Ergodic risk-sensitive admission control of Markovian multi-server queues with abandonment

NICOLA G. SASSONE* AND GUODONG PANG*

Abstract. We consider the ergodic risk-sensitive admission control problem for a Markovian multi-server queueing system with abandonment, where costs are incurred for server idleness, customer abandonment, and rejecting incoming arrivals. We first derive the Bellman optimality equation for this problem and show that a threshold policy—one that rejects incoming arrivals whenever the system-size exceeds a threshold—is optimal among all admissible control policies. We then propose a policy iteration algorithm to identify the optimal threshold, where we prove, under certain conditions on the problem's parameters, that the algorithm will terminate at the optimal threshold level. We also characterize the effect of risk-sensitivity on the optimal threshold, proving that this threshold monotonically decreases with respect to the sensitivity parameter and converges to the average-cost optimal threshold from below as the sensitivity parameter tends to zero.

1. Introduction

When designing customer service systems, it is important to maximize server utilization while also keeping customer congestion and abandonment to a minimum. Given this trade-off, a crucial question arises: is there an optimal way to choose how many people to let into the system?

One approach to address this problem is through admission control, where at any time a decision-maker can decide how many customers can enter the system by rejecting incoming arrivals. In settings where it is important to minimize customer abandonment, it may be advantageous to expend effort to reject customers rather than risk them joining and increasing the queue just to leave before service. On the other hand, rejecting incoming arrivals may also contribute to the total time that the system is idle. There has been a large body of works that address the admission control problem for queueing systems by formulating it as a Markov Decision Process (MDP) in order to find an optimal admission control policy. However, all of these works optimize with respect to an objective function that is risk-neutral (such as the finite-horizon, infinite-horizon discounted, and long run average cost criteria). These cost criteria do not account for higher order moments of the cumulative objective cost associated with the queueing process, which limits decision-makers from considering factors that contribute to large fluctuations around the average behavior in the queueing dynamics and associated costs (see [11, Sections 1 and 2]).

In this work, we consider the ergodic risk-sensitive admission control problem for an $M/M/N+M$ queueing system where there are costs associated with server idleness, customer abandonment, and rejecting incoming arrivals. For the objective function, we use the ergodic risk-sensitive cost criterion, which consists of an exponential utility function applied to the accrued cost in order to account for all moments of the queueing process and to penalize large fluctuations of customer abandonment as well as long idling periods. This cost criterion incorporates a risk-sensitivity parameter $\gamma > 0$, which allows for flexibility in weighing the role of higher order moments of the system (see Remark 2.1

*Department of Computational Applied Mathematics and Operations Research, George R. Brown School of Engineering and Computing, Rice University, Houston, TX 77005

E-mail addresses: ngs6@rice.edu, gdpang@rice.edu.

Date: November 6, 2025.

Key words and phrases. Ergodic risk-sensitive admission control, Markovian multi-server queues with abandonment, M/M/n+M queues, threshold policy.

for a more in depth discussion on this). It is known that as $\gamma \to 0$, the optimal value of ergodic risk-sensitive control problems converges to the optimal value of corresponding average cost problems. For this reason, we frequently refer to the average cost setting as the risk-neutral setting.

Because we do not assume an upper bound on the system's capacity size, the probability transition rates are unbounded, which prevents us from analyzing the problem by constructing a uniformizing chain. Instead, we frame our problem as the ergodic risk-sensitive control (ERSC) of a continuous-time Markov chain taking values in a countable state space (see [9, 29]). We begin by presenting the Bellman equation for our problem (which is multiplicative instead of additive), and prove that having a solution pair to the multiplicative Bellman equation guarantees the optimality of a threshold-type policy among all admissible control policies (see Theorem 2.1). We remark that this claim is not at all immediate, and the multiplicative structure in the cost functional poses significant challenges. Our methods greatly differ from those used in the average cost setting (where the authors proved the optimality of a threshold policy among all stationary policies; see [16, Theorem 3.1]).

Since it is unknown whether a solution pair of the multiplicative Bellman equation exists, we turn our attention to solving for an optimal threshold within the class of threshold policies, with the intention that this will enable us to construct a solution pair to the multiplicative Bellman equation, and hence compute and optimal threshold (see Remark 2.3 for a discussion on this and a summary of our approach in solving the multiplicative Bellman equation). As a result, we propose a policy iteration-type algorithm (see Algorithm 1) that computes the ERSC cost corresponding to a sequence of increasing threshold levels until a local minimum is found. Under the condition that the cost per abandoned customer is greater than a ratio involving the cost per rejected customer and the sensitivity parameter γ, we prove that the first local minimum of the sequence is a global minimum. Using the fact that the first local minimum is indeed a global minimum, we show that this can be used to construct a solution pair to the multiplicative Bellman equation, and thus show that the first local minimum corresponds to an optimal threshold policy within the entire class of admissible policies (see Theorem 2.2). This result, as well as the condition on the cost parameters described above, can both be seen as generalizations of those from the associated risk-neutral setting in [16], and our proofs involve establishing various analogous properties of the solution to multiplicative Bellman equation. In the case that a local minimum is not found in the sequence, we prove in Theorem 2.3 that the ERSC costs corresponding to the algorithm's iterates converge to the optimal cost. The multiplicative structure in the ERSC cost functional prevents us from using classical techniques in average cost control in proving this, and so we instead make use of the stability of the process to derive a suitable limit that solves the Bellman equation on the entirety of the state space.

We also address the following questions on how the risk-sensitivity relates with the admission control problem:

1. How does the sensitivity parameter affect the optimal threshold size?
2. How does the risk-sensitive optimal threshold compare with the risk-neutral optimal threshold?

We show that the optimal threshold is monotonically decreasing with respect to the sensitivity parameter, and that the optimal threshold from the risk-sensitive setting converges to the optimal threshold from the risk-neutral setting (from below) as the sensitivity parameter tends to 0 (see Proposition 2.1 and Theorem 2.4). Our results indicate an important managerial result for capacity-sizing in service systems: a decision-maker that is risk-sensitive should be more cautious towards abandonment by lowering the admission threshold (but never intentionally idling servers). The proof of threshold policy convergence involves embedding the set of all threshold policies within the space of binary sequences, and then applying a compactness argument.

We also comment that the approximating diffusion control problem that arises in the Halfin-Whitt many-server limit regime is a mathematically rich and practically relevant aspect of the risk-sensitive admission control problem. Due to the novelty of the associated Hamilton-Jacobi-Bellman equation

(which is multiplicative instead of additive and involves a Robin boundary condition), as well as the technical challenges in proving asymptotic optimality for ERSC problems (see [2, 3]), we plan to address the approximating diffusion control problem in a subsequent paper.

1.1. **Literature review.** Admission control in queueing systems is an area which has been extensively studied since it was first proposed by Naor [21]. Here, the author investigated the problem for an $M/M/1$ queue in which departing customers yield a fixed reward and there are costs associated with customer congestion and rejecting potential arrivals, and considered both the self-optimization problem (in which individual customers are maximizing their own profit) and the social-welfare optimization problem (in which the system is seeking to maximize the long run average expected profit). Knudsen [15] extended the model in [21] to multi-server systems which consider a more general running cost. For a comprehensive survey of some of the early works studying admission control without customer abandonment, we refer the reader to [25]. The admission control problem has also been investigated for systems with multiple customer classes [1, 18] and time-varying parameters [19, 33].

There have been a growing body of works addressing the admission control problem with an MDP analysis for queues with customer abandonment. Koçağa and Ward [16] investigated the admission control problem for a multi-server queue with abandonment when there are costs associated with idling servers, customer abandonment, and rejecting incoming arrivals. Working under the average cost criterion, the authors characterized the optimality of a threshold policy among all stationary (and potentially random) policies, and proposed an iterative algorithm which is shown to compute the optimal threshold level under the assumption that the cost per abandoned customer is greater than the cost per arriving customer. We remark that the model and results in [16] directly inspired the current paper. Ayhan [5] considered the admission control problem for a multi-server queue with abandonment and a maximum capacity, and proved the optimality of a threshold policy under the average reward criterion for the cases that rewards are either processed at the time of arrival or at service completion. Wu and Ayhan [31] recently investigated the admission control problem for a multiclass multi-server system with abandonment and a maximum buffer-size, and proved the optimality of a double-threshold policy (where customers of a class are only admitted if the number in the system falls within the pair of thresholds). Zayas-Cabán and Lewis [34] studied the admission control problem for a non-stationary multi-server queue with two customer classes and customer abandonment, and show the optimality of a threshold policy under the discounted cost criterion.

In the context of single-server queues with abandonment, we mention the following articles. In [8], the authors investigate a production-inventory control problem for a single server system with abandonment. Working under the discounted cost criterion, the authors show the existence of an optimal policy consisting of two thresholds (one for determining when production takes place, and the other for determining when orders should be admitted). The authors in [17] investigated a content delivery problem for a single server processing system where jobs may abandon while waiting to be served, and the authors prove the optimality of a threshold policy under the average cost criteria. In [7], both the self-optimization and social-welfare problem for a single server system with abandonment were investigated, with the surprising result that the individual threshold joining strategy is independent of the abandonment rate, and in particular it is equal to the threshold strategy in a queue without abandonment.

There are also a number of works which consider an asymptotic analysis of the admission control problem for queues with abandonment. Ward and Kumar [26] considered the admission control problem for a single server queue with abandonment in heavy traffic. Working under the discounted cost criterion, the authors explicitly solve for an optimal threshold policy in the approximating diffusion control problem and use this to construct an asymptotically optimal policy for the original system. Similarly, Ghosh and Weerasinghe [12] investigated the coupled service rate and admission control problem for a single server queue with abandonment in heavy traffic. Koçağa and Ward [16],

in addition to their MDP analysis of the admission control problem, solved the approximating diffusion control problem that arises in the Halfin-Whitt many-server regime (under the average cost criterion), and used the solution to construct an asymptotically optimal threshold policy. Weerasinghe and Mandelbaum [27] also studied a closely related problem to [16] but under the discounted cost criterion. Liu and Weerasinghe [20] investigated the admission control problem for a double-ended queue with two customer classes and abandonments, where they solved for an optimal threshold policy under the discounted cost criterion and proved its asymptotic optimality.

We remark that nearly all of the existing works on the admission control of queues optimize with respect to risk-neutral cost criteria. To our knowledge, [24] is the first (and only) work which investigated a version of the admission control problem with risk-sensitivity. Here, the authors considered the problem of admitting packet transmissions in a stochastic wireless network with a finite maximum buffer, where the system incurs a cost whenever it attempts to transmit a packet or loses packets due to buffer overflow, and receives a reward for successful packet transmission. The authors incorporated risk-sensitivity into the decision-making by using the ergodic risk-sensitive cost criterion. Because the system has a finite maximum buffer, the authors approached the problem by constructing a uniformizing chain, and proved that the optimal policy is of a threshold structure, increasing with the transmission cost parameter, and can be explicitly solved for.

1.2. **Organization of the paper.** We introduce the necessary notation in the next subsection. In Section 2, we introduce the risk-sensitive admission control problem and present our main results. In Section 3, we provide numerical implementations to illustrate our main results. In Section 4, we collect the proofs of our main results.

