, which is also proportional to the average waiting time by Little's Law. We have,

$$L^u = \liminf_{T\to\infty} \frac{\mathrm{E}\left[\int_0^T X^u(t)dt\right]}{T}.$$

The long run average profit is then equal to

$$g^u = p_0\lambda\left[(V+C)A^u - C\right] - \liminf_{T\to\infty}\frac{1}{T}\mathrm{E}\left[\int_0^T c(X^u(t))dt\right], \tag{2}$$

which holds from the conservation of flow of τ customers. In particular, with linear waiting cost where $c(x) = c\times x$, this equation reduces to $g^u = p_0\lambda(V+C)A^u - c\,L^u - p_0\lambda C$ which is linear in A^u and L^u with coefficients of opposite signs. We show in Section 5 how this formulation leads to an accuracy/congestion efficient frontier.

3 Optimal Control Policy

In our setting achieving the best decision accuracy under congestion consists in maximizing the long run average profit. In the following, our goal is to determine the optimal policy u^* that maximizes g^u. We formulate the problem as a Markov Decision Process (MDP). The system state is given by $(x,k)\in\mathbb{Z}^+\times\mathbb{Z}^+$ where x indicates the current number of customers in the system and k the number of cues that have been elicited so far. To that end, we uniformize the system and assume, without loss of generality, that $\lambda+\mu=1$. The optimality equations that the optimal average profit g^* and value function J^* need to satisfy are then equivalent to (see Appendix B),

$$g + J(x,k) = \begin{cases} \max\{ & g + J(x-1,0) - p_kC, \\ & \pi(x,k) + \lambda J(x+1,k) \quad + p_k\mu\beta_k J(x-1,0) \\ & \qquad\qquad + (1-p_k\beta_k)\mu J(x,k+1)\}, \quad x>0 \\ & \lambda J(1,0) + \mu J(0,0), \qquad\qquad x=0 \end{cases} \tag{3}$$

with expected instantaneous profit defined as $\pi(x,k) \equiv \mu p_k \beta_k (V+C) - c(x)$, which is non-increasing in x and k (from Assumption 1).

The analysis of this problem is quite intricate. We present the complete proof in Appendix C, where we first consider the expected total discounted profit of the infinite horizon MDP associated with policy u and discount rate γ. We characterize the optimal policy for the discounted system by showing that the marginal increase in x of the optimal discounted profit is decreasing in x (see Lemma 5 in Appendix C). We then show that the optimal policy is of the non-increasing threshold type (see Lemma 6 in Appendix C). Next, we extend this characterization to the average profit case by letting γ go to zero (see Appendix C.2). This leads to our main theoretical result,

Theorem 1 *The optimal average profit g^* satisfies optimality equation (3). Further, the optimal policy is characterized by a queue-length dependent threshold $\hat{k}(x)$ such that the service provider stops eliciting cues if and only if $k \geq \hat{k}(x)$. Moreover, $\hat{k}(x)$ is non-increasing in x.*

The structure of the optimal threshold $\hat{k}(x)$ is quite complex in general and we present heuristics that are cognitively less demanding in Section 6. We can nonetheless provide a bound on $\hat{k}(x)$ using the myopic policy which elicits additional cues as long as $\pi(x,k) \geq 0$. Define then $k^{myo}(x)$ as the smallest k such that $\pi(x,k) < 0$, where $k^{myo}(x)$ is non-increasing in x. We have,

Proposition 1 *$\hat{k}(x) \leq k^{myo}(x)$ for all x.*

The result holds from Lemma 5 in Appendix C and from iterating on the value function.

Figure 1 depicts an example of the optimal policy for $V = 500$, $C = 0$, $p_0 = 0.7$, $\beta_k = 0.5$ for all k, and $\rho = 0.1$. The optimal thresholds are represented by the black line above the shaded area. The light line depicts the myopic thresholds which are larger than the optimal ones consistent with Proposition 1. The service provider continues to elicit cues as long as the current state of the system lies in the shaded area, which also corresponds to the recurrent states. When only one customer is present in the system, the service provider elicits up to 10 cues. This number falls to 7 when 2 customers are present. When there are 40-67 customers, the decision is based on one cue only. The maximum number of customers allowed in the system is then 67. Note that the maximum number of elicited cues is decreasing in x, in accordance with Theorem 1.

When $\hat{k}(x)$ is null for all x, the policy is degenerate. In this case, the long run average profit is negative if the system engages in serving customers. The following result provides necessary and sufficient conditions for not eliciting a single cue.

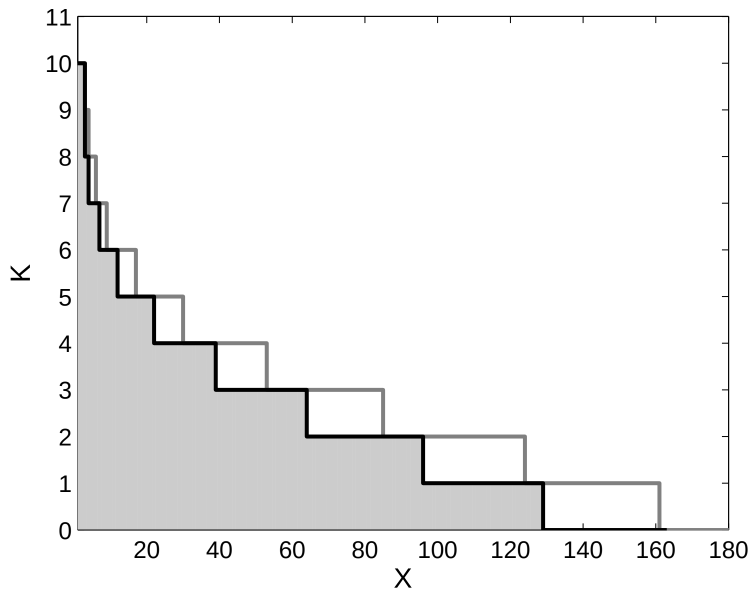

Figure 1: Optimal Policy and Myopic Thresholds

Proposition 2 *It is optimal not to serve customers (i.e., not to elicit any cue) if and only if* $\mu p_0 \beta_0 (V + C) \leq c(1)$.

The proof can be found in Appendix D.1.

In other words, if the expected revenue $p_0\beta_0(V + C)$ is lower than the average cost incurred during an elicitation period with one customer $c(1)/\mu$, no customer should be accepted in the system.

The previous structure of the optimal policy appears natural. It has, however, some surprising implications for the management of judgement accuracy under congestion, as discussed in the next section.

4 Improving Service Capability

The firm has several levers to improve service capability. For example, an increase in the elicitation rate or cue validity should make it possible to identify τ customers more quickly. In this section, we explore analytically and numerically the impact of such changes on accuracy and congestion.

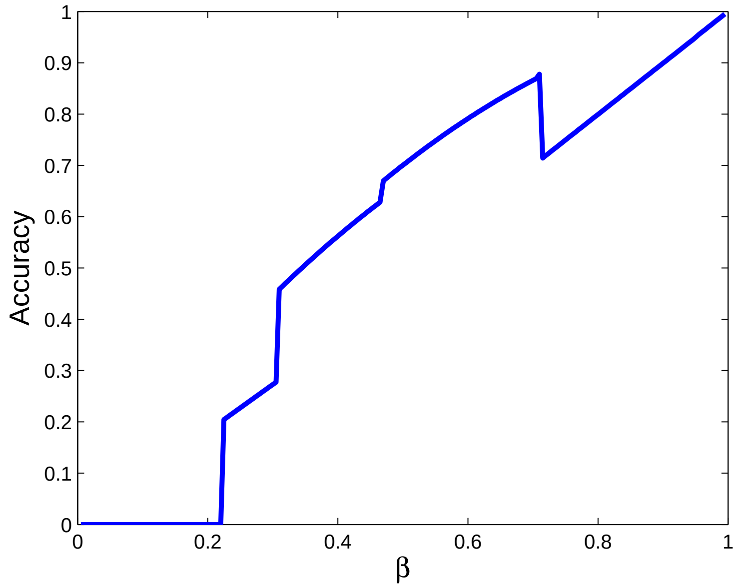

Figure 2: Impact of Cue Validity on Accuracy.

4.1 Cue Validity

4.1.1 Accuracy

The organization may seek to improve its service capability in order to provide higher accuracy. In particular, the firm can train its employees or introduce new pieces of equipment to increase cue validity. Enabling employees to better identify τ customers should increase accuracy as well. In fact, this effect can easily be shown for a simpler system where customers never wait (which is equivalent to considering an infinite number of employees). The systems reduces then to a classical stopping time problem for which we have,

Proposition 3 *For a system without congestion, accuracy always increases in cue validity.*

The proof is presented in Appendix D.2.