1.3. **Notation.** We use $(\Omega, \mathcal{F}, \mathbb{P})$ to denote the underlying abstract probability space with $\mathbb{E}$ as the associated expectation. The set of nonnegative real numbers (integers) is denoted by $\mathbb{R}_+$ ($\mathbb{Z}_+$), $\mathbb{N}$ stands for the set of natural numbers, and $\mathbb{1}_A(\cdot)$ denotes the indicator function corresponding to set A. The minimum (maximum) of two real numbers a and b is denoted by $a \wedge b$ ($a \vee b$), and $a^{\pm} = (\pm a) \vee 0$. We define $S \doteq \{0, 1, 2, \dots\}$, which we also refer to as the state space. For a function $W : S \to (0, \infty)$, we let $\mathcal{O}^+(W)$ denote the class of all strictly positive functions on S which are of order W, *i.e.,*

$$\mathcal{O}^+(W) = \left\{ f : S \to (0, \infty) \, : \limsup_{i \to \infty} \frac{f(i)}{W(i)} < \infty \right\}.$$

2. Model and Results

2.1. **Queue dynamics and the admission control problem.** We consider $M/M/N + M$ queues with admission control. Both the queue dynamics and sources of the running cost in our admission control problem are modeled in the same way as in the risk-neutral setting [16]. We first discuss the dynamics of the system when there is no admission control. We denote the arrival rate, service rate (for each individual server), and abandonment rate by λ, μ, and θ, respectively. The number of customers in the system at time t can be represented by the initial number of customers, plus the number who have arrived, subtracted by the number who have departed from service and the number who have abandoned. Letting $X(t)$ represent the number of customers in the system at time t and X_0 the initial number of customers in the system, we have that

$$X(t) = X_0 + P_a(\lambda t) - P_s\left(\mu \int_0^t (N \wedge X(s))\, ds\right) - P_r\left(\theta \int_0^t (X(s) - N)^+ \, ds\right). \tag{2.1}$$

where P_a, P_s and P_r are unit-rate Poisson processes that are mutually independent, for arrival, service and abandonment/reneging, respectively. We now describe an admission control policy and the dynamics of the controlled system. For any $t \geq 0$, let $U_t \in \{0, 1\}$ denote the decision at time t on whether to accept ($U_t = 0$) or reject ($U_t = 1$) an incoming arrival. Note that admission

control decisions are only implemented at the arrival times of customers. We define an admission control policy as a process $U = \{U_t\}_{t\geq 0}$ which is adapted to the filtration $\sigma\big(X_s, U_s : s \leq t\big)$ for all $t \geq 0$, augmented to include all null sets. A control policy U is Markov if $U_t = v(t, X(t))$ for some $v : \mathbb{R}_+ \times S \to \{0, 1\}$. U is referred to as a stationary policy if v is independent of t. We let $\mathfrak{U}$ denote the set of all admission control policies, and $\mathfrak{U}_{sm}$ denote the set of all stationary Markov control policies. Furthermore, $v \in \mathfrak{U}_{sm}$ is said to be a threshold policy that regulates at level $\ell \in S$ if

$$v(i) = \begin{cases} 0, \, i \leq \ell - 1, \\ 1, \, i \geq \ell \, . \end{cases} \tag{2.2}$$

In other words, ℓ is the threshold for the number of customers in the system for when the system will begin to reject arrivals. Let $\mathfrak{U}_{thresh}$ denote the set of all threshold control policies. Note that any stationary policy is equivalent to a threshold policy when the system starts empty. However, in our analysis we do not assume the system starts empty. Furthermore, it is worth noting that the controlled process corresponding to a threshold level ℓ acts as an $M/M/N/\ell + M$ queue.

Under a control policy U, note that $\int_0^t U_s \, dP_a(\lambda s)$ counts the number of rejected customers up to time t. As a result, the number of customers in the system under a control policy U is given by

$$X(t) = X_0 + P_a(\lambda t) - P_s\bigg(\mu \int_0^t (N \wedge X(s)) \, ds\bigg) - P_r\bigg(\theta \int_0^t (X(s) - N)^+ \, ds\bigg) - \int_0^t U_s \, P_a(\lambda s) \, . \tag{2.3}$$

From (2.1) and (2.3), we can characterize the controlled transition rate matrix. Let $q(i, j, u)$ denote the transition rate of $X(t)$ going from state i to state j under control $u \in \{0, 1\}$. Let $\overline{\mu}(n) = \mu(n \wedge N) + \theta(n - N)^+$ represent the state dependent death rate. Then for $n \in S$, we have that

$$q(n, m, 0) = \begin{cases} \lambda, & m = n + 1, \\ \overline{\mu}(n), & m = n - 1, \\ -(\lambda + \overline{\mu}(n)), & m = n, \\ 0, & \text{otherwise}, \end{cases} \qquad q(n, m, 1) = \begin{cases} 0, & m = n + 1, \\ \overline{\mu}(n), & m = n - 1, \\ -\overline{\mu}(n), & m = n, \\ 0, & \text{otherwise}. \end{cases} \tag{2.4}$$

We now discuss our running cost model. Let $h_I > 0$ represent the cost per unit of time for each server that is idling, $\kappa > 0$ the cost of rejecting an arriving customer, and $a > 0$ the cost of each abandoned customer. We define the instantaneous idling and abandonment cost if there are i customers in the system as

$$h(i) \doteq h_I(N - i)^+ + a\theta(i - N)^+ \, . \tag{2.5}$$

The running cost per unit of time is given by $c : S \times \{0, 1\} \to \mathbb{R}_+$, where

$$c(i, u) \doteq \begin{cases} h(i), & u = 0, \\ h(i) + \lambda\kappa, & u = 1 \, . \end{cases} \tag{2.6}$$

Given an admissible control policy U and an initial state $x_0 \in S$, the corresponding cost functional for the ergodic risk-sensitive control (ERSC) problem, with a sensitivity parameter $\gamma > 0$, is given by

$$\mathcal{E}^{\gamma}_{x_0}(U) \doteq \limsup_{T\to\infty} \frac{1}{\gamma T} \log \mathbb{E}^U_{x_0}\bigg[\exp\bigg(\gamma \int_0^T c(X(t), U(t)) \, dt\bigg)\bigg], \tag{2.7}$$

where $\mathbb{E}^U_{x_0}$ denotes the expectation when the underlying control policy is U and $X_0 = x_0$. The objective of the risk-sensitive admission control problem is to minimize (2.7) over all control policies and initial states, *i.e.,* to find the optimal value

$$\rho^{\gamma} \doteq \inf_{x_0 \in S} \inf_{U \in \mathfrak{U}} \mathcal{E}^{\gamma}_{x_0}(U) \tag{2.8}$$

and to characterize optimal control policies.

Remark 2.1. *In the long-run average cost formulation of the admission control problem (see* [16, Section 2]*), the cost functional corresponding to an admissible policy U and initial state x_0 is given by*

$$\mathcal{E}^{avg}_{x_0}(U) \doteq \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}^U_{x_0}\left[\int_0^T c(X(t),U(t))\,dt\right], \tag{2.9}$$

and the objective is to solve the following problem:

$$\rho_{avg} := \inf_{x_0\in S}\inf_{U\in\mathfrak{U}} \mathcal{E}^{avg}_{x_0}(U)\,. \tag{2.10}$$

We now discuss the difference between the two cost formulations (2.7) *and* (2.9)*. For small γ, applying Taylor's formula to the exponential in* (2.7)*, and then again to the natural logarithm, we have*

$$\begin{aligned}&\frac{1}{\gamma}\log\mathbb{E}^U_{x_0}\left[\exp\left(\gamma\int_0^T c(X(t),U(t))\,dt\right)\right]\\ &\quad=\mathbb{E}^U_{x_0}\left[\int_0^T c(X(t),U(t))\,dt\right]+\frac{\gamma}{2}\mathbb{E}^U_{x_0}\left[\left(\int_0^T c(X(t),U(t))\,dt\right)^2\right]+\mathcal{O}(\gamma^2)\,.\end{aligned} \tag{2.11}$$

This shows how the sensitivity parameter γ controls the weight of the higher order moments of the cumulative cost in the objective functional (2.7)*, and thus adds risk-sensitivity to the optimization problem in* (2.8)*. Moreover, taking $\gamma\to 0$ in* (2.11)*, one can see that the risk-sensitive cost functional* (2.7) *approaches the average cost functional* (2.9) *(see Theorem 2.4). Hence, we refer to the average cost formulation in* (2.9) *as "risk-neutral."*

We also comment that the choice of using the exponential in our cost formulation can be seen as using a certain type of utility function in the objective function (instead of just the expected value), which has origins in economics as a means of modeling individuals' risk-preferences and valuations in uncertain environments (see [32, Section 1] *for a more thorough discussion on this). We make the choice of using an exponential utility function since it is amenable to dynamic programming arguments (see the discussion in* [9, Section 1]*).*

2.2. **Solving for the optimal policy.** The starting point in our analysis is characterizing optimal policies. In particular, we provide a condition demonstrating how within the space of all admissible control policies, there exists a threshold policy which is optimal. This requires the following condition on the risk-sensitivity parameter.

Assumption 2.1. *Define $\overline{\gamma} \doteq \frac{1}{2a}$. The risk-sensitivity parameter γ satisfies*

$$\gamma < \overline{\gamma}\,.$$

Assumption 2.1 is important in the sense that it ensures the running cost from (2.6)—which is also unbounded—does not grow "too fast" with respect to the process $X(t)$. This can be viewed as an explicit characterization of the "small cost" condition, which is commonly employed in studying the ERSC problem for processes which take values in an unbounded space (see [9]). In particular, the small cost condition is used in conjunction with the uniform stability condition of the controlled process $\{X(t)\}$, which is shown to be satisfied in the following lemma. Note that we define a function $g : S \to \mathbb{R}$ to be *norm-like* if for any $\varepsilon > 0$, the level set $\{i \in S : g(i) \le \varepsilon\}$ is finite or empty.

Lemma 2.1. *There exist a norm-like and increasing function $\mathcal{V} : S \to [1,\infty)$, a constant $\hat{C}$, a finite set $\mathcal{K} \subseteq S$, and a norm-like function g, all independent of the sensitivity parameter γ, such that the following holds:*

$$\max_{u\in\{0,1\}}\sum_{j\in S}\mathcal{V}(j)q(i,j,u) \le \hat{C}\mathbb{1}_{\mathcal{K}}(i) - g(i)\mathcal{V}(i), \quad \forall i\in S\,. \tag{2.12}$$

Moreover, $g(i) - \gamma\max_{u\in\{0,1\}} c(i,u)$ is norm-like under Assumption 2.1.