However, this result does not hold for accuracy under congestion, and the previous intuition may mislead companies in their effort to improve capability. Indeed, Figure 2 depicts the impact of β, with $\beta = \beta_k$ for all k, $\rho = 0.1$, $p_0 = 0.1$, and $(V + C)/c = 50$ with linear waiting cost $c(x) = c \times x$. When β increases from 0.7 to 0.75, accuracy decreases from 0.86 to 0.75. This reflects a decrease in the number of elicited cues for the corresponding optimal policies. The effect of cue validity on accuracy presented in Figure 2 is not rare. We have tested 1000 cases by varying $(V + C)/c$, ρ and p_0, where $(V + C)/c$ takes 10 values $(10, 20, 30, 50, 70, 100, 150, 200, 300, 500)$, and ρ and p_0 each take

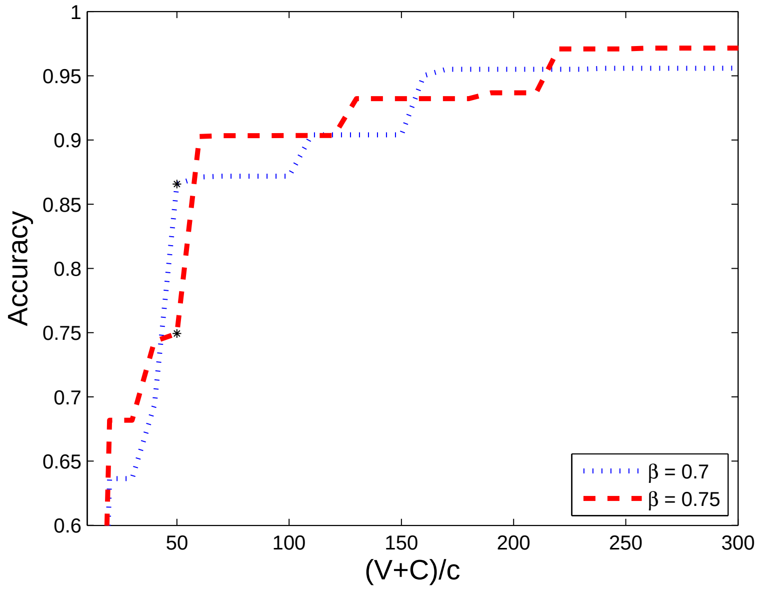

Figure 3: Effect of $(V + C)/c$ on Accuracy for different values of β

10 values $(0.05, 0.15, 0.25, \ldots, 0.95)$. For 28.6% of all cases (or 286 cases), accuracy decreases at least once in β. This, along with Proposition 3, leads to the following counter-intuitive observation.

Observation 1 *In the presence of congestion effects, increasing cue validity can decrease accuracy.*

The reason for this effect is that an increase in β may make it beneficial to elicit fewer cues per customer on average in order to reduce congestion. This happens when the value of identifying a type τ customer is not much higher than the waiting costs. Otherwise, trading accuracy for reduced congestion is never advantageous. Figure 3 illustrates this point and depicts accuracy as a function of $(V + C)/c$ for $\beta = 0.7$ and $\beta = 0.75$ when $\rho = 0.1$ and $p_0 = 0.1$. When $(V+C)/c \leq 220$, the two accuracy curves cross several times. In particular, when $(V + C)/c = 50$, the increase of β from 0.7 to 0.75 causes a decrease of accuracy from 0.86 to 0.75. However, when $(V + C)/c \geq 220$, accuracy for $\beta = 0.75$ always dominates accuracy when $\beta = 0.7$.

4.1.2 Congestion

Cues reveal a type τ customer more often as cue validity increases. Identifying a type τ customer may then require fewer cues. This might lead service firms to conclude that

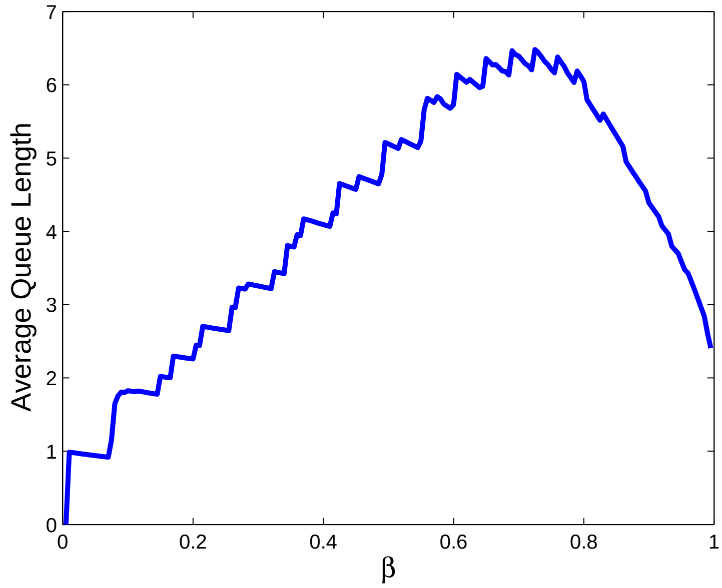

Figure 4: Impact of Cue Validity on Congestion.

increasing validity reduces congestion. The following proposition, however, shows that this effect is not systematic.

Proposition 4 *Increases in cue validity can lead to increases in the average number of customers in the system.*

Proof See Appendix D.3

Figure 4 depicts the average number of customers as a function of β, for $(V+C)/c = 300$, $\rho = 0.7$ and $p_0 = 0.9$. When β is small, accuracy and congestion are low. Hence, an increase in β brings the highest relative benefit when used to improve accuracy. On the other hand, congestion increases. When β becomes large, accuracy and congestion are high. It becomes more beneficial to reduce congestion. The average number of customers then decreases in β. Congestion reaches a peak at $\beta = 0.73$.

This behavior almost always appears in our numerical study. Congestion can increase in β for more than 94% of the tested cases.

4.2 Cue Elicitation Rate

A more direct approach to reduce congestion consists in increasing the cue elicitation rate. Indeed, systems with congestion typically seek to improve their service capacity in order to reduce waiting times. This may be achieved in our setting by hiring more experienced

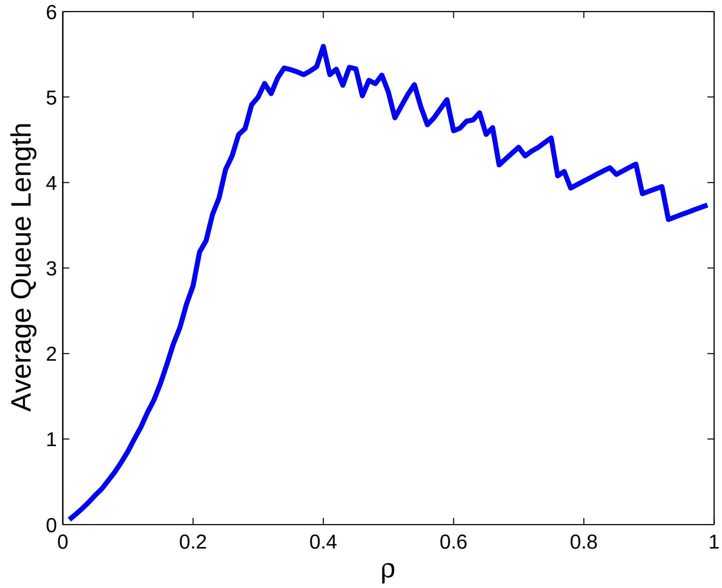

Figure 5: Impact of ρ on Congestion.

employees, training them, improving technology, etc. However, the following observation shows that improving service capacity can in fact intensify congestion.

Figure 5 depicts the impact of ρ on congestion for $\beta = 0.35$, $(V + C)/c = 500$ and $p_0 = 0.8$. When $\rho > 0.4$, the average number of customers has a decreasing trend in ρ. The intuition behind this effect is that an increase in μ allows one to elicit more cues. This can improve accuracy but also may intensify congestion. Such a phenomenon is also quite typical. Congestion can increase in ρ for more than 79% of the tested cases. We can therefore make the following observation:

Observation 2 *Increasing the cue elicitation rate can increase the average number of customers in the system.*

Hopp et al. (2007) have observed a similar effect in their numerical study for systems with service time-dependant revenues. They observed in their model that the phenomenon that higher μ increases congestion is more prevalent when the maximum potential revenue is high. In our numerical study, however, we observe that this phenomenon is typical for small values of $(V + C)/c$. Figure 6 depicts such an example with $\beta = 0.75$ and $p_0 = 0.75$. When ρ increases from 0.7 to 0.9, the congestion decreases when $(V + C)/c$ is relatively low, and always increases when $(V + C)/c > 170$.

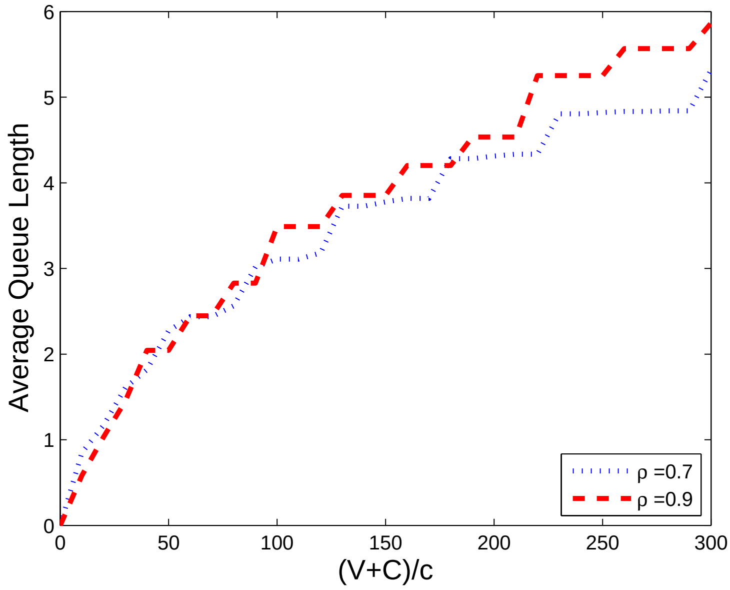

Figure 6: Impact of ρ on Congestion for different values of $(V + C)/c$.

4.3 Base Rate Effects

The firm does not directly control the underlying base rate of type τ customers. Nonetheless, everything else being equal, the organization may expect to offer greater accuracy when serving populations (or markets) with higher type τ base rates. This is because type τ customers are more frequent and thus more easily identifiable. This result is actually always true for systems without congestion in which customers never wait. Nonetheless, accuracy can decrease as the base rate increases in congested systems. However, this effect is much less prevalent than in the cue validity case. In our numerical study, we again tested 1000 cases, with $(V+C)/c$ taking 10 values $(10, 20, 30, 50, 70, 100, 150, 200, 300, 500)$, ρ and β each takes 10 values respectively $(0.05, 0.15, 0.25, \ldots, 0.95)$. We find that accuracy decreases with base rate in fewer than 1% of all cases.

On the other hand, type τ customers are more frequent at higher base rates. Therefore, fewer cues may be required to identify these customers. This should lead to less congestion. The next proposition, however, shows that this is not always the case.

Proposition 5 *Increasing the base rate of type τ customers can increase the average number of customers in the system.*

See Appendix D.3 for the proof.

Figure 7 shows, for instance, that when p_0 increases from 0.1 to 0.5, the average number of customers increases from 1.14 to 2.47, for $(V + C)/c = 300$, $\rho = 0.5$ and $\beta = 0.85$. This

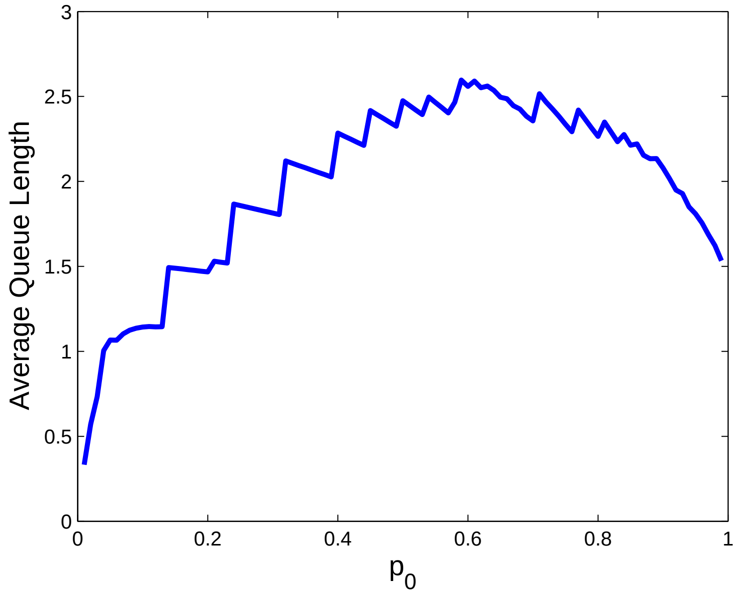

Figure 7: Impact of Base Rate on Congestion.

behavior is similar to Proposition 4. Numerically, we find that in almost all cases (996 out of 1000) the average queue length can increase with p_0.

In short, the effect of p_0 on congestion and accuracy are similar (albeit less significant) to those of β. The underlying reasons driving such results are also similar.

5 Efficient Frontier

In this section we present a complementary alternative approach based on an accuracy/congestion efficient frontier. With linear waiting costs, our model naturally produces an efficient frontier for the accuracy/congestion tradeoff. More formally, given system parameters $\{\beta_k\}$, ρ, p_0, function $A = e_{\beta,\rho,p_0}(L)$ determines the highest possible level of accuracy A that can be achieved in a system with congestion intensity L. This frontier is characterized by the following proposition,

Proposition 6 *When $c(x)$ is a linear function, Efficient Frontier $A = e_{\beta,\rho,p_0}(L)$ exists and is increasing concave in L.*

Proof: See Appendix D.4.

In particular, no control policy can achieve an accuracy level such that $A > e_{\beta,\rho,p_0}(L)$. A manager may therefore choose an A/L tradeoff on this frontier and implement the corresponding policy (which can be evaluated by retrieving the corresponding value of $(V+C)/c$).

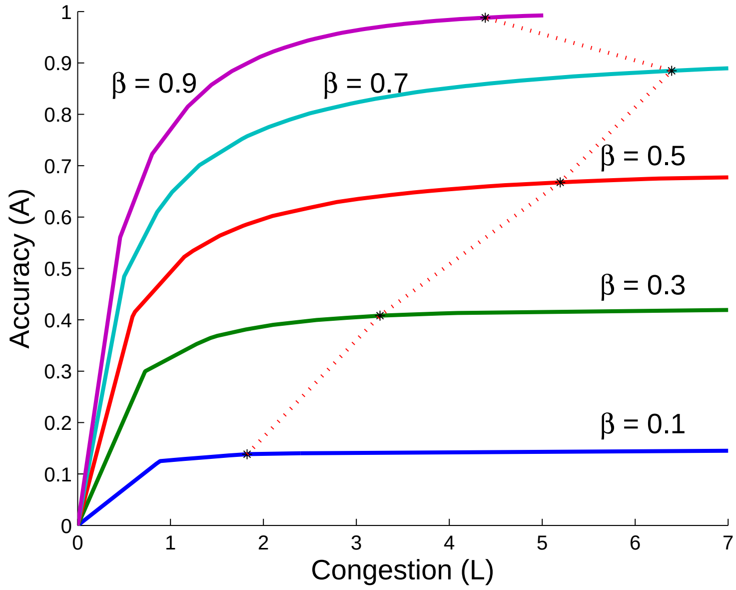

Figure 8: Impact of Base Rate on Congestion.

Further, the efficient frontier improves with the system parameters:

Proposition 7 *The efficient frontier* $A = e_{\beta,\rho,p_0}(L)$ *is point-wise increasing in* β *and* p_0 *and point-wise decreasing in* ρ.

The proof is in Appendix D.5.

Proposition 7 affords a different perspective on the insights of the previous section. For example, Figure 8 depicts the efficient frontiers for five different values of β, with $\rho = 0.7$ and $p_0 = 0.9$. Consistent with the previous result, the efficient frontier increases in β. Further, the dashed line intersects with these different frontiers. Each intersection point represents the A/L tradeoff that the optimal policy generates when $(V + C)/c = 300$, for the corresponding β. While accuracy always improves along the dash line (as β increases) congestion deteriorates from $\beta = 0.1$ to $\beta = 0.7$ and improves only from $\beta = 0.7$ to $\beta = 0.9$. This is consistent with Figure 4.

6 Simple Cognitive Heuristics

In this section, we study cognitively simple decision rules for decision making under congestion. Our representation of the decision process is a special case of the sequential elicitation of probabilistic cues that psychologists typically study, and elucidates the structure of simple and efficient cognitive rules that are adapted to environments with congestion.

Psychologists have argued that in some circumstances simple cognitive heuristics can be very accurate while outperforming other policies in speed. For instance, Gigerenzer and Goldstein (1996) introduced the "Take-the-Best" heuristic which elicits cues one at a time and stops the process as soon as a cue discriminates among alternatives [1]. Such "one-reason" decision processes appear to be very efficient (Gigerenzer and Goldstein 1996), especially under high (exogenous) time pressure (Payne, Bettman, and Johnson 1988).

However, this literature in decision making has focused on isolated decision tasks where time pressure is imposed exogenously in the form of either deadlines or opportunity costs. By contrast, we are concerned with decisions made in environments typical of service delivery systems. In particular, the service provider is subject to endogenous time pressure: the time spent on a given customer determines the level of congestion and hence the level of time pressure experienced in subsequent decisions.

Cognitive heuristics proposed by psychologists typically rely on a fraction of the available information. These rules perform well if the selected pieces of information are the most relevant ones. In our settings, the minimal fraction of information available to a decision maker is either the level of congestion (as measured by the current number of customers in the system) or the number of cues that has been elicited so far. More demanding cognitive processes may track both types of information, but use an elementary rule to combine them.

6.1 Heuristics Based on a Minimal Fraction of Available Information

One possible candidate for a fast and frugal heuristic under congestion is to rely solely on the number of elicited cues. This means disregarding congestion effects and treating this endogenous time pressure as aggregate opportunity cost. This form of exogenous time pressure is actually consistent with prior research in psychology (Payne et al. 1996). We refer to such heuristics as ignoring-the-queue policies (IQ). This class of policies limits the maximum number of cues while ignoring the actual number of customers in the system. This corresponds to thresholds such that $k(x) = \kappa$ for all x. We set the maximum number of elicited cues κ so as to maximize profit within this class of policies. The problem reduces then to a single decision task where elicitation costs are cue-dependant.

Alternatively, the decision maker may base her decision on one cue only and monitor the queue length. Eliciting a single cue corresponds in our framework to "one-reason" heuristics such as Take-The-Best proposed for exogenous time pressure. We refer to this rule as the first impression policy (1^{st}Imp). (Hogarth and Karelaia 2007 also consider this heuristic for a general Brunswick lens model without time pressure, which is referred to as the Single

[1] Take-the-Best actually requires that cues are elicited according to their validity. This corresponds in our model to decreasing β_k, which satisfies Assumption 1.

β	0.1	0.3	0.5	0.7	0.9	0.99		
λ/μ	0.01	0.05	0.1	0.3	0.5	0.7	0.9	0.99
$V+C$	10	30	50	100	300	500		
p_0	0.1	0.3	0.5	0.7	0.9	0.99		

Table 1: Model parameters.

Variable rule in their settings.) This corresponds to thresholds such that $k(x) = 1$ for $x \geq \chi$ and $k(x) = 0$ otherwise. Again, the maximum number of customers χ is set so as to maximize profit within this class of policies. (We present closed form solutions for IQ and 1^{st}Imp policies in Appendix E.)

IQ and 1^{st}Imp rely on a minimal fraction of available information. We are interested in identifying under which circumstances the level of congestion or the number of cues affords sufficient information to make accurate decisions. To that end we first characterize systems for which a non-degenerated first impression policy is optimal.

Proposition 8 *The first impression policy is optimal and not degenerated if and only if*

$$\mu\beta_1 p_1(V+C) \leq c(1) < \mu\beta_0 p_0(V+C) \ . \tag{4}$$

The proof can be found in Appendix D.6.

Proposition 8 is surprising as congestion does not play a direct role in (4). This result is actually equivalent to the optimality condition for eliciting at most one cue, under exogenous time pressure with opportunity cost $c(1)$ per unit of time. In other words the performance of the 1^{st}Imp should be robust to the type of time pressure. In particular, if the average elicitation cost with one customer ($c(1)/\mu$) is larger than the expected value of the first cue, the first impression policy is optimal. But if the cost of eliciting a cue is too large, the system stops serving customers and the policy is degenerate.