Remark 2.2. *Lemma 2.1 takes the form of a Lyapunov stability condition, which is commonly employed in analyzing the ERSC problem for unbounded processes; see* [10, Assumption 3.4] *for the general form of this condition in the context of continuous-time Markov chains and* [9, Sections 2 and 3] *for the appropriate variants in the context of discrete-time Markov chains and diffusion processes. Furthermore, Lemma 2.1 ensures that the ERSC cost functional in (2.7) is uniformly bounded over all admissible control policies and initial states (see the proof of* [10, Lemma 3.10]*).*

Theorem 2.1. *Let Assumption 2.1 hold, and suppose there exists a pair* $(\rho^*, V^*) \in \mathbb{R}_+ \times \mathcal{O}^+\big(\mathcal{V}^\delta\big)$ *for some* $\delta \in (0,1)$*, which satisfies the Bellman equation:*

$$\gamma \rho^* V^*(i) = \min_{u \in \{0,1\}} \left[\gamma c(i,u) V^*(i) + \sum_{j \in S} V^*(j) q(i,j,u) \right], \quad \forall i \in S\,. \tag{2.13}$$

Then $\rho^* = \rho^\gamma$*, the optimal cost given in (2.8). Furthermore, there exists a threshold policy which is optimal among all admissible control policies* $\mathfrak{U}$*, and the optimal threshold corresponds to the first state in* S *where the minimum in (2.13) is attained at* $u = 1$*.*

Remark 2.3. *The primary reason we do not a priori have the existence of a solution pair to (2.13) stems from the structure of the admission control problem, since the implementation of admission control causes the state space to be truncated, which essentially means that the process is no longer globally irreducible on all of* S*. In particular, the Dirichlet eigenvalue approach given in* [10, Section 3] *(where one computes an eigenpair of the Bellman operator in (2.13) on an exhausting sequence of finite subsets of* S *and then takes a limit of the corresponding eigenpairs) does not appear to extend well to the admission control setting. The effect of exercising admission control in the Bellman operation causes the Dirichlet eigenfunctions to not be strictly positive (at some fixed point), which prevents one from taking a suitable limit of the pairs as the Dirichlet subsets increase to* S*. In* [13]*, the authors are able to prove the existence of a solution pair to the multiplicative Bellman equation without the global irreducibility condition, but they require a stronger Lyapunov stability condition than the one in Lemma 2.1 since they require that* $\mathcal{K} = \{i_0\}$ *for some* $i_0 \in S$*.*

Although the existence of a solution pair to (2.13) is not known a priori, we emphasize that Theorem 2.1 states that the Bellman optimality condition (which characterizes an optimal threshold) holds if there exists a solution pair (ρ^*, V^*) *which solve the multiplicative Bellman equation in (2.13). In Theorems 2.2 and 2.3, we provide conditions which allow us to explicitly construct a solution pair to the multiplicative Bellman equation. For that purpose, we rewrite the multiplicative Bellman equation in a form more amenable to the admission control problem and that is equivalent to Theorem 2.1 in Lemma 2.2.*

We remark that in the risk-neutral setting of this problem, one also does not a priori have of the existence of a solution to the average cost Bellman equation. Hence, the authors instead work with the iterates of the n-terminating problem, and provide conditions which allows for the construction of a solution pair to the additive Bellman equation and thus characterize an optimal threshold policy (see [16, Theorems 3.2 and 3.3]*).*

Remark 2.4. *In the risk-neutral setting* [16]*, the Bellman optimality condition takes the following form: there exists a pair* (ρ_{avg}, V_{avg})*, where* $\rho_{avg} \geq 0$ *and* $V_{avg} : S \to \mathbb{R}$*, which solves the additive Bellman equation*

$$\rho_{avg} = \min_{u \in \{0,1\}} \left[c(i,u) + \sum_{j \in S} V_{avg}(j) q(i,j,u) \right], \quad \forall i \in S\,. \tag{2.14}$$

From here, we see that the risk-sensitive Bellman equation in (2.13) is structurally different from the risk-neutral version, as the risk-sensitive value function V^* *is multiplied to both the running cost and optimal value terms (in addition to appearing in the generator term). In comparison, the risk-neutral value function* V_{avg} *only appears in the generator term. Furthermore, it is clear that*

multiplying V^* by a positive constant does not affect (2.13), while adding a constant to V_{avg} does not affect (2.14). It is for these reasons that the risk-sensitive Bellman equation is commonly referred to as the "multiplicative Bellman equation," while the risk-neutral Bellman equation is referred to as the "additive Bellman equation." We additionally comment that the multiplicative form of (2.13) allows for the interpretation of V^* as a strictly positive eigenvector of the Bellman operator on the right side, and ρ^* as its corresponding eigenvalue (we refer the reader to [9] for an overview of exploiting this eigenvalue interpretation as a means to prove the existence of this pair).

Finally, since the next steps in our approach involve analyzing the multiplicative Bellman equation (2.13), the subsequent results (and their proofs) tend to take a different form than those from the risk-neutral setting, with some notable comparisons (see Remarks 2.5 and 2.7).

The structure of the admission control problem as well as the multiplicative nature of the risk-sensitive Bellman equation encourages us to rewrite (2.13) in a more tractable form, which will greatly aid us in the rest of this section. The following calculations can be seen as the multiplicative analogue to those in [16, Pg. 280].

For $i \in S$, let $y^*(i) \doteq \frac{V^*(i)}{V^*(i-1)}$ (which is possible since $V^* > 0$). Substituting the running cost and transition rates from (2.6) and (2.4) into (2.13) results in

$$\gamma\rho^* V^*(i) = \min\left\{ \begin{array}{c} V^*(i)(\gamma h(i) - \lambda - \overline{\mu}(i)) + \lambda V^*(i+1) + \overline{\mu}(i)V^*(i-1), \\ V^*(i)(\gamma h(i) + \gamma\lambda\kappa - \overline{\mu}(i)) + \overline{\mu}(i)V^*(i-1) \end{array} \right\}, \quad i \geq 1\,.$$

Dividing by $V^*(i)$ on both sides then implies

$$\begin{aligned} \gamma\rho^* &= \min\left\{ \begin{array}{c} \gamma h(i) - \lambda - \overline{\mu}(i) + \lambda\frac{V^*(i+1)}{V^*(i)} + \overline{\mu}(i)\frac{V^*(i-1)}{V^*(i)}, \\ \gamma h(i) + \gamma\lambda\kappa - \overline{\mu}(i) + \overline{\mu}(i)\frac{V^*(i-1)}{V^*(i)} \end{array} \right\} \\ &= \gamma h(i) - \overline{\mu}(i) + \overline{\mu}(i)\frac{V^*(i-1)}{V^*(i)} + \lambda\min\left\{\frac{V^*(i+1)}{V^*(i)} - 1, \gamma\kappa\right\} \\ &= \gamma h(i) - \overline{\mu}(i) + \frac{\overline{\mu}(i)}{y^*(i)} + \lambda\min\left\{y^*(i+1) - 1, \gamma\kappa\right\}. \end{aligned}$$

Furthermore, the last equation reveals that the optimal admission threshold n^* is the first n such that $y^*(n+1) \geq \gamma\kappa + 1$. For $i = 0$, a similar analysis gives us

$$\gamma\rho^* = \gamma h(0) + \lambda\min\left(y^*(1) - 1, \gamma\kappa\right).$$

A reversal of these arguments shows how if we are given (ρ^*, y^*), then by setting $V^*(0) = 1$, we can clearly reconstruct the original value function V^* so that (2.13) is satisfied. Moreover, if y^* is assumed to be uniformly bounded on S (by some $\eta > 0$), then it follows that $V^*(i) \leq \eta^i$, and as a consequence $V \in \mathcal{O}^+(\mathcal{V}^\delta)$ for some $\delta < 1$ (see Remark 4.1). We summarize these calculations and findings in the following lemma.

Lemma 2.2. *The Bellman equation in (2.13) can be equivalently written as a pair (ρ^*, y^*), where $\rho^* \geq 0$ and $y^* : S \to (0, \infty)$ is uniformly bounded on S, such that this pair satisfies*

$$\gamma\rho^* = \gamma h(i) - \overline{\mu}(i) + \frac{\overline{\mu}(i)}{y^*(i)} + \lambda\min\left(y^*(i+1) - 1, \gamma\kappa\right), \quad \forall i \geq 1, \tag{2.15}$$

$$\gamma\rho^* = \gamma h(0) + \lambda\min\left(y^*(1) - 1, \gamma\kappa\right). \tag{2.16}$$

Furthermore, an optimal threshold n^* can be characterized by

$$n^* = \min\left\{n \geq 0 : y^*(n+1) \geq 1 + \gamma\kappa\right\}. \tag{2.17}$$

Remark 2.5. *In the risk-neutral setting of this problem [16], the authors defined $y_{avg}(i) \doteq V_{avg}(i) - V_{avg}(i-1)$, which allowed them to equivalently write the additive Bellman equation as*

$$\rho_{avg} = h(i) - \overline{\mu}(i)y_{avg}(i) + \lambda\min\left(y_{avg}(i+1), \kappa\right).$$

From here, it follows that the risk-neutral optimal threshold n_{avg} is characterized by

$$n_{avg} = \min\{n \geq 0 \,:\, y_{avg}(n+1) \geq \kappa\},$$

which was given in the paragraph immediately following [16, Theorem 3.1]. *Thus, for the risk-sensitive setting, the condition in* (2.17) *can be seen as a multiplicative analogue.*

Theorem 2.1 and Lemma 2.2 show that the optimality of a threshold policy is characterized by the eigenpair (ρ^*, y^*) satisfying the multiplicative Bellman equation. However, these are primarily necessary conditions, as the existence of a pair (ρ^*, y^*) is not readily known to us, and so it is unclear how to compute an optimal threshold policy (if it even exists).

We approach this issue by proposing a policy iteration-type algorithm. This can be seen as a risk-sensitive extension of the n-terminating problem from the risk-neutral setting as described in [16, Sections 3.2 and 3.3]. This is described in the next paragraph.

Consider fixing a threshold policy with a threshold level of n, that is, a policy that only allows a maximum of n people in the system. Under this threshold, the state space of the controlled process is $\{0, 1, \ldots, n\}$, and it is straightforward to verify that the controlled process is ergodic on $\{0, 1, \ldots, n\}$. As such, we can view the system under a threshold level of n as a finite state ERSC MDP (with the same running cost as Section 2.1). By [28, Theorem 3.2], it can be shown that there exists a pair (ρ_n, V_n) with $\rho_n \geq 0$ and $V_n : \{0, \ldots, n\} \to \mathbb{R}_+$, which satisfies

$$\gamma\rho_n V_n(i) = \gamma V_n(i)h(i) + \sum_{j\in S} V_n(j)q(i,j,0), \quad i \leq n-1, \tag{2.18}$$

$$\gamma\rho_n V_n(n) = \gamma V_n(n)\big(h(n) + \lambda\kappa\big) + \sum_{j\in S} V_n(j)q(n,j,1)\,. \tag{2.19}$$

Furthermore, ρ_n is the ERSC cost corresponding to a threshold level of n. Since the above is a finite-state MDP, it is straightforward to see how (ρ_n, V_n) can be computed, such as by uniformizing the chain and then applying risk-sensitive relative value iteration (see [23, Ch. 11.4] for a background on uniformization for continuous-time MDPs and [24, Section 3] for an example of this in the ERSC problem). Setting

$$y_n(i) \doteq \begin{cases} \frac{V_n(i)}{V_n(i-1)}, & 1 \leq i \leq n, \\ \gamma\kappa + 1, & i = n+1, \end{cases}$$

one can apply the same analysis as in the proof of Lemma 2.2 to show that the pair (ρ_n, y_n) satisfies the following set of equations:

$$\gamma\rho_n = \gamma h(i) - \overline{\mu}(i) + \overline{\mu}(i)\frac{1}{y_n(i)} + \lambda y_n(i+1) - \lambda, \quad i = 1, \ldots, n, \tag{2.20}$$

$$\gamma\rho_n = \gamma h(0) + \lambda y_n(1) - \lambda, \tag{2.21}$$

$$y_n(n+1) = \gamma\kappa + 1\,. \tag{2.22}$$

Remark 2.6. *Since a threshold at n blocks arrivals from happening at state n, we emphasize that V_n is not defined at $n+1$. This, in conjunction with the fact that there is no minimum operation in* (2.20)-(2.21)*, means that we do not know whether this pair satisfies the minimum operation in the multiplicative Bellman equation.*

We are now in a position to present the algorithm for computing an optimal threshold.