These insights are confirmed in our numerical study. We have evaluated $1,728$ combinations of model parameters (see Appendix E for a short description of the algorithm we use). The choices of parameters are summarized in Table 1. A total of 160 combinations satisfy the conditions of Proposition 2 under which it is optimal not to serve any customer. We focus our study on the remaining $1,568$ cases.

To study the performance of a heuristic policy h, we evaluate the corresponding long run average profit g^h starting with an empty system ($x = k = 0$ at $t = 0$). We assume linear waiting cost and normalize $V + C$ such that $c(x) = x$. This means that $V + C$ represents the relative value of the rewards compared to a unit cost. For each value of $V + C$, we test the two extreme cases $V = 0$ and $C = 0$. Furthermore, $\beta_k = \beta$, for all k which also satisfies

	Percentiles					
	5^{th}	10^{th}	50^{th}	90^{th}	95^{th}	Mean
IQ	0%	0%	5.39%	100%	100%	30.22%
1^{st}Imp	0%	0%	13.74%	63.14%	69.86%	23.15%
Best of 2	0%	0%	0.79%	28.13%	40.96	8.37%

Table 2: Distribution of 1^{st}Imp and IQ relative gaps when $V > 0$ and $C = 0$.

	Percentiles					
	5^{th}	10^{th}	50^{th}	90^{th}	95^{th}	Mean
IQ	0%	0%	7.18%	36.08%	51.66%	12.99%
1^{st}Imp	0%	0%	12.56%	87.16%	94.16%	30.12%
Best of 2	0%	0%	2.34%	13.44%	18.05	4.81%

Table 3: Distribution of 1^{st}Imp and IQ relative gaps when $C > 0$ and $V = 0$.

Assumption 1. When $V > 0$ and $C = 0$, we maximize profit, and set the performance metric to be the relative optimality gap defined as $(g^* - g^h)/g^*$. When $V > 0$ and $C = 0$, g^* and g^h are negative. We use the performance metric $(g^h - g^*)/g^h$ so that it also takes values between 0 and 1.

Tables 2 and 3 indicate that these simple policies are not robust in general, with a maximum optimality gap larger than 50% for at least 10% of cases for both IQ and 1^{st}Imp. This is not surprising as both rules rely on a minimum fraction of the available information. Their performances are, however, reasonable for the 50% of cases that have a gap of less than 5.39% and 13.74%, respectively, when $C = 0$, and 7.18% and 13.44% when $V = 0$. Moreover, both heuristics are optimal for at least 10% of the tested scenarios no matter whether $C = 0$ or $V = 0$.

Studies in decision making under exogenous time pressure have emphasized the importance of relying on the most relevant piece of information for "one-reason" rules to perform well. If the wrong piece of information is used, the performance can dramatically deteriorate (see, for example, Payne, Bettman, and Johnson 1993). In our setting, the most relevant piece of information is determined by whether IQ or 1^{st}Imp performs best. Table 2 indicates that the best of both policies performs quite well. The average gap drops to 8.21% and optimality is nearly achieved for at least 50% of the cases. This is because IQ and 1^{st}Imp are complementary to some extent and perform well in different scenarios.

Figure 9 illustrates this point in further detail and depicts the efficient frontiers of the optimal policy as well as the best of 1^{st}Imp and IQ, for $\rho = .5$, $\beta = .8$ and $p_0 = .9$. The

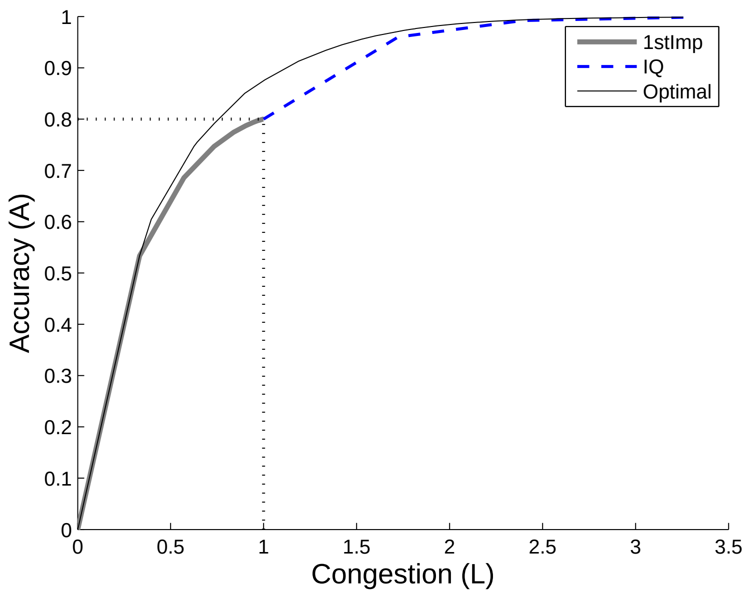

Figure 9: IQ and 1^{st}Imp Efficient Frontiers, $\rho = .5$, $\beta = .8$ and $p_0 = .9$

1^{st}Imp efficient frontier dominates that of IQ when the desired accuracy and congestion levels are below .8 and 1, respectively. Indeed, the best possible accuracy level that 1^{st}Imp can achieve is β. This happens when all customers are served, and no limit is put on queue length. The system corresponds then to an M/M/1 for which the average number of customers is equal to $\rho/(1-\rho)$, that is 1 with $\rho = .5$. This congestion level constitutes the maximum queue length that can be achieved with 1^{st}Imp. But 1^{st}Imp with no limit on queue length is equivalent to IQ with one elicited cue only. Hence β and $\rho/(1-\rho)$ correspond to the minimum accuracy and congestion levels that IQ can achieve, respectively. This point on the curve is also where the distance to the optimal policy frontier is the greatest.

Hence, relying only on the most relevant piece of information appears sufficient to achieve a reasonable level of performance in most cases. On the other hand, using the wrong piece can dramatically damage the system performance. This finding is consistent with the general claim advanced by research programs in decision making, such as the fast and frugal heuristics framework (Gigerenzer and Todd 1999).

6.2 Threshold Rules

More demanding cognitive heuristics may rely on all available pieces of information, but use simple rules to combine them. One of the main difficulties of applying an optimal policy is to remember and track queue-length dependant thresholds. As an alternative, we propose to consider fixed-threshold heuristics (FixT), which are characterized by thresholds χ and κ

	Percentiles					
	5%	10^{th}	50^{th}	90^{th}	95^{th}	Mean
FixT	0%	0%	0.08%	3.69%	5.16%	0.99%

Table 4: Distribution of FixT gap when $C = 0$ and $V > 0$.

	Percentiles					
	5^{th}	10^{th}	50^{th}	90^{th}	95^{th}	Mean
FixT	0%	0%	0.37%	5.43%	7.96%	1.86%

Table 5: Distribution of FixT gap when $V = 0$ and $C > 0$.

on the numbers of customers and elicited cues, respectively. Again, we consider the optimal policy within this class.

Threshold-type policies are often used to represent cognitive decision processes (see, for instance, Bearden and Connolly 2007 and references therein). These heuristics are also common in the queueing literature. In particular, our fixed-threshold policy can be thought of as a special case of the double-threshold policy (DT) proposed by (Hopp et al. 2007) (which has two thresholds for the service time and one for the queue length).

Tables 4 and 5 show that FixT performs extremely well, especially for the case with $V > 0$, where the average relative gap is below 1%, and for more than 90% of the tested cases, the gap is below 4%. The robustness of FixT is also surprising. As a comparison, DT requires more parameters to achieve a similar level of performance.

Because the optimal thresholds are decreasing, the system is in state $(x, 0)$ for large x or state $(1, k)$ for large k with low probabilities under the optimal policy (see Figure 1). FixT essentially cuts at least one of these tails, which does not significantly affect the system performance. On the other hand, FixT does not work well when both tails have high probabilities to be reached. This, however, does not happen in most cases. Out of the $3,136$ tested cases, the gap is greater than 10% in only 85 cases. These situations are not easy to characterize but it appears that FixT does not perform well in general for mild to high β's and high $(V + C)/c$'s.

Figure 10 illustrates the robustness of FixT. For the same parameters as in Figure 9, the efficient frontier of FixT is almost exactly equal to the optimal one. When A is lower (resp. larger) than .8, FixT mimics 1^{st}Imp (resp. IQ) which performs very well. For A around .8, neither 1^{st}Imp nor IQ works well. FixT, on the other hand, relies on all available information and adjust both thresholds χ and κ.

The excellent performance of FixT shown in Figure 10 prevails in our study. FixT

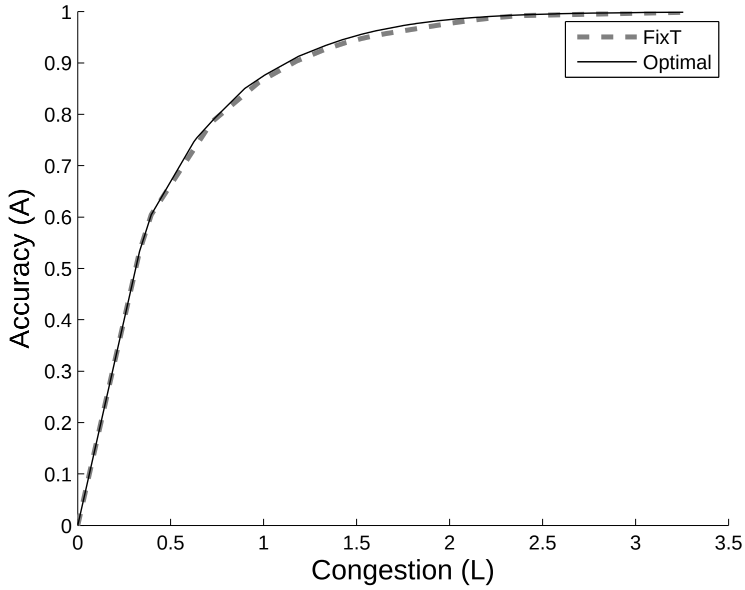

Figure 10: FixT Efficient Frontier, $\rho = .5$, $\beta = .8$ and $p_0 = .9$

policies not only generate close-to-optimal profits, but also achieve the most efficient accuracy/congestion tradeoffs almost perfectly.