The idea behind Algorithm 1 is that we are searching for a local minimum of the sequence $\{\rho_n\}_{n\geq 0}$. The reason that we are starting at N instead of 0 is because it is never optimal to idle a server, *i.e.*, $\rho_0 > \rho_1 > \cdots > \rho_N$ (see Lemma 4.5 in Section 4.3).

In Theorem 2.2, we show that $a \geq \frac{\kappa}{\gamma\kappa+1}$ is a sufficient condition to guarantee that the first local minimum produced by Algorithm 1 is in fact a global minimum. Furthermore, we are able to construct a solution to the multiplicative Bellman equation in Lemma 2.2 to show that controlling

Algorithm 1: PI Algorithm for risk-sensitive admission control

1. Initialize $n = N$.
2. Compute (ρ_n, y_n) and (ρ_{n+1}, y_{n+1}) as the solutions to (2.20)-(2.22).
3. If $\rho_n \leq \rho_{n+1}$, terminate and return n as the desired threshold. Otherwise, set $n \longleftarrow n+1$ and return to **2.**

with this threshold is optimal among all admissible control policies. As mentioned earlier, this result can be seen as a generalization of [16, Theorem 3.2], however recall that the authors proved optimality within the class of all stationary policies (including random policies).

Theorem 2.2. *Suppose $a \geq \frac{\kappa}{\gamma\kappa+1}$, and let ρ_m denote the first local minimum of the sequence $\{\rho_n\}_{n=0}^{\infty}$,* i.e., *$\rho_m$ satisfies $\rho_N > \rho_{N+1} > \cdots > \rho_m$ and $\rho_m \leq \rho_{m+1}$. Then $\rho_m = \rho^\gamma$, the optimal cost. Moreover, the threshold policy which regulates at level m is optimal among all admissible control policies.*

Remark 2.7. *We note that Theorem 2.2 can also be interpreted as a convexity result on the cost iterates $\{\rho_n\}_{n\geq 0}$, since the first local minimum corresponds to a global minimum. The proof of this result is done by method of contradiction, where if we assume ρ_m is not a global minimum, then we can construct an auxiliary ERSC problem and use the condition $a \geq \frac{\kappa}{\gamma\kappa+1}$ to force a contradiction on the assertion that ρ_m is not a global minimum (see Section 4.3).*

In the risk-neutral setting, the condition for guaranteeing optimality of a local minimum is $a > \kappa$ (see [16, Theorem 3.2]*). Since $\gamma > 0$, the condition $a > \kappa$ is actually stronger than the condition present in the above theorem, demonstrating that our condition is a natural generalization of the one present in the risk-neutral setting. Additionally, combining the condition $a \geq \frac{\kappa}{\gamma\kappa+1}$ with our assumption $a\gamma < \frac{1}{2}$ from Assumption 2.1, it necessarily follows that $a \geq \frac{\kappa}{2}$.*

Note that if the sufficient condition above is not satisfied, then our results cannot answer whether a local minimum of $\{\rho_n\}_{n\geq 0}$ is in fact a global minimum. This is also the case in the risk-neutral setting. Furthermore, even if it is given that we have a global minimum at ρ_m, it is not clear if it is possible to construct a solution to the optimality equations (2.15) without the condition $a \geq \frac{\kappa}{\gamma\kappa+1}$, and so it is not clear whether controlling with a threshold m is optimal among all admissible control policies.

While Theorem 2.2 guarantees that the first local minimum in $\{\rho_n\}_{n\geq 0}$ corresponds to an optimal threshold, it is possible that a local minimum may not exist. In this case, we prove that the sequence converges to the optimal cost.

Theorem 2.3. *Suppose $\{\rho_n\}_{n\geq 0}$ is strictly decreasing. Then we have $\rho_n \to \rho^\gamma$ as $n \to \infty$. Moreover, the policy which never exercises admission control is optimal among all admissible control policies.*

Remark 2.8. *As a consequence of Theorems 2.1, 2.2, and 2.3, under the condition $a \geq \frac{\kappa}{\gamma\kappa+1}$, we have the existence of an optimal threshold policy among all admissible control policies $\mathfrak{U}$. If this condition does not hold, then there still exists a threshold policy which is optimal within the smaller class of policies $\mathfrak{U}_{thresh}$, since there must exist a global minimum within the set $\{\rho_n\}_{n\geq 0} \cup \{\rho_\infty\}$ (a similar result holds in the risk-neutral setting). However, it is unclear whether we can characterize/compute such a threshold using the methods in this work. Additionally, it is unclear whether this threshold is optimal over all admissible policies $\mathfrak{U}$.*

Remark 2.9. *Since it is unknown whether Algorithm 1 will terminate in finite time (as $\{\rho_n\}_{n\geq 0}$ may or may not have a local minimum), we provide a few guidelines on when to stop Algorithm 1. Recall that n^* denotes the (potentially infinite) optimal threshold corresponding to sensitivity γ (which is fixed).*

(i) *In Section 2.3, Theorem 2.4(ii) states that $n^* \leq n_{avg}$, where n_{avg} denotes the risk-neutral optimal threshold. As a result, in the case that n_{avg} is already known to the user (and finite), we have that Algorithm 1 must terminate by the n_{avg}th iterate.*

(ii) *In the case that $a \geq \frac{\kappa}{\gamma\kappa+1}$, our numerical experiments so far indicate that before a local minimum is reached, the error rates $\varepsilon_n \doteq \rho_n - \rho_{n+1}$ are strictly decreasing in n. As a result, one may initialize Algorithm 1 with a tolerance of $\varepsilon > 0$, and terminate whenever $\varepsilon_n \leq \varepsilon$ to return a threshold which is "nearly optimal."*

(iii) *In the case that $a < \frac{\kappa}{\gamma\kappa+1}$, it is not guaranteed whether the first local minimum in $\{\rho_n\}_{n\geq 0}$ is optimal (see Remark (2.7)). However, our numerical experiments so far indicate that $\{\rho_n\}_{n\geq 0}$ is decreasing in this case. As a result, for this case we conjecture that the optimal policy is the one which never exercises admission control. We hope to investigate this conjecture more carefully in a follow-up paper by studying the approximating risk-sensitive diffusion control problem which arises in the Halfin-Whitt many-server regime. For now, we suggest using a large (or infinite) threshold in the case that $a < \frac{\kappa}{\gamma\kappa+1}$.*

We also note that the issue in (iii) similarly arises in the risk-neutral setting (but under the condition $a \leq \kappa$). The authors address this by investigating the approximation diffusion control problem, where they prove that if $a \leq \kappa$ then the policy which never exercises admission control is asymptotically optimal (see [16, Theorems 4.1 and 5.2]*).*

2.3. **Sensitivity properties of the optimal policy.** In this section, we address the question of how the risk-sensitivity affects the decision-making in the context of the admission control problem. More precisely, we seek to explicitly characterize the effect that the sensitivity parameter $\gamma > 0$ has on an optimal threshold policy, and how an optimal threshold from the risk-sensitive setting compares with the optimal threshold from the risk-neutral setting in [16]. In order to ensure that there is always an optimal threshold policy in both the risk-sensitive and risk-neutral setting, we assume that $a > \kappa$. However, we remark that the results of this section also hold if $a < \frac{\kappa}{\gamma\kappa+1}$, but the optimal threshold is only optimal within the class of $\mathfrak{U}_{thresh}$ (see Remark 2.8).

We address these questions by first proving that the risk-sensitive optimal thresholds satisfy a certain monotonicity property with respect to the sensitivity parameter. Recall from (2.8) that ρ^γ denotes the optimal ERSC cost corresponding to the sensitivity parameter γ, and we let n^γ represent the corresponding (potentially infinite) optimal threshold (for n^γ to be uniquely defined, we denote it as the smallest optimal threshold which attains the optimal cost ρ^γ).

Proposition 2.1. *Let $\gamma_1, \gamma_2 \in (0, \overline{\gamma})$ be two sensitivity parameters such that $\gamma_1 > \gamma_2$. It follows that $\rho^{\gamma_1} \geq \rho^{\gamma_2}$ and $n^{\gamma_1} \leq n^{\gamma_2}$.*

Remark 2.10. *Proposition 2.1 can be summarized as follows: the higher the risk-sensitivity parameter, the smaller the optimal threshold. One potential explanation for this is that a higher sensitivity parameter causes the ERSC cost functional in (2.7) to penalize larger fluctuations about the mean, and customer abandonment can cause the queue size to fluctuate considerably. As a result, optimizing with respect to the ERSC cost criteria results in a smaller optimal capacity size in an effort to reduce the fluctuations in the cost driven by customer abandonment.*

Recall from Remark 2.1 that ρ_{avg} is defined to be the optimal cost in the risk-neutral admission control problem. Also, let n_{avg} denote the (possibly infinite) optimal threshold (which we assume to be minimal as we did for n^γ).

We prove the following convergence results as the risk sensitivity parameter $\gamma \downarrow 0$.

Theorem 2.4. *The following hold: as $\gamma \downarrow 0$ (i) $\rho^\gamma \downarrow \rho_{avg}$ and (ii) $n^\gamma \uparrow n_{avg}$.*

Remark 2.11. *The convergence of the risk-sensitive optimal cost to the risk-neutral optimal cost is proven for discrete-time Markov chains taking values in a finite state space (see* [6, Theorem 3.1]

and the references therein). However, this does not immediately apply to our setting for proving Theorem 2.4(ii), since our controlled process may not necessarily take values in a finite space.

3. Numerical Implementations

In this section, we provide implementations for computing an optimal threshold policy through Algorithm 1, and we compare the optimal threshold sizes for the risk-sensitive and risk-neutral case. Let n^* denote an optimal threshold produced by the algorithm, ρ^* the optimal cost, and n_{avg} the optimal threshold for the risk-neutral case. In particular, we are interested in understanding the optimal threshold level across various levels of γ when the system is critically loaded and staffed appropriately. In other words, we consider cases in the Halfin-Whitt regime when $N \approx \frac{\lambda}{\mu} + \beta\sqrt{\frac{\lambda}{\mu}}$. Here, β is the quality of service parameter and is connected to the probability of delay (the probability that an arriving customer has to wait before service; see [30]).

In the following implementation, we use $\lambda = 50, \mu = 1$, and $\theta = 1$. By using $\beta = 1$, we have $P(Delay) \approx 0.7$, and the square-root staffing rule suggests that N should be approximately 57. To get a range on the different optimal thresholds when the system is slightly (over/under)loaded, we let N range between 45, 60, and 75. Table 1 compares the optimal cost and threshold for varying N and risk-sensitivity parameter $\gamma > 0$.

N	γ	n^*	ρ^*	n_{avg}	ρ_{avg}
45	1×10^{-2}	45	32.957	52	11.733
60	1×10^{-2}	60	47.479	94	13.422
75	1×10^{-2}	75	68.207	114	31.253
45	1×10^{-4}	51	12.010	52	11.733
60	1×10^{-4}	73	14.078	94	13.422
75	1×10^{-4}	75	32.221	114	31.253
45	1×10^{-6}	52	11.736	52	11.733
60	1×10^{-6}	82	13.429	94	13.422
75	1×10^{-6}	85	31.264	114	31.253
45	1×10^{-8}	52	11.733	52	11.733
60	1×10^{-8}	89	13.422	94	13.422
75	1×10^{-8}	93	31.253	114	31.253

Table 1. Comparison of optimal costs and threshold values for varying N and γ ($a = 2$, $\theta = 1$, $\kappa = h_I = 1.25$, $\mu = 1$, $\lambda = 50$.

The first notable result is that when the sensitivity parameter γ is large, the system is highly sensitive to the cost associated with customer abandonment. For $\gamma = 0.01$, the optimal threshold equals the number of servers because the system does not want to risk any customers abandoning. As the sensitivity decreases, we see an increase in the optimal system capacity past the number of servers. The optimal threshold continues to increase towards the risk-neutral optimal threshold from below, verifying Theorem 2.4. The numerical figures also verify that the risk-sensitive costs converge to the risk-neutral costs as $\gamma \to 0$.

The following plot demonstrates how the optimal cost changes with γ, for the parameters above with $N = 45, 60, 75$. Figure 1 shows that the optimal cost ρ is increasing with γ but tapers off as γ grows sufficiently large. It also shows how the rate that the risk-sensitive optimal cost decreases (as $\gamma \to 0$) varies, depending on the parameters and number of servers.

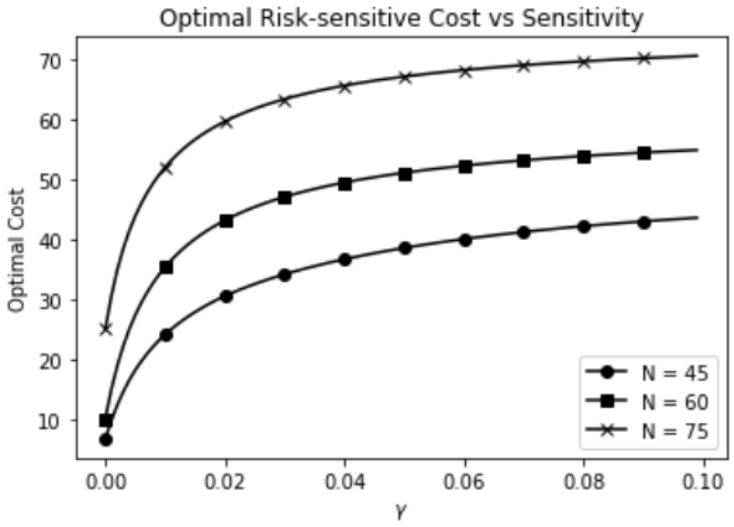

FIGURE 1. Plot of the optimal cost ρ as a function of γ ($a = 2$, $\theta = 0.1$, $\kappa = h_I = \mu = 1$, $\lambda = 50$)

We also investigate when the risk-sensitive problem has a finite optimal threshold, but the risk-neutral setting has a potentially infinite optimal threshold. In particular, consider the case when $a < \kappa$. For the risk-neutral setting, this means that we no longer know whether a local minimum in the cost iterates is optimal, but we know that the policy which never exercises admission control is asymptotically optimal when the queue operates in the Halfin-Whitt regime (see [16, Theorem 5.2]). Our numerical experiments show that this correlates with the risk-neutral cost iterates decreasing as the threshold size increases. In the risk-sensitive setting, we can choose the sensitivity parameter γ so that $a \geq \frac{\kappa}{\gamma\kappa+1}$, which meets the criteria of Theorem 2.2 so that a produced local minimum is optimal.

N	γ	n^*	ρ^*	n_{avg}	ρ_{avg}
45	1×10^{-2}	45	60.702	> 500	9.950
60	1×10^{-2}	60	52.012	> 500	1.012
75	1×10^{-2}	75	44.511	> 500	2.500
45	5×10^{-3}	48	47.875	> 500	9.950
60	5×10^{-3}	63	35.465	> 500	1.012
75	5×10^{-3}	79	24.514	> 500	2.500
45	2.6×10^{-3}	50	36.322	> 500	9.950
60	2.6×10^{-3}	67	19.781	> 500	1.012
75	2.6×10^{-3}	84	4.800	> 500	2.500

TABLE 2. Comparison of optimal costs and threshold levels for varying N when the risk-neutral setting has an infinite threshold ($a = 1.99, \kappa = 2$, $\theta = 0.01$, $h_I = 0.1$, $\mu = 1$, $\lambda = 50$).

For the parameter specifications in Table 2, we set $\gamma \geq 2.6 \times 10^{-3}$ to ensure that $a \geq \frac{\kappa}{\gamma\kappa+1}$. For larger γ, we see that the threshold is equal to the number of servers, as was the case in Table 1. However, as γ decreases, we see that the risk-sensitive optimal costs slowly decrease towards the risk-neutral costs and the corresponding risk-sensitive optimal thresholds increases. In addition to

numerically verifying Proposition 2.1 and Theorem 2.4, this reveals that it is indeed possible to have a risk-sensitive optimal threshold that is finite, but a risk-neutral optimal threshold that is potentially infinite.

4. Proofs

In this section, we provide the proofs for the results in Sections 2.2 and 2.3.

4.1. **Proof of Lemma 2.1.**

Proof. Let $\mathcal{V}(i) = e^{\alpha i}$ for some $\alpha > 0$, which we will specify at the end of the proof. Then for $i \geq 1$, observe that

$$\sum_{j\in S} \mathcal{V}(j)q(i,j,0) = \mathcal{V}(i)\Big[\lambda e^{\alpha} + \overline{\mu}(i)e^{-\alpha} - \big(\lambda + \overline{\mu}(i)\big)\Big].$$

Let $i \geq N$ so that $\overline{\mu}(i) = \mu N + \theta(i-N)$. Observe that

$$\begin{aligned}\sum_{j\in S} \mathcal{V}(j)q(i,j,0) &= \mathcal{V}(i)\Big[\lambda e^{\alpha} + e^{-\alpha}\big(\mu N + \theta i - \theta N\big) - \lambda - \mu N - \theta i + \theta N\Big] \\ &= i\mathcal{V}(i)\Big[\theta(e^{-\alpha}-1) + \frac{1}{i}\big(\lambda e^{\alpha} + e^{-\alpha}\mu N - e^{-\alpha}\theta N - \lambda - \mu N + \theta N\big)\Big]. \qquad (4.1)\end{aligned}$$

Clearly $e^{-\alpha} < 1$ for any $\alpha > 0$. Now, let $\beta := \frac{1}{2}\theta(1-e^{-\alpha}) > 0$. By (4.1), we have

$$\sum_{j\in S} \mathcal{V}(j)q(i,j,0) = \beta i\mathcal{V}(i)\Big[-2 + \frac{1}{\beta i}\big(\lambda e^{\alpha} + e^{-\alpha}\mu N - e^{-\alpha}\theta N - \lambda - \mu N + \theta N\big)\Big]. \qquad (4.2)$$

Next, we define the set $\mathcal{K}_1$ by

$$\mathcal{K}_1 \doteq \{i \in S : \beta i \leq \lambda e^{\alpha} + e^{-\alpha}\mu N - e^{-\alpha}\theta N - \lambda - \mu N + \theta N\} \cup \{1,\dots,N\}. \qquad (4.3)$$

Clearly $\mathcal{K}_1$ is finite. Combining (4.3) with (4.2), it follows that

$$\sum_{j\in S} \mathcal{V}(j)q(i,j,0) \leq -\beta i\mathcal{V}(i), \quad \forall i \in \mathcal{K}_1^c.$$

For the case when $u = 1$, we see that

$$\sum_{j\in S} \mathcal{V}(j)q(i,j,1) = -\overline{\mu}(i)\mathcal{V}(i)\big(1-e^{-\alpha}\big) = -\frac{2\beta}{\theta}\overline{\mu}(i)\mathcal{V}(i), \quad \forall i \geq 1.$$

But observe that $\frac{2}{\theta}\overline{\mu}(i) = \frac{2\mu}{\theta}(i\wedge N) + 2(i-N)^+ > i$ whenever $i > N\max\big(1, \frac{2(\theta-\mu)}{\theta}\big)$. Hence, setting

$$\begin{aligned}\mathcal{K}_2 &\doteq \left\{i \in S : i \leq N\cdot\max\left(1, \frac{2(\theta-\mu)}{\theta}\right)\right\}, \\ \mathcal{K} &\doteq \mathcal{K}_1 \cup \mathcal{K}_2,\end{aligned}$$

it follows that

$$\max_{u\in\{0,1\}} \sum_{j\in S} \mathcal{V}(j)q(i,j,u) \leq \hat{C}\mathbb{1}_{\mathcal{K}} - \beta i\mathcal{V}(i), \quad \forall i \in S$$

for some $\hat{C} > 0$. Thus, setting $g(i) = \beta i$, the proof will be complete upon showing that $g(i) - \gamma\max_{u\in\{0,1\}} c(i,u)$ is norm-like. From Assumption 2.1, we have that $a\gamma < 1$, which immediately implies $a\gamma < 1 - e^{-\alpha}$ for some large enough $\alpha > 0$. Hence, selecting such an α, we then have $\beta i - a\gamma\theta i > 0$, which implies that $g(i) - \gamma(h(i) + \lambda\kappa)$ is norm-like. This completes the proof. □

Remark 4.1. *Since $\mathcal{V}(i) = e^{\alpha i}$ for some large enough $\alpha > 0$ (whose exact value is not important), it follows that for any function $f(i) = \eta^i$, $\eta > 0$, we can always choose α to be large enough so that $\frac{\eta}{e^\alpha} < 1$. This implies that f is of small order $\mathcal{V}$. Furthermore, this implies that there exists a $\delta \in (0,1)$ such that f is $\mathcal{O}(\mathcal{V}^\delta)$.*

4.2. **Proof of Theorem 2.1.** Here we provide the proofs for the ERSC Bellman equation for the admission control problem, given in Theorem 2.1. For the remainder of this subsection, we fix $\gamma = 1$ (which can be easily done by absorbing γ into the running cost), and refer to the pair in Theorem 2.1 as (ρ, V).

To prove Theorem 2.1, we will show that by having access to such a pair (ρ, V), we have the following upper bound on ρ:

$$\rho \leq \inf_{i \in S} \inf_{U \in \mathfrak{U}} \mathcal{E}_i(U)$$

as well as the following lower bound on ρ:

$$\rho \geq \inf_{i \in S} \mathcal{E}_i(v)$$

for some stationary policy v. Upon establishing this, it follows that $\rho = \inf_{i \in S} \inf_{U \in \mathfrak{U}} \mathcal{E}_i(U) = \inf_{i \in S} \mathcal{E}_i(v)$ for this policy v.

In order to establish both of these bounds, we require the following result.