7 Discussion and Conclusion

This paper presents the first study of dynamic judgement processes under congestion. We have shown that, because of congestion effects, judgement accuracy can decrease with cue validity and, to a lesser extent, with the base rate. Improving either the elicitation rate or validity of cues may also increase congestion. Initiatives to improve service capabilities that appear natural can lead to counter-productive results. Further, we have explored cognitively simple decision rules adapted to environments with congestion, in which time pressure is endogenously determined. Fixed-Threshold policies, which rely on all available information but use a simple combining rule, perform extremely well. We have also identified environments to which heuristics relying on a small fraction of the information are adapted. All these decision rules decrease the number of elicited cues as the number of customers increases, which is also the structure of the optimal policy.

Our model captures the fundamental dynamics of forming a judgement with congestion effects. Its simplicity allows one to isolate important insights while keeping the problem tractable. To the best of our knowledge, this is the first representation of the dynamic accuracy/congestion tradeoff, which also constitutes a valuable contribution of the paper.

Natural extensions of our system should help explore many other questions related to forming judgements under congestion.

For example, our study is mostly relevant for systems with few servers. When the number of agents are large, such as in call centers, the previous decision/accuracy tradeoff is embedded inside a staffing plan. A typical approach to set staffing levels consists in first determining the time an agent should spend serving each customer. The number of employees are then set to limit congestion. When the objective is to provide a high level of accuracy, this first step corresponds to defining the number of cues that should be elicited per customer. This, in turn, determines the service time. Our results can help shed some light onto the performance of such hierarchical procedures. With n servers, the system can be roughly approximated by a single server system with elicitation rate $n\mu$. The approach described above essentially leads to a system controlled by an IQ policy. We know from Section 6, however, that this rule does not work well in general. This suggests that the previous hierarchical approach will likely need to overstaff the system in order to contain congestion.

Alternatively, FixT can easily be extended to a multi-server system: if the queue-length reaches a threshold on x the server that has elicited the most cues releases its customer. An intriguing question is then how the number of servers should increase with the offered load ρ. In particular, one of the main results of queueing theory states that staffing levels should be set at capacity plus a square-root term in the offered load. Given the counter intuitive results established in this paper, the square-root structure of the staffing rule might not be adapted to the accuracy/congetsion tradeoff. A heavy-traffic analysis which approximates the queue length with a diffusion process could help answer this important question.

We also consider systems which focus solely on judgement tasks, where potential subsequent services are delivered elsewhere. In many instances however, employees need both to make the judgement and take follow-up actions. Our main findings, especially those presented in Section 4, should still hold for this more general case. The non-diagnostic tasks may, however, produce congestion effects that could override our results. This is especially the case when the corresponding service time is longer than the judgement process. This raises the important design problem of establishing the right balance between diagnosing and providing subsequent services.

Finally, our approach constitutes a very promising framework to understand how individuals make actual decisions when tasks can accumulate. 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, congestion phenomena which endogenize the level of time pressure have systematically been ignored. Our model naturally lends itself to experimental studies. The experimental design essentially extends existing frameworks by assigning the subject to a random stream of decision tasks that eventually accumulate, instead of considering isolated tasks. Further, our results indicate that fixed-threshold policies perform extremely well and are very robust. They are not too cognitively demanding, which suggests that individuals could use related decision rules. Individuals may well be adapted to environments with congestion. This, at least, constitutes a first hypothesis worth testing.

References

Anand, K., M. Pac, and S. Veeraraghavan (2008). Quality-speed conundrum: tradeoffs in labor-intensive services. *Wharton School working paper*.

Bearden, J. and T. Connolly (2007). Multi-attribute sequential search. *Organizational Behavior and Human Decision Processes 103*(1), 147–158.

Bertsekas, R. (2007a). *Dynamic Programming and Optimal Control*, Volume I. Athena Scientific.

Bertsekas, R. (2007b). *Dynamic Programming and Optimal Control*, Volume II. Athena Scientific.

Bertsimas, D. and J. Tsitsiklis (1997). *Introduction to Linear Optimization.* Athena Scientific, Belmont, MA.

de Véricourt, F. and Y. Zhou (2005). Managing Response Time in a Call-Routing Problem with Service Failure. *Operations Research 53*(6), 968.

George, J. and J. Harrison (2001). Dynamic control of a queue with adjustable service rate. *Operations Research 49*(5), 720–731.

Gerdtz, M. and T. Bucknall (2001). Triage nurses' clinical decision making. An observational study of urgency assessment. *Journal of Advanced Nursing 35*(4), 550–561.

Gigerenzer, G. and D. Goldstein (1996). Reasoning the Fast and Frugal Way: Models of Bounded Rationality. *Psychological Review 103*, 650–669.

Gigerenzer, G. and P. Todd (1999). *Simple heuristics that make us smart.* Oxford University Press New York.

Goldstein, D. and G. Gigerenzer (1999). The recognition heuristic, how ignorance makes us smart. In G. Gigerenzer, P. Todd, A. R. Group, and I. NetLibrary (Eds.), *Simple heuristics that make us smart*, pp. 37–58. Oxford University Press New York.

Gross, D. and C. Harris (1998). *Fundamentals of queueing theory.* John Wiley & Sons, Inc. New York, NY, USA.

Guide, V. J. and L. Wassenhove (2001). Managing product returns for remanufacturing. *Production and Operations Management 10*(2), 142–155.

Hammond, K. and T. Stewart (2001). *The Essential Brunswick: Beginnings, Explications, Applications.* New York, NY, USA: Oxford University Press.

Hogarth, R. and N. Karelaia (2007). Heuristic and linear models of judgment: Matching rules and environments. *Psychological review 114*(3), 733–758.

Hopp, W., S. Iravani, and G. Yuen (2007). Operations systems with discretionary task completion. *Management Science 53*(1), 61.

Institute of Medicine (2001). *Crossing the Quality Chasm: A New Health System for the 21st Century*. Institute of Medicine. National Academies Press.

Lippman, S. (1975). Applying a new device in the optimization of exponential queueing systems. *Operations Research 23*, 687–710.

Payne, J., J. Bettman, and E. Johnson (1988). Adaptive strategy selection in decision making. *Journal of experimental psychology. Learning, memory, and cognition 14*(3), 534–552.

Payne, J., J. Bettman, and E. Johnson (1993). *The Adaptive Decision Maker*. Cambridge University Press.

Payne, J., J. Bettman, and M. Luce (1996). When Time Is Money: Decision Behavior under Opportunity-Cost Time Pressure. *Organizational Behavior and Human Decision Processes 66*(2), 131–152.

Rieskamp, J. and U. Hoffrage (1999). When do people use simple heuristics, and how can we tell? In G. Gigerenzer, P. Todd, A. R. Group, and I. NetLibrary (Eds.), *Simple heuristics that make Us smart*, pp. 37–58. Oxford University Press New York.

Sennott, L. I. (1999). *Stochastic dynamic programming and the control of queueing systems*. New York, NY, USA: Wiley-Interscience.

Travers, D. (1999). Triage: How long does it take? How long should it take? *Journal of Emergency Nursing 25*(3), 238–240.

Wallsten, T., T. Pleskac, and C. Lejuez (2005). Modeling behavior in a clinically diagnostic sequential risk-taking task. *Psychological review 112*(4), 862–880.

Wang, X., L. Debo, A. Scheller-Wolf, and S. Smith (2008). Design and Analysis of Diagnostic Service Centers. Technical report, Working paper, Carnegie Mellon University, Pittsburgh, USA.

A Optimal Sequence of Cues

In this section, we prove that the optimal sequence of cues correspond to ordering β_k a monotonically decreasing function in k.

Suppose a sequence $\{\beta_k\}$ is not monotonically decreasing in k. Consider a random variable $\mathcal{K}$, which indicates the number of cues that are elicited before identifying a type τ customer, given that that the customer is of type τ. The cumulative probability distribution is equal to $P(\mathcal{K} \leq \kappa) = 1 - \prod_{k=0}^{\kappa-1}(1-\beta_k)$. Now consider $\{\beta_{(k)}\}$ an ordered sequence of $\{\beta_k\}$ such that $\beta_{(k)} \geq \beta_{(k+1)}$ for all k. Similarly, define for this new system $\hat{\mathcal{K}}$ the random number of elicited cues needed to identify a type τ customer. The cumulative distribution of $\hat{\mathcal{K}}$ is $P(\hat{\mathcal{K}} \leq \kappa) = 1 - \prod_{k=0}^{\kappa-1}(1-\beta_{(k)})$.

Since $\{\beta_{(k)}\}$ is decreasing in k, we have $P(\mathcal{K} \leq \kappa) \leq P(\hat{\mathcal{K}} \leq \kappa)$ for all k. In other words, $\mathcal{K}$ stochastically dominates $\hat{\mathcal{K}}$. In particular, we can then consider sample paths such that $\ell \geq \hat{\ell}$, in which ℓ and $\hat{\ell}$ are realizations of $\mathcal{K}$ and $\hat{\mathcal{K}}$, respectively.

Denote π to be the optimal policy following the system with sequence $\{\beta_k\}$. We now construct a new policy $\hat{\pi}$ for the system with ordered sequence $\{\beta_{(k)}\}$, and demonstrate that it is at least as good as π. Specifically, whenever policy π stops the search, we also let $\hat{\pi}$ stop the search. Whenever policy π continues eliciting cues according to the sample path in the $\{\beta_k\}$-system, so does policy $\hat{\pi}$ for the corresponding sample path generated in the $\{\beta_{(k)}\}$-system, even if a type τ customer has already been identified. This is possible because $\ell > \hat{\ell}$. Policy $\hat{\pi}$ generates then an average profit at least as good as the average profit corresponding to π. It follows that the optimal profit for the $\{\beta_{(k)}\}$ -system is at least as high as the optimal profit of the $\{\beta_k\}$-system.