Lemma 4.1. *Let $U \in \mathfrak{U}$ and $i \in S$ with $X_0 = i$.*

(i) *Let τ_n be the stopping time defined by*

$$\tau_n \doteq \inf\{t \geq 0 : X_t \notin \{1, \ldots, n-1\}\} = \inf\{t \geq 0 : X_t = n\}.$$

Then we have that $\tau_n \to \infty$ a.s. as $n \to \infty$.

(ii) *Let $\mathcal{B} = \{0 \ldots, B\}$ be a finite set such that $\mathcal{K} \subseteq \mathcal{B}$, and for $i \notin \mathcal{B}$, let $\breve{\tau}(\mathcal{B})$ be the stopping time defined by*

$$\breve{\tau}(\mathcal{B}) \doteq \inf\{t \geq 0 : X_t \in \mathcal{B}\}.$$

Then we have that $\breve{\tau}(\mathcal{B}) < \infty$ a.s.

Proof. We first prove part (i). From Lemma 2.1, since $g : S \to \mathbb{R}_+$ is norm-like (and strictly increasing), it is clear that $\mathcal{K}$ can be enlarged (without loss of generality) so that $g(i)$ is bounded away from 0 outside of $\mathcal{K}$. In other words, there exists $\alpha > 0$ such that $g(i) > \alpha$ for all $i \notin \mathcal{K}$. Following the proof of Lemma 2.1, we have

$$\sup_{u \in \{0,1\}} \sum_{j \in S} \mathcal{V}(j) q(i,j,u) \leq -g(i)\mathcal{V}(i) \leq -\alpha \mathcal{V}(i), \quad \forall i \notin \mathcal{K}.$$

As a result, by increasing $\hat{C}$ if necessary, it follows that

$$\sup_{u \in \{0,1\}} \sum_{j \in S} \mathcal{V}(j) q(i,j,u) \leq \hat{C} \mathbb{1}_{\mathcal{K}} - \alpha \mathcal{V}(i), \quad \forall i \in S. \tag{4.4}$$

From here, by applying Dynkin's formula to the function $e^{\alpha t}\mathcal{V}(X_t)$ up to the stopping time $T \wedge \tau_n$ we have

$$\begin{aligned}
\mathbb{E}_i^U \left[e^{\alpha(T \wedge \tau_n)} \mathcal{V}(X_{T \wedge \tau_n}) \right] &= \mathcal{V}(i) + \mathbb{E}_i^U \left[\int_0^{T \wedge \tau_n} e^{\alpha t} \left(\alpha \mathcal{V}(X_t) + \sum_{j \in S} \mathcal{V}(j) q(i,j,u) \right) dt \right] \\
&\leq \mathcal{V}(i) + \hat{C} \mathbb{E}_i^U \left[\int_0^{T \wedge \tau_n} e^{\alpha t} \, dt \right] \\
&\leq \mathcal{V}(i) + \frac{\hat{C}}{\alpha} \left(e^{\alpha T} - 1 \right).
\end{aligned} \tag{4.5}$$

To obtain the second line in (4.5), we applied (4.4), and to obtain the third line we used the non-negativity of the integrand combined with the fact that $T \wedge \tau_n \leq T$.

Now, observe that

$$\begin{aligned}\mathbb{E}_i^U\left[e^{\alpha(T\wedge\tau_n)}\mathcal{V}(X_{T\wedge\tau_n})\right] &= \mathbb{E}_i^U\left[e^{\alpha\tau_n}\mathcal{V}(n)\mathbb{1}_{\{\tau_n\leq T\}}\right] + \mathbb{E}_i^U\left[e^{\alpha T}\mathcal{V}(X_T)\mathbb{1}_{\{T<\tau_n\}}\right]\\ &\geq \mathbb{E}_i^U\left[e^{\alpha\tau_n}\mathcal{V}(n)\mathbb{1}_{\{\tau_n\leq T\}}\right]\\ &\geq \mathcal{V}(n)\mathbb{P}_i^U(\tau_n\leq T)\,.\end{aligned}$$

Combining this with the last line of (4.5), we have

$$\mathbb{P}_i^U(\tau_n\leq T) \leq \frac{\mathcal{V}(i)+\frac{\hat{C}}{\alpha}\big(e^{\alpha T}-1\big)}{\mathcal{V}(n)}\,. \tag{4.6}$$

Taking $n\to\infty$ in (4.6) together with the fact that $\mathcal{V}(n)\to\infty$, we have

$$\lim_{n\to\infty}\mathbb{P}_i^U(\tau_n\leq T)=0\,. \tag{4.7}$$

Since τ_n is monotonically increasing (as X_t is a birth-death process), we can apply the monotone convergence theorem to (4.7) to obtain

$$\mathbb{P}_i^U\big(\lim_{n\to\infty}\tau_n\leq T\big)=0\,.$$

Since the above holds for all $T>0$, this implies that $\lim_{n\to\infty}\tau_n=\infty$ a.s. for any initial state i and control U. This completes the proof of part (i).

We now prove part (ii). Set n to be large enough such that $i<n$, and let τ_n be as defined in the statement of the lemma. By applying Dynkin's formula to $e^{\alpha t}\mathcal{V}(X_t)$ up to the stopping time $T\wedge\tau_n\wedge\check{\tau}(\mathcal{B})$, it follows from (4.4) and the fact that $i\in\mathcal{B}^c\subseteq\mathcal{K}^c$ that

$$\mathbb{E}_i^U\left[e^{\alpha(T\wedge\tau_n\wedge\check{\tau}(\mathcal{B}))}\mathcal{V}(X_{T\wedge\tau_n\wedge\check{\tau}(\mathcal{B})})\right]\leq\mathcal{V}(i)\,. \tag{4.8}$$

From here, using the fact that $\mathcal{V}$ is increasing, it follows that $\mathcal{V}(X_{T\wedge\tau_n\wedge\check{\tau}(\mathcal{B})})\geq\mathcal{V}(B)$. Moreover, both finiteness of $\mathcal{B}$ and Lemma 2.1 imply $1\leq\mathcal{V}(B)<\infty$. Combining these with (4.8), we have

$$\mathbb{E}_i^U\left[e^{\alpha(T\wedge\tau_n\wedge\check{\tau}(\mathcal{B}))}\right]\leq\frac{\mathcal{V}(i)}{\mathcal{V}(B)}\,. \tag{4.9}$$

From here, taking $T\to\infty$ followed by $n\to\infty$ and using the fact that $\tau_n\to\infty$ a.s. from part (i), we have $\mathbb{E}_i^U\big[e^{\alpha\check{\tau}(\mathcal{B})}\big]\leq\frac{\mathcal{V}(i)}{\mathcal{V}(B)}$. From here, the desired result follows as a consequence of Jensen's inequality together with facts that $\mathcal{V}$ is increasing and $\mathcal{V}\geq 1$. This completes the proof of part (ii). □

We show the upper bound in the next lemma.

Lemma 4.2. *Suppose there exists a pair (ρ, V) which satisfies the conditions in Theorem 2.1. Then we have $\rho\leq\inf_i\inf_{U\in\mathfrak{U}}\mathcal{E}_i(U)$.*

Proof. Our approach is heavily inspired by the one used in the proof of [22, Theorem 3.2]. Since our setting deals with a Markov process taking values in a discrete state space, and our running cost is unbounded, we present the adapted proof for completeness.

Fix $U\in\mathfrak{U}$ and $i\in S$. For n larger than i, we define the stopping time τ_n as in Lemma 4.1(i). By applying Dynkin's formula to the function $e^{\int_0^t(c(X_s,U_s)-\rho)\,ds}V(X_t)$ up to the stopping time $T\wedge\tau_n$,

we have

$$\begin{aligned}
&\mathbb{E}_i^U\left[e^{\int_0^{T\wedge\tau_n}[c(X_t,U_t)-\rho]dt}V\big(X_{T\wedge\tau_n}\big)\right]\\
&= V(i)+\mathbb{E}_i^U\left[\int_0^{T\wedge\tau_n} e^{\int_0^t[c(X_s,U_s)-\rho]ds}\bigg(\sum_{j\in S}V(j)q(X_t,j,U_t)+\big[c(X_t,U_t)-\rho\big]V(X_t)\bigg)\,dt\right]\\
&\geq V(i), \qquad (4.10)
\end{aligned}$$

where (4.10) follows from the fact that (ρ, V) satisfies (2.13). Next, observe that

$$\mathbb{E}_i^U\left[e^{\int_0^{T\wedge\tau_n}[c(X_t,U_t)-\rho]dt}V\big(X_{T\wedge\tau_n}\big)\right] = \mathbb{E}_i^U\left[e^{\int_0^{T}[c(X_t,U_t)-\rho]dt}V\big(X_{T}\big)\mathbb{1}_{\{T<\tau_n\}}\right] \qquad (4.11)$$

$$+\,\mathbb{E}_i^U\left[e^{\int_0^{\tau_n}[c(X_t,U_t)-\rho]dt}V(n)\mathbb{1}_{\{\tau_n\leq T\}}\right]. \qquad (4.12)$$

We argue that (4.12) vanishes as $n\to\infty$. To do this, first note that the Lyapunov stability condition in (2.12) together with the fact that $\mathcal{V}\geq 1$ implies

$$\max_{u\in\{0,1\}}\sum_{j\in S}\mathcal{V}(j)q(i,j,u)\leq\big(\hat{C}-g(i)\big)\mathcal{V}(i),\quad \forall i\in S\,. \qquad (4.13)$$

Applying Dynkin's Formula on the function $e^{\int_0^t\left(g(X_t)-\hat{C}\right)g(X_t)\,ds}\mathcal{V}(X_t)$ up to the stopping time $T\wedge\tau_n$, it follows that

$$\begin{aligned}
&\mathbb{E}_i^U\left[e^{\int_0^{T\wedge\tau_n}\left(g(X_t)-\hat{C}\right)dt}\mathcal{V}(X_{T\wedge\tau_n})\right]\\
&= \mathcal{V}(i)+\mathbb{E}_i^U\left[\int_0^{T\wedge\tau_n} e^{\int_0^t(g(X_s)-\hat{C})ds}\bigg(\sum_{j\in S}\mathcal{V}(j)q(X_t,j,U_t)+\big(g(X_t)-\hat{C}\big)\mathcal{V}(X_t)\bigg)\,dt\right]\\
&\leq \mathcal{V}(i), \qquad (4.14)
\end{aligned}$$

where the inequality in (4.14) follows from (4.13). Because the bound in (4.14) is independent of n, it follows from Fatou's Lemma and Lemma 4.1(i) that

$$\mathbb{E}_i^U\left[e^{\int_0^{T}\left(g(X_t)-\hat{C}\right)dt}\mathcal{V}(X_T)\right]\leq\liminf_{n\to\infty}\,\mathbb{E}_i^U\left[e^{\int_0^{T\wedge\tau_n}\left(g(X_t)-\hat{C}\right)dt}\mathcal{V}(X_{T\wedge\tau_n})\right]\leq\mathcal{V}(i)\,. \qquad (4.15)$$

Multiplying the constant $\hat{C}$ in (4.15) to the other side, we have shown

$$\mathbb{E}_i^U\left[e^{\int_0^{T}g(X_t)\,dt}\mathcal{V}(X_T)\right]\leq e^{\hat{C}T}\mathcal{V}(i)\,. \qquad (4.16)$$