B Optimality Equations

The set of actions is $A = \{0, 1, \ldots, x\}$, which represents the number of customers to be simultaneously removed from the system. Action $a = 0$ also states to keep the customer in service. Let $J(x,k)$ represent the value function. Following the uniformization procedure ($\lambda + \mu = 1$), the optimality equation for the long run average case becomes

$$g + J(x,k) = \max_{a \in A} Q(x,k,a) \ ,$$

in which

$$\begin{aligned}
Q(x,k,0) &= \pi(x,k) + \lambda J(x+1,k) + \mu\left(p_k\beta_k J(x-1,0) + (1-p_k\beta_k)J(x,k+1)\right) \ , \\
Q(x,k,a) &= \pi(x-a,0) - [(a-1)p_0 + p_k]C + \lambda J(x-a+1,0) \\
&\quad + \mu\left(p_0\beta_0 J(x-a-1,0) + (1-p_0\beta_0)J(x-a,1)\right), \text{ for } 1 \leq a < x, \\
Q(x,k,x) &= -[(x-1)p_0 + p_k]C + \lambda J(1,0) + \mu J(0,0), \text{ for } a = x.
\end{aligned}$$

Note that we have $Q(x,k,a)+p_kC = Q(x-1,0,a-1)$, which further implies that for any $k \geq 0$,

$$\begin{aligned} g + J(x-1,0) &= \max_{0 \le a \le x-1} Q(x-1,0,a) \\ &= \max_{0 \le a \le x-1} Q(x,k,a+1) + p_k C \\ &= \max_{1 \le a \le x} Q(x,k,a) + p_k C \end{aligned}$$

Hence the optimality equation is equivalent to

$$\begin{aligned} g + J(x,k) &= \max\{g + J(x-1,0) - p_k C, Q(x,k,0)\} \qquad \text{for } x > 0 \\ g + J(0,0) &= \lambda J(1,0) + \mu J(0,0) \quad \text{for } x = 0. \end{aligned}$$

C Characterization of the Optimality Policy

According to Equation 2, we can limit our analysis to systems with value $V + C$ and no mismatch cost C. Therefore, we assume in the following that $C = 0$.

C.1 Discounted Profit

To analyze the discounted profit case, we first uniformize the system and assume, without loss of generality, that $\lambda+\mu+\gamma = 1$. Given control policy u and initial state (x,k), the total discount profit is equal to $J^u(x,k) = \mathrm{E}_{x,k}\left[\int_0^{+\infty} V e^{-\gamma t} dN^u(t) - \int_0^{+\infty} c(X^u(t))\, e^{-\gamma t}\, dt\right]$.

The optimality equations are then equivalent to

$$J(x,k) = \Gamma J(x,k) \tag{5}$$

with

$$\Gamma J(x,k) = \begin{cases} \max\left\{J(x-1,0),\ TJ(x,k)\right\} & \text{for } \ x > 0\ ; \\ \lambda J(1,0)/(\lambda+\gamma)\ , & \text{for } \ x = k = 0\ . \end{cases} \tag{6}$$

where

$$TJ(x,k) = \pi(x,k) + \lambda J(x+1,k) + \mu\left(p_k \beta_k J(x-1,0) + (1 - p_k \beta_k) J(x,k+1)\right) \tag{7}$$

The next result shows the existence of an optimal value function $J^*(x,k)$. Define recursively $\Gamma^{(k)} J = \Gamma(\Gamma^{(k-1)} J)$ (with $\Gamma^{(1)} = \Gamma$) and the null function $J_0(x,k) = 0$ for all (x,k). We have,

Lemma 1 *There exists a unique optimal value function J^* such that $J^* = \Gamma J^*$ and*

$$\lim_{k \to \infty} \Gamma^{(k)} J_0 = J^*.$$

Proof: Consider the equivalent minimization problem where the instantaneous cost is given by $-\pi(x,k)$. Following Chapter 3.1 of Bertsekas (2007a),Vol. II, we can then add V to the $-\pi(x,k)$ so that the minimization problem satisfies the Positivity Assumption (see Bertsekas 2007a, Vol. II, equation (3.1)) since $V-\pi(x,k)\geq 0$. The result then follows directly from Proposition 3.1.6 in Bertsekas (2007a), Vol. II.

■

We can immediately deduce that the optimality function is non-negative and non-decreasing in the number of customers when no cues are elicited as shown by the following result. (In the following, notations $\searrow$ and $\nearrow$ indicate that the function is non-increasing and non-decreasing, respectively.)

Proposition 9 *$J^*(x,0)\nearrow x$ for all x and $J^*(x,k)\geq 0$ for all (x,k).*

Proof: From (5) we directly have $J^*(x+1,0)=\max\{J^*(x,0),\ TJ^*(x+1,0)\}\geq J^*(x,0)$. For the second part, consider a feasible but potentially suboptimal policy that never elicits any cue. Under this policy, the system is always empty and the profit is zero, with the assumption that $c(0)=0$.

■

Given the existence of $J^*(x,0)$ from Lemma 1, we can consider a revised instantaneous cost defined for $x>0$ as

$$\bar{c}(x)\equiv c(x)+\gamma J^*(x-1,0)$$

with $\bar{c}(0)=0$. Proposition 9 implies $\bar{c}(x)\nearrow x$. The following instantaneous profit is then well defined,

$$\bar{\pi}(x,k)\equiv \mu p_k\beta_k V-\bar{c}(x),$$

with $\bar{\pi}(x,k)\searrow x,k$.

For any real valued function $\Delta(x,k)$ on $x>0$ and $k\geq 0$, we consider operators Ψ and Υ such that

$$\Psi\Delta\equiv\left[\Upsilon\Delta(x,k)\right]^+ \ , \tag{8}$$

where

$$\Upsilon\Delta(x,k)\equiv\bar{\pi}(x,k)+\lambda\big(\Delta(x+1,k)+\Delta(x,0)\big)+\mu(1-p_k\beta_k)\Delta(x,k+1) \tag{9}$$

We will show later that function Δ can be interpreted as the marginal increase of J in x.

This new formulation allows us to show by iteration that the optimal value function is increasing and concave in the number of customers in the system. More precisely, define set $\mathcal{D}$ of functions $\Delta(x,k)$ for $x>0$ such that $\Delta\in\mathcal{D}$ if

C1 $\Delta\geq 0$,

C2 $\Delta \searrow x$,

Operator Ψ preserves conditions C1-C2 as stated by the following lemma,

Lemma 2 *$\Delta \in \mathcal{D}$ implies $\Psi\Delta \in \mathcal{D}$.*

Proof: Consider $\Delta \in \mathcal{D}$. $\Psi\Delta \geq 0$ by definition. For Condition C2, $\Upsilon\Delta \searrow x$ since $\bar{\pi}(x,k) \searrow x$ and $\Delta \searrow x$. It follows that $\Psi\Delta \searrow x$.

■

The next result justifies the interpretation of Δ^* solving (8) as the marginal difference in x of J^* solving (5).

Lemma 3 *There exists a unique function Δ^* such that*

$$\Delta^* = \Psi\Delta^* \ , \quad \textit{and} \tag{10}$$

$$\lim_{k\to\infty} \Psi^{(k)} J_0 = \Delta^* \ , \tag{11}$$

such that $\Delta^*(x,k)$ is non-negative and non-increasing in x.
Further, the optimal value functions J^* and Δ^* satisfy

$$J^*(x,k) = \Delta^*(x,k) + J^*(x-1,0) \ , \tag{12}$$

and an optimal policy stops searching in state (x,k) if and only if $\Upsilon\Delta^*(x,k) < 0$.

Proof: The proof of existence of Δ^* and Equation (11) is similar to the one of Lemma 1 by noting that the instantaneous cost $-\bar{\pi}(x,k)$ is bounded from below since $J^*(x,0) \geq 0$ from Proposition 9. We can then deduce from Lemma 2 that $\Delta^* \in \mathcal{D}$.
The second part of the theorem directly follows from the optimality equations (5) and the definition of Ψ.

■

We can now show that the optimal policy is of threshold type in k. First we show that for any x it is optimal to stop for all $k > k^{myo}(x)$, where $k^{myo}(x) \equiv \arg\min_k\{\pi(x,k) < 0\}$ is the myopic threshold. Define set $\mathcal{D}_1$ to be such that if $J \in \mathcal{D}_1$,

C3 $TJ(x,k) \leq J(x-1,0)$ for all $k \geq k^{myo}(x)$;

C4 $0 \leq J(x,0) \leq \Gamma J(x-1,0)$

C5 $J(x,k) \leq \Gamma J(x,k)$

Notice that $J_0(x,k) \equiv 0$ belongs to $\mathcal{D}_1$. We have,

Lemma 4 *If $J \in \mathcal{D}_1$, then $\Gamma J \in \mathcal{D}_1$.*

Proof: Consider state x, k such that $k > k^{myo}(x)$.

$$\begin{aligned}
T\Gamma J(x,k) &= \pi(x,k) + \lambda\Gamma J(x+1,k) + \mu\left(p_k\beta_k\Gamma J(x-1,0) + (1-p_k\beta_k)\Gamma J(x,k+1)\right) \\
&= \pi(x,k) + \lambda J(x,0) + \mu\left(p_k\beta_k\Gamma J(x-1,0) + (1-p_k\beta_k)J(x-1,0)\right) \\
&\leq \pi(x,k) + \lambda J(x,0) + \mu\Gamma J(x-1,0) \\
&\leq \pi(x,k) + (\lambda+\mu)\Gamma J(x-1,0) \\
&\leq \Gamma J(x-1,0)
\end{aligned}$$

where the second equality holds since $\Gamma J(x,k+1) = J(x-1,0)$ and $\Gamma J(x+1,k) = J(x,0)$ for $k > k^{myo}(x)$, the first inequality holds from C5 in $x-1$, the second one from C4 and the last one since $\pi(x,k) < 0$, $\lambda + \mu < 1$ and $\Gamma J(x-1,0) \geq 0$. Therefore we propagated C3.