Now, recall from Lemma 2.1 that $g(i)-\max_u c(i,u)$ is norm-like. As a result, it follows that there exists a positive constant C_1 such that

$$\max_u c(i,u)-\rho\leq\max_u c(i,u)\leq g(i)+C_1,\quad\forall i\in S\,. \qquad (4.17)$$

Using (4.16) and $\mathcal{V}\geq 1$, it is easy to check that

$$\mathbb{E}_i^U\left[e^{\int_0^{T}(g(X_t)+C_1)\,dt}\right]\leq\mathbb{E}_i^U\left[e^{\int_0^{T}(g(X_t)+C_1)\,dt}\mathcal{V}(X_T)\right]\leq e^{(\hat{C}+C_1)T}\mathcal{V}(i)\,. \qquad (4.18)$$

Now, observe that

$$\mathbb{E}_i^U\left[e^{\int_0^{T\wedge\tau_n}(g(X_t)+C_1)\,dt}\right] \qquad (4.19)$$

$$= \mathbb{E}_i^U\left[e^{\int_0^{T}(g(X_t)+C_1)\,dt}\mathbb{1}_{\{T<\tau_n\}}\right]+\mathbb{E}_i^U\left[e^{\int_0^{\tau_n}(g(X_t)+C_1)\,dt}\mathbb{1}_{\{\tau_n\leq T\}}\right]. \qquad (4.20)$$

Since the integrand in (4.19) is positive, we apply the monotone convergence theorem together with Lemma 4.1(i) to obtain

$$\lim_{n\to\infty} \mathbb{E}_i^U\left[e^{\int_0^{T\wedge\tau_n}(g(X_t)+C_1)\,dt}\right] = \mathbb{E}_i^U\left[e^{\int_0^{T}(g(X_t)+C_1)\,dt}\right] < \infty,$$

where finiteness follows from (4.18). By a similar analysis, we also have

$$\lim_{n\to\infty} \mathbb{E}_i^U\left[e^{\int_0^{T}(g(X_t)+C_1)\,dt}\mathbb{1}_{\{T<\tau_n\}}\right] = \mathbb{E}_i^U\left[e^{\int_0^{T}(g(X_t)+C_1)\,dt}\right] < \infty\,.$$

Hence, taking $n\to\infty$ in (4.19)-(4.20), it follows that

$$\lim_{n\to\infty} \mathbb{E}_i^U\left[e^{\int_0^{\tau_n}(g(X_t)+C_1)\,dt}\mathbb{1}_{\{\tau_n\le T\}}\right] = 0,$$

and so by (4.17) we must have

$$\lim_{n\to\infty} \mathbb{E}_i^U\left[e^{\int_0^{\tau_n}(c(X_t,U_t)-\rho)\,dt}\mathbb{1}_{\{\tau_n\le T\}}\right] = 0\,. \tag{4.21}$$

We are now in a position to argue that (4.12) vanishes as $n\to\infty$. Our approach involves splitting the level sets of V, which was a technique that was originally introduced in [4] (see the proof of Theorem 4.3). As such, we fix $k\in\mathbb{Z}^+$. For each n, we either have $V(n)<k$ or $V(n)\ge k$. Since V is $\mathcal{O}(\mathcal{V}^\delta)$, we know there exists $\overline{C}>0$ such that $V\le \overline{C}\,\mathcal{V}^\delta$ on S. Hence, it follows that

$$\begin{aligned}
\mathbb{E}_i^U\left[e^{\int_0^{\tau_n}[c(X_t,U_t)-\rho]dt}V(n)\mathbb{1}_{\{\tau_n\le T\}}\right] &= \mathbb{E}_i^U\left[e^{\int_0^{\tau_n}[c(X_t,U_t)-\rho]dt}V(n)\mathbb{1}_{\{\tau_n\le T\}}\mathbb{1}_{\{V(n)<k\}}\right]\\
&\quad+\mathbb{E}_i^U\left[e^{\int_0^{\tau_n}[c(X_t,U_t)-\rho]dt}V(n)\mathbb{1}_{\{\tau_n\le T\}}\mathbb{1}_{\{V(n)\ge k\}}\right]\\
&< k\mathbb{E}_i^U\left[e^{\int_0^{\tau_n}[c(X_t,U_t)-\rho]dt}\mathbb{1}_{\{\tau_n\le T\}}\right]\\
&\quad+\overline{C}\mathbb{E}_i^U\left[e^{\int_0^{\tau_n}[c(X_t,U_t)-\rho]dt}\mathcal{V}(n)^\delta\mathbb{1}_{\{\tau_n\le T\}}\mathbb{1}_{\{V(n)\ge k\}}\right]\\
&\le k\mathbb{E}_i^U\left[e^{\int_0^{\tau_n}[c(X_t,U_t)-\rho]dt}\mathbb{1}_{\{\tau_n\le T\}}\right]\\
&\quad+\overline{C}^{2-\delta}k^{\delta-1}\mathbb{E}_i^U\left[e^{\int_0^{\tau_n}[c(X_t,U_t)-\rho]dt}\mathcal{V}(n)\mathbb{1}_{\{\tau_n\le T\}}\right]\\
&\le k\mathbb{E}_i^U\left[e^{\int_0^{\tau_n}[c(X_t,U_t)-\rho]dt}\mathbb{1}_{\{\tau_n\le T\}}\right] &(4.22)\\
&\quad+\overline{C}^{2-\delta}k^{\delta-1}e^{(\hat{C}+C_1)T}\mathcal{V}(i), &(4.23)
\end{aligned}$$

where in the above display we used the property that V is $\mathcal{O}(\mathcal{V}^\delta)$ for some $\delta\in(0,1)$ in the fourth line and (4.18) to obtain (4.23). Note that (4.22) vanishes if we take $n\to\infty$ by (4.21), and (4.22) vanishes if we take $k\to\infty$ since $\delta<1$. Thus, taking $n\to\infty$ followed by $k\to\infty$ in (4.22)-(4.23), we have (4.12) vanishes as $n\to\infty$. Applying this to (4.11)-(4.12), we have shown that for any $\varepsilon>0$, we can choose n sufficiently large so that

$$V(i)\le \mathbb{E}_i^U\left[e^{\int_0^{T}[c(X_t,U_t)-\rho]dt}V\big(X_T\big)\mathbb{1}_{\{T<\tau_n\}}\right]+\varepsilon\,. \tag{4.24}$$

We now focus on bounding (4.24) in terms of the finite horizon risk-sensitive cost. First, note that $c(i,u)\le g(i)$ for any $u\in\{0,1\}$ for all i outside some finite set $\mathcal{D}$. Combining this with the

Lyapunov condition in (2.12), we see

$$\sum_{j\in S}\mathcal{V}(j)q(i,j,u)\le -g(i)\mathcal{V}(i)\le -c(i,u)\,\mathcal{V}(i)\quad \forall u\in\{0,1\},\,\forall i\notin\mathcal{K}\cup\mathcal{D}\doteq\overline{\mathcal{K}}\,.$$

Hence, there exists a constant $\hat{C}_2>0$ such that

$$\sum_{j\in S}\mathcal{V}(j)q(i,j,u)\le \hat{C}_2\mathbb{1}_{\overline{\mathcal{K}}}-c(i,u)\,\mathcal{V}(i),\quad \forall u\in\{0,1\},\,\forall i\in S\,. \tag{4.25}$$

We next apply Dynkin's formula to $e^{\int_0^T c(X_t,U_t)dt}\mathcal{V}\big(X_T\big)\mathbb{1}_{\{T\le\tau_n\}}$. Note that this is justified from the fact that $\mathbb{E}\big[e^{\int_0^T c(X_t,U_t)\,dt}\mathcal{V}(X_T)\big]$ is finite (see (4.16) and the proceeding arguments, replacing g with c and the original Lyapunov drift condition with (4.25)), together with the fact that the mean of the generator of this process is finite.

Hence, using $V\le C\mathcal{V}^\delta\le C\mathcal{V}$, applying Dynkin's formula to $e^{\int_0^T c(X_t,U_t)dt}\mathcal{V}\big(X_T\big)\mathbb{1}_{\{T\le\tau_n\}}$, and using (4.25), it follows that

$$\begin{aligned}
&\mathbb{E}_i^U\left[e^{\int_0^T c(X_t,U_t)dt}\,V\big(X_T\big)\mathbb{1}_{\{T\le\tau_n\}}\right] &(4.26)\\
&\le \overline{C}\mathbb{E}_i^U\left[e^{\int_0^T c(X_t,U_t)dt}\,\mathcal{V}\big(X_T\big)\mathbb{1}_{\{T\le\tau_n\}}\right]\\
&=\overline{C}\mathcal{V}(i)+\overline{C}\mathbb{E}_i^U\left[\int_0^T e^{\int_0^t c(X_s,U_s)\,ds}\left(\sum_{j\in S}\mathcal{V}(j)q(X_t,j,U_t)+c(X_t,U_t)\mathcal{V}(X_t)\right)\mathbb{1}_{\{T<\tau_n\}}\,dt\right]\\
&\le\overline{C}\mathcal{V}(i)+\overline{C}\mathbb{E}_i^U\left[\int_0^T e^{\int_0^t c(X_s,U_s)\,ds}\Big(\hat{C}_2-c(X_t,U_t)\mathcal{V}(X_t)+c(X_t,U_t)\mathcal{V}(X_t)\Big)\mathbb{1}_{\{T<\tau_n\}}\,dt\right]\\
&\le\overline{C}\big(\mathcal{V}(i)+\hat{C}_2T\big)\,\mathbb{E}_i^U\left[e^{\int_0^T c(X_t,U_t)\,dt}\right], &(4.27)
\end{aligned}$$

where in the last line we used the monotonicity of the integrand to pull out T and crudely bounded the indicator function. By multiplying both sides of the inequality in (4.26)-(4.27) by $e^{-\rho T}$, we see that

$$\mathbb{E}_i^U\left[e^{\int_0^T\big(c(X_t,U_t)-\rho\big)dt}\,V\big(X_T\big)\mathbb{1}_{\{T\le\tau_n\}}\right]\le\overline{C}\big(\mathcal{V}(i)+\hat{C}_2T\big)e^{-\rho T}\,\mathbb{E}_i^U\left[e^{\int_0^T c(X_t,U_t)\,dt}\right].$$

Combining this with (4.24), it follows that

$$\begin{aligned}
V(i)&\le\mathbb{E}_i^U\left[e^{\int_0^T\big(c(X_t,U_t)-\rho\big)dt}\,V\big(X_T\big)\mathbb{1}_{\{T\le\tau_n\}}\right]+\varepsilon\\
&\le\overline{C}\big(\mathcal{V}(i)+\hat{C}_2T\big)e^{-\rho T}\,\mathbb{E}_i^U\left[e^{\int_0^T c(X_t,U_t)\,dt}\right]+\varepsilon\,.
\end{aligned}$$

From here, taking the natural logarithm of both sides, dividing by T, and taking $T\to\infty$, we have shown

$$\rho\le\limsup_{T\to\infty}\frac{1}{T}\log\mathbb{E}_i^U\left[e^{\int_0^T c(X_t,U_t)\,dt}\right].$$

Since i and U were arbitrary, the desired result follows. □

We move onto proving the lower bound and showing that the optimal ERSC cost is attained at a stationary policy.