Propagation of C4 and C5 follows from the monotonicity of Operator Γ, a standard result in Dynamic Programming.

■

We can deduce that,

Lemma 5 *$J^* \in \mathcal{D}_1$.*

Proof: This holds from $J_0 \in \mathcal{D}_1$ and by iterating on the value function.

■

Hence, from Lemma 5 there exists an upper bound on x, denoted by $\bar{X}$, such that it is optimal to stop eliciting cues for all states (x,k) with $x > \bar{X}$.

We show next that the optimal policy is of threshold type in k.

Lemma 6 *For all x, there exists a non-increasing threshold $\bar{k}(x)$ such that it is optimal to stop eliciting cues if and only if $k < \bar{k}(x)$.*

Proof:

We first show that the optimal policy is of threshold type in x, such that it is optimal to elicit a cue if and only if x is at or below a threshold $\bar{x}(k)$. From Lemma 3, $\Delta^*(x,k)$ is non increasing in x. Equation (10) implies then that it is optimal to elicit a cue if and only if x is at or below a threshold $\bar{x}(k)$, for all k. We now show by contradiction that $\bar{x}(k)$ is non-increasing. Assume then that $\bar{x}(k)$ is increasing in k. This means that there exists some k such that there exists a $\kappa > k$ with $\bar{x}(\kappa) > \bar{x}(k)$. In particular, define $k' \equiv \max\{\arg\max_{\kappa:\ \kappa > k} \bar{x}(\kappa)\}$. We have then $\bar{x}(k') > \bar{x}(k)$ and it is optimal to stop eliciting cues in state (x', k) with $x' \equiv \bar{x}(k') \leq \bar{X}$.

By definition of x' and k', it is optimal to elicit a cue in state (x', k') while it is optimal to stop in both states $(x'+1, k)$ and $(x', k+1)$. In particular we have $J^*(x', k'+1) = J^*(x'-1, 0)$ and $J^*(x'+1, k') = J^*(x', 0)$. It follows then that in state (x', k'),

$$\begin{aligned} TJ^*(x', k') &= \pi(x', k') + \lambda J^*(x'+1, k') + \mu\left(p_{k'}\beta_{k'} J^*(x'-1, 0) + (1 - p_{k'}\beta_{k'})J^*(x', k'+1)\right) \\ &= \pi(x', k') + \lambda J^*(x', 0) + \mu J^*(x'-1, 0) > J^*(x'-1, 0) \ . \end{aligned} \tag{13}$$

where the last inequality holds because $TJ^*(x', k') > J^*(x'-1, 0)$ since it is optimal to elicit a cue in state (x', k').

On the other hand we have,

$$\begin{aligned} TJ^*(x', k) &= \pi(x', k) + \lambda J^*(x'+1, k) + \mu(p_k\beta_k J^*(x'-1, 0) + (1 - p_k\beta_k)J^*(x', k+1)) \\ &\geq \pi(x', k) + \lambda J^*(x', 0) + \mu J^*(x'-1, 0) && (14) \\ &\geq \pi(x', k') + \lambda J^*(x', 0) + \mu J^*(x'-1, 0) && (15) \\ &= TJ^*(x', k') > J^*(x'-1, 0). && (16) \end{aligned}$$

where (14) holds since $J^*(x, k) \geq J^*(x-1, 0)$ for any state x, k from the optimality equation, (15) holds from Assumption 1 and (14) from Equation 13. This implies that it is optimal to elicit a cue in (x', k), which contradicts the previous assumption that it is optimal to stop in state (x', k). Hence $\bar{x}(k)$ is non-increasing in k and the result follows directly.

■

C.2 Proof of Theorem 1

To extend Lemma 6 to the average profit case, we consider the limit of discount optimal policies as the discount factor after uniformization $\nu = (\lambda+\mu)/(\lambda+\mu+\gamma)$ (Lippman 1975) approaches one from below (or γ approaches zero from above). Sennott (1999) presents three assumptions (SEN1-3) that guarantee the convergence of the discount model to the average cost model, as ν (resp. γ) approaches 1 (resp. 0) from above. We denote J^γ and Δ^γ to represent the optimal value function and marginal increase in x, respectively, when the discount factor is γ.

We first check that the optimal value function J^γ of a γ discounted problem satisfies the SEN assumptions for State $(0, 0)$ (Sennott 1999, Chapter 7.2). Note first that the degenerated policy which consists of deciding without searching is a $(0, 0)$ standard policy (Sennott 1999, Definition 7.5.2) which provides the null value function. (In particular, $(0, 0)$ is the only recurrent state under this policy.) This implies that (SEN1-2) hold for state $(0, 0)$ (Sennott 1999, Proposition 7.5.3).

For (SEN3), we need to show that for all $\mathbf{s} = (x, k)$, $J^\gamma(\mathbf{s}) - J^\gamma(0,0)$ is bounded from above. Using value iteration, we can easily show that $\Delta^*(x,k) \leq V$ for all x, k. We can then deduce that $J^\gamma(x,k) - J^\gamma(0,0)$ is uniformly bounded from above,

$$J^\gamma(x,k) - J^\gamma(0,0) = \Delta^\gamma(x,k) + \sum_{i=1}^{x-1} \Delta^\gamma(i,0) \leq (\bar{x}-1)V \ ,$$

with $\Delta^*(x,k) = 0$ for all $x > \bar{x}$. Therefore, (SEN3) holds. From Sennott (1999, Theorem 7.2.3), the optimal long run profit satisfies optimality equation, $\gamma J^\gamma(x-1,0)$, $\Upsilon\Delta^\gamma$ and $\Upsilon\Delta^\gamma$ converge and all results for the discount case of the previous section hold for the average case as well.

■

D Other Proofs

D.1 Proof of Proposition 2

As in Appendix C, we assume that $C = 0$ without loss of generality. If $\mu p_0 \beta_0 V \leq c(1)$, then from Proposition 1 $\hat{k}(1) = 0$ which, from Theorem 1, implies that $\hat{k}(x) = 0$ for all x. Assume now that $\mu p_0 \beta_0 V > c(1)$, consider the control policy with threshold $\hat{k}(1) = 1$ and $\hat{k}(x) = 0$ otherwise. The corresponding average profit is $(\mu p_0 \beta_0 V - c(1))/(\lambda + \mu)$ which is larger than not asking any question.

D.2 Proof of Propositions 3

The problem becomes the classical search problem (Bertsekas 2007a) with search cost $\pi(1,k)$ for which the system elicits the $(k+1)^{th}$ cue if and only if $c(1) \leq p_k \beta_k (V + C)$ (see for instance Bertsekas 2007a). Hence, the maximum number of elicited cues increases with β_k and p_0 (since p_k increases in p_0). Given a maximum number of elicited cues $\bar{k}$, accuracy is equal to $1 - \prod_{i=1}^{\bar{k}} (1 - \beta_i p_i)$, which is increasing in $\bar{k}$.

D.3 Proof of Propositions 4 and 5

For simplicity, assume that $\beta = \beta_k$ for all k. Consider a system where the optimal policy is characterized by $\bar{x} > 0$ such that $\bar{k}(x) = 1$ if and only if $x < \bar{x}$. Such systems exist from Proposition 4. Denote $L(\bar{x})$ to be the corresponding expected queue length and $q(\bar{x})$ the probability $Pr(X) = \bar{x}$ (where X is the number of customer in the queue as defined in Section 2). The system becomes a classical $M/M/1/\bar{x}$ with limited buffer size. In particular $L(\bar{x})$ and $q(\bar{x})$ are increasing and decreasing in $\bar{x}$, respectively (see, for example, Gross and

Harris 1998). The corresponding value function is then equal to

$$g(\beta, \bar{x}) = \lambda p_0 \beta V(1 - q(\bar{x})) - cL(\bar{x})$$

where $g(\beta, \bar{x})$ is supermodular in β and $\bar{x}$ since

$$\frac{\partial g}{\partial \beta} = \lambda p_0 V(1 - q(\bar{x})) > 0 \ ,$$

which increases in $\bar{x}$ since $q(\bar{x})$ is decreasing in $\bar{x}$. Hence the optimal $\bar{x}$ that maximizes g increases as β increases and so does $L(\bar{x})$. The proof for p_0 is similar.

D.4 Proof of Proposition 6

Denote the set of policies to be U. For each given policy $u \in U$, we have (A^u, L^u) to be the pair of accuracy and congestion measures. A particular pair (A, L) is *achievable* if it is the convex combination of some (A^u, L^u)'s, reflecting that it is achieved by a randomized policy. Formally, the efficient frontier is obtained from solving the following linear optimization problem with varying L.

$$e_{\beta,\rho,p_0}(L) = \max_{y \in \mathcal{R}^{|U|}} \sum_{u \in U} y_u A^u$$

subject to

$$\sum_u y_u L^u = L \ ; \quad y \geq 0 \ \text{ and} \sum_u y_u = 1 \ .$$

Linear programming sensitivity results (see, for example, Bertsimas and Tsitsiklis 1997, Theorem 5.1) imply that $e_{\beta,\rho,p_0}(L)$ is a concave function of L. Now consider a policy $\hat{u}$ that never let a customer go without positively identifying type τ. Obviously $A^{\hat{u}} = 1$ and $L^{\hat{u}} = \infty$. Therefore $e_{\beta,\rho,p_0}(L)$ is monotonically non-decreasing.

D.5 Proof of Proposition 7

Consider a control policy with a decreasing threshold $\hat{k}(x)$. Consider a realization of the arrival poisson process $A(t)$ which counts the total number of customers that have arrived up to time instant t. For Customer i, consider the realization of cue value according to the probabilities p_k for $k \leq \hat{k}(0)$ until type τ is identified or $\hat{k}(0)$. Denote by k_i the resulting number of elicited cues. This realization corresponds to the number of elicited cues for Customer i if she were alone in the system. For the previous realizations of the departure and cue value processes, define two systems S^a and S^b with elicitation rates μ_a and μ_b, respectively. For Customer i in system S^c, consider k_i elicitation times s_k^c, $c = a, b$ where $s_k^a \leq s_k^b$ for $k \in [1 \cdots, k_i]$.

Define now the departure processes $D^a(t)$ and $D^b(t)$ which count the total numbers of customers that have left the system up to time t for systems a and b, respectively. Consider then a realization of departure times t_n^a and t_n^b which indicate the time instants of the n^{th} customer departure in systems a and b. Further $x^a(t)$ and $x^b(t)$ denote the queue lengths at time t for system S^a and S^b, respectively. Finally, consider τ_{n+1}^a and τ_{n+1}^b, the time instants at which systems S^a and S^b starts serving customer $n+1$.

We set $t_0^a = t_0^b = 0$. We next show by iteration that

(P1) $t_n^a \leq t_n^b$ for all n (which is equivalent to $D^a(t) \geq D^b(t)$)

(P2) $x^a(\tau_n^a) \leq x^b(\tau_n^b)$

(P3) $x^a(t_n^a) \leq x^b(t_n^b)$

Assume the properties are true up to the n^{th} departure in system S^a. Let θ_{n+1} be the time instant at which the $(n+1)^{st}$ customer joins the system. Since $\tau_{n+1}^a = \theta_{n+1}$ when customer $n+1$ joins an empty system and $\tau_{n+1}^a = t_n^a$ otherwise, we deduce that $\tau_{n+1}^a = \max(t_n^a, \theta_{n+1})$. Similarly $\tau_{n+1}^b = \max(t_n^b, \theta_{n+1})$, where τ_{n+1}^b is the time instant at which System S^a starts serving customer $n+1$. It follows that $\tau_{n+1}^a \leq \tau_{n+1}^b$ and $D^a(\tau_{n+1}^a) = n \geq D^b(\tau_{n+1}^a)$.

Now, if $t_{n+1}^a \leq \tau_{n+1}^b$ then $t_{n+1}^a \leq t_{n+1}^b$ and (P1) holds. Otherwise $t_{n+1}^a \geq \tau_{n+1}^b$. In this case, $D^a(\tau_{n+1}^b) = n = D^b(\tau_{n+1}^b)$ and $x^a(\tau_{n+1}^b) = A(\tau_{n+1}^b) - D^a(\tau_{n+1}^b) = A(\tau_{n+1}^b) - D^b(\tau_{n+1}^b) = x^b(\tau_{n+1}^b)$. For $t \in [\tau_{n+1}^b\,,\, \min(t_{n+1}^a, t_{n+1}^b)]$ only arrivals occur and $x^a(t) = x^b(t)$ so that $\hat{k}(x^a(t)) = \hat{k}(x^b(t))$. Both systems elicit the same number of cues for customer $n+1$, but System S^a is faster since $s_k^a \leq s_k^b$ for $k \in [1 \cdots, k_{n+1}]$. It follows that $t_{n+1}^a \leq t_{n+1}^b$ which shows the first property. Further, $x^a(\tau_{n+1}^a) = A(\tau_{n+1}^a) - D^a(\tau_{n+1}^a) = A(\tau_{n+1}^a) - n \leq A(\tau_{n+1}^b) - n = A(\tau_{n+1}^b) - D^b(\tau_{n+1}^b) = x^b(\tau_{n+1}^b)$ since $\tau_{n+1}^a \leq \tau_{n+1}^b$ and $D^a(\tau_{n+1}^a) = n = D^b(\tau_{n+1}^b)$.

Since $A(t_n^a) \leq A(t_n^b)$ and $D^a(t_n^a) = D^b(t_n^b) = n$, we have $x^a(t_n^a) = A(t_n^a) - D^a(t_n^a) \leq A(t_n^b) - D^b(t_n^b) = x^b(t_n^b)$, which is (P3).

From (P1) we can deduce that $x^a(t) \leq x^b(t)$ for all t which means that congestion decreases in μ. For accuracy, (P3) along with the decreasing threshold $\hat{k}(x)$ implies that by the time that any customer n is forced to leave the systems a or b, it must be that more number of searches have been done to her in system a. In case she leaves system b due to a correct identification in one system according, the same must be true for system a as well following our construction of the Bernoulli process. It follows that accuracy increases in μ.

D.6 Proof of Proposition 8

From Propositions 2 the system elicit at least one cue if and only if $c(1) < \mu\beta_0 p_0 V$. Assume now the second inequality $\mu\beta_1 p_1 V \leq c(1)$ holds. The system elicits then exactly one cues for $x = 1$ from Proposition 1 with $k^{myo}(1) = 1$. Since the optimal thresholds decrease in x from Theorem 1, the system elicits at most one cue and 1^st Imp is optimal.

For the other direction, assume 1^{st}Imp is optimal. Equation (3) implies then that $J(x,k) = J(0,0) - p_k C$ and $J(0,0) - p_k C + g \geq Q(x,k,0)$ for $k \geq 1$. In particular for state $(1,1)$, we have,

$$J(0,0) - p_1 C + g \geq Q(1,1,0) = \pi(1,1) + \lambda J(1,0) + \mu J(0,0) - (\lambda + \mu(1 - \beta_1 p_1))p_1 C \quad (17)$$

where the inequality holds from $J(2,1) = J(0,0) - p_1 C$, $J(1,2) = J(0,0) - p_2 C \geq J(0,0) - p_1 C$, and $\lambda + \mu = 1$. Further, $J(0,0) + g = \lambda J(1,0) + \mu J(0,0)$ from (3). It follows that

$$0 \geq -\mu\beta_1 p_1 p_1 C \geq \pi(1,1) + J(0,0) \quad (18)$$

which yields the result.

E Algorithms

In this Appendix, we present the computational algorithms we used to evaluate the different policies introduced in the paper.

E.1 Policies

The optimal policy is evaluated using the Relative Value Iteration Algorithm for long run average problems (Bertsekas 2007b).

Under an IQ policy with at most κ cues for each customer, the system is equivalent to an M/G/1 queue (see Gross and Harris 1998). Denote random variable $\tilde{k}$ to represent the number of elicited cues that are needed to reach a decision for a given customer. We have

$$\begin{aligned} p(k) \equiv P(\tilde{k} = k) &= \prod_{i=0}^{k-2}(1 - p_i\beta_i)p_{k-1}\beta_{k-1} \quad \text{for } k < \kappa \\ p(\kappa) \equiv P(\tilde{k} = k) &= \prod_{i=0}^{\kappa-2}(1 - p_i\beta_i) \quad \text{otherwise.} \end{aligned}$$

(Here $\prod_{i=a}^{b}(1 - p_i\beta_i) = 1$ if $a > b$.) We deduce,

$$\begin{aligned} L_\kappa^{\text{IQ}} &= \bar{\rho} + \frac{\bar{\rho}^2 + \lambda^2\sigma_s^2}{2(1-\bar{\rho})} \\ A_\kappa^{\text{IQ}} &= 1 - \prod_{i=0}^{\kappa}(1 - \beta_i) \\ g_\kappa^{\text{IQ}} &= \lambda(V + C)p_0 A_\kappa^{\text{IQ}} - cL_\kappa^{\text{IQ}} - \lambda p_0 C \ . \end{aligned}$$

for linear waiting costs $c(x) = cx$, with $\bar{\rho} = \rho \mathrm{E}[\tilde{k}]$ and, using the variance decomposition formula,

$$\sigma_s^2 = \frac{1}{\mu^2}\left(\mathrm{E}[\tilde{k}] + \sum_{k=1}^{\kappa} p(\kappa)k^2 - \mathrm{E}[\tilde{k}]^2\right)$$

The algorithm to evaluate IQ consists then in finding κ^* which optimizes g_κ^{IQ}.

Similarly, under a 1^{st}Imp which allows a maximum of χ customers in the queue, the system is is equivalent to an M/M/1/χ queue (with buffer size χ).

$$\begin{aligned}
L_\chi^{1^{st}\mathrm{Imp}} &= \frac{\rho}{1-\rho} - \frac{\chi\rho^{\chi+1}+1}{1-\rho^{\chi+1}} \\
A_\chi^{1^{st}\mathrm{Imp}} &= \beta\frac{1-\rho^\chi}{1-\rho^{\chi+1}} \\
g_\chi^{1^{st}\mathrm{Imp}} &= \lambda p_0(V+C)A_\chi^{1^{st}\mathrm{Imp}} - cL_\chi^{1^{st}\mathrm{Imp}} - \lambda p_0 C\ .
\end{aligned}$$

for linear waiting costs $c(x) = cx$. The algorithm consists then in finding χ^* which optimizes $g_\chi^{1^{st}\mathrm{Imp}}$.

On the other hand, FixT does not lead itself to a well-established queueing system. Closed form solutions do not seem to exist. We therefore evaluate a policy with threshold (χ, κ) by iterating on the value function. We then identify the best thresholds by enumeration (more efficient algorithms might exist).

E.2 Efficient Frontier

In order to generate the efficient frontier for optimal policies, one can first solve the optimization problem with different $(V+C)/c$ values using the Relative Value Iteration Algorithm. After solving the optimization problem for each choice of $(V+C)/c$, record the optimal policy. Then evaluate A and L of this policy, which provides a point on the efficient frontier. The whole efficient frontier is obtained by connecting the points. The efficient frontier of the square policy is obtained in a similar manner.

The efficient frontier from either the 1^{st}Imp or IQ policies are easier to obtain, since a 1^{st}Imp (IQ) policy with any χ (κ) is on the efficient frontier. (We leave the proof of such a claim to the readers.) Therefore one just need to evaluate A and L of the policies one by one and connect the points.