Lemma 4.3. *Suppose there exists a pair (ρ,V) which satisfies the conditions in Theorem 2.1. Then there exists a stationary policy $v\in\mathfrak{U}_{sm}$ such that $\rho\ge\inf_{i\in S}\mathcal{E}_i(v)$.*

Proof. Our approach here is similar to the lower bound proof for optimality for the ERSC problem in the case that the running cost satisfies the near-monotone condition (see [10, Section 3.1]). Note that because the running cost $\min_{u\in\{0,1\}} c(i,u) = h(i)$ is monotonically increasing and unbounded, we can choose a finite subset $\mathcal{B} = \{0,\dots,B\}$ of S such that $\mathcal{B}$ contains $\mathcal{K}$ and

$$h(i) - \rho \geq 0, \quad \forall i \notin \mathcal{B}\,.$$

Next, fix $i \notin \mathcal{B}$, and for $X_0 = i$, we define the stopping time $\check{\tau}(\mathcal{B})$ as in Lemma 4.1(ii). Let τ_n be defined as in Lemma 4.2(i), and choose n sufficiently large so that $\mathcal{B}\cup\{i\} \subseteq \{0,\dots,n\}$. Now, let v be the stationary policy which is defined as a minimizer to (2.13). By applying Dynkin's formula to $\int_0^t e^{c(X_s,v(X_s))-\rho)\,ds} V(X_t)$ up to the stopping time $\check{\tau}(\mathcal{B})\wedge\tau_n\wedge T$, we have

$$\begin{aligned}
&\mathbb{E}_i^v\left[e^{\int_0^{\check{\tau}(\mathcal{B})\wedge\tau_n\wedge T}\left(c(X_t,v(X_t))-\rho\right)\,dt} V(X_{\check{\tau}\wedge\tau_n\wedge T})\right] \qquad (4.28)\\
&= V(i) + \mathbb{E}_i^v\Bigg[\int_0^{\check{\tau}(\mathcal{B})\wedge\tau_n\wedge T} e^{\int_0^t (c(X_s,v(X_s))-\rho)\,ds}\Bigg(\sum_{j\in S} V(j)q(X_t,j,v(X_t))\\
&\qquad\qquad + \big(c(X_t,v(X_t))-\rho\big)V(X_t)\Bigg)\,dt\Bigg]\\
&= V(i), \qquad (4.29)
\end{aligned}$$

where (4.29) follows because (ρ, V) satisfies (2.13) and our choice of v. Now, taking $n\to\infty$ in (4.28)-(4.29), from Fatou's Lemma together with Lemma 4.1(i) we have

$$\begin{aligned}
V(i) &= \liminf_{n\to\infty} \mathbb{E}_i^v\left[e^{\int_0^{\check{\tau}(\mathcal{B})\wedge\tau_n\wedge T}\left(c(X_t,v(X_t))-\rho\right)\,dt} V(X_{\check{\tau}\wedge\tau_n\wedge T})\right]\\
&\geq \mathbb{E}_i^v\left[e^{\int_0^{\check{\tau}(\mathcal{B})\wedge T}\left(c(X_t,v(X_t))-\rho\right)\,dt} V(X_{\check{\tau}\wedge T})\right]. \qquad (4.30)
\end{aligned}$$

Taking $T\to\infty$ in (4.30) in the same way and using $\check{\tau}(\mathcal{B}) < \infty$ a.s. from Lemma 4.1(ii), we get

$$V(i) \geq \mathbb{E}_i^v\left[e^{\int_0^{\check{\tau}(\mathcal{B})}\left(c(X_t,v(X_t))-\rho\right)\,dt} V(X_{\check{\tau}})\right] \geq V(B), \qquad (4.31)$$

where the last inequality follows because the exponential is larger than one by our choice of $\mathcal{B}$ and $X(t)$ is a birth and death process. But $V > 0$ and $i \notin \mathcal{B}$ was arbitrary, and so (4.31) implies that $\inf_S V > 0$, *i.e.,* V has a strictly positive lower bound on S.

Now, fix $i \in S$. We apply Dynkin's formula in a manner similar as above to $e^{\int_0^t \left(c(X_s,v(X_s))-\rho\right)\,ds} V(X_t)$ up to the stopping time $T\wedge\tau_n$, and then take $n\to\infty$ as before to obtain that

$$V(i) \geq \mathbb{E}_i^v\left[e^{\int_0^T \left(c(X_t,v(X_t))-\rho\right)\,dt} V(X_T)\right] \geq \left(\inf_S V\right)\mathbb{E}_i^v\left[e^{\int_0^T \left(c(X_t,v(X_t))-\rho\right)\,dt}\right].$$

From here, we take natural log of both sides, divide by T, and take $T\to\infty$ to obtain

$$\rho \geq \limsup_{T\to\infty} \frac{1}{T}\log \mathbb{E}_i^v\left[e^{\int_0^T c(X_t,v(X_t))\,dt}\right] = \mathcal{E}_i(v)\,.$$

This completes the proof. □

Completing the proof for Theorem 2.1. Given the pair (ρ, V) which satisfies the conditions in Theorem 2.1, the results of Lemmas 4.2 and 4.3 imply that

$$\rho = \inf_{i\in S}\inf_{U\in\mathfrak{U}} \mathcal{E}_i(U) = \inf_{i\in S}\mathcal{E}(v),$$

where v is the minimizing stationary policy in (2.13). This shows that ρ is the optimal cost and v is an optimal stationary policy.

To show this implies the optimality of a threshold policy, there are two cases: (i) v never exercises admission control (*i.e.,* $v(i) = 0$ for all $i \in S$), or (ii) v exercises admission control at some threshold (*i.e.,* $v(i) = 1$ for some $i \in S$).

For (i), it is clear that v is of the form of a threshold policy which regulates with an infinite threshold. As for (ii), let ℓ denote the minimum state i where $v(i) = 1$. Clearly (ρ, V) satisfy the following equations:

$$\rho V(i) = c(i,0)V(i) + \sum_{j\in S} q(i,j,0)V(j), \quad i \le \ell - 1,$$
$$\rho V(\ell) = c(\ell,1)V(\ell) + \sum_{j\in S} q(\ell,j,1)V(j)\,.$$

However, these are the dynamic programming equations for a finite state ERSC problem corresponding to a threshold policy regulating at ℓ (see [28, Theorem 3.2]). By uniqueness, it follows that $\rho = \rho_\ell$, where ρ_ℓ is the ERSC cost corresponding to a threshold ℓ. From the inequalities involving ρ above, we then have

$$\rho_\ell = \inf_{i\in S}\inf_{U\in\mathfrak{U}} \mathcal{E}_i(U) = \inf_{i\in S} \mathcal{E}(v_\ell),$$

where v_ℓ denotes the threshold policy regulating at ℓ. And so it follows that the threshold policy regulating at ℓ is optimal. □

4.3. **Proof of Theorem 2.2.** First, we establish a useful monotonicity property between the cost ρ_n and the corresponding value function y_n associated with a threshold policy of level n, which will be used frequently in this and later sections. This is similar in structure to the risk-neutral version presented in [16, Lemma 3.1].

Lemma 4.4. *Suppose that $\rho_{n_1} > \rho_{n_2}$ for some n_1, n_2. Then $y_{n_1}(i) > y_{n_2}(i)$ for $i = 1, \dots, \min\{n_1, n_2\}+1$. The above also holds if "$>$" is replaced by "$\ge$".*

Proof. Assume $\rho_{n_1} > \rho_{n_2}$. Rewriting (2.16) gives us

$$\gamma\rho_{n_1} = \gamma h(0) + \lambda y_{n_1}(1) - \lambda,$$
$$\gamma\rho_{n_2} = \gamma h(0) + \lambda y_{n_2}(1) - \lambda\,.$$

Subtracting the second equation from the first, it is clear that $y_{n_1}(1) > y_{n_2}(1)$. Next, rewriting (2.15) for $i = 1$ gives us

$$\gamma\rho_{n_1} = \gamma h(1) - \overline{\mu}(1) + \overline{\mu}(1)\frac{1}{y_{n_1}(1)} + \lambda y_{n_1}(2) - \lambda,$$
$$\gamma\rho_{n_2} = \gamma h(1) - \overline{\mu}(1) + \overline{\mu}(1)\frac{1}{y_{n_2}(1)} + \lambda y_{n_2}(2) - \lambda\,.$$

Subtracting the second equation from the first, and using the fact that $y_{n_1}(1) > y_{n_2}(1)$ gives us $y_{n_1}(2) > y_{n_2}(2)$. The same analysis applies by induction to give us $y_{n_1}(i) \ge y_{n_2}(i)$ for $i = 1, \dots, \min\{n_1, n_2\} + 1$. □

The following result shows that it is never optimal to intentionally idle available servers (just as in the risk-neutral setting, see [16, Lemma 3.2]), and justifies Algorithm 1 starting with $n = N$.

Lemma 4.5. *We have $\rho_0 > \rho_1 > \cdots > \rho_N$.*

Proof. We proceed by inducting on $n \le N - 1$.

Base case: $\rho_0 > \rho_1$.

Observe that $\gamma\rho_0 = \gamma h_I N + \lambda\gamma\kappa$, as this involves letting no people into the system. We show that $\rho_1 < \rho_0$. Let $y = y_1(1) > 0$. Observe that ρ_1 satisfies

$$\gamma\rho_1 = \gamma h_I N + \lambda(y-1),$$
$$\gamma\rho_1 = \gamma h_I(N-1) - \mu + \mu\frac{1}{y} + \lambda\gamma\kappa\,.$$

Setting the above equations equal to each other results in

$$\gamma h_I N + \lambda(y-1) = \gamma h_I(N-1) - \mu + \mu\frac{1}{y} + \lambda\gamma\kappa\,.$$

Multiplying both sides by y and collecting terms results in

$$\lambda y^2 + \big(\gamma h_I + \mu - \lambda(\gamma\kappa+1)\big)y - \mu = 0.$$

Let $b = \gamma h_I + \mu - \lambda(\gamma\kappa+1)$. The above equation is a quadratic equation, with two solutions for y (one negative and one positive). Since y is strictly positive, it must satisfy

$$y = \frac{-b + \sqrt{b^2 + 4\lambda\mu}}{2\lambda}. \tag{4.32}$$

Now, observe that

$$\gamma(\rho_0 - \rho_1) = h_I N + \kappa - (h_I N + y - 1) = \lambda\gamma\kappa - \lambda y + \lambda.$$

Hence, it suffices to show that $y < 1 + \gamma\kappa$ in order to prove that $\rho_1 < \rho_0$. From equation (4.32), we know that

$$\begin{aligned} & y < 1 + \gamma\kappa \\ \iff & \frac{-b + \sqrt{b^2 + 4\lambda\mu}}{2\lambda} < 1 + \gamma\kappa \\ \iff & \sqrt{b^2 + 4\lambda\mu} < b + 2\lambda(\gamma\kappa + 1) \\ \iff & \mu < (\gamma\kappa+1)(b + \lambda + \lambda\gamma\kappa)\,. \end{aligned}$$

But substituting $b = \gamma h_I + \mu - \lambda(\gamma\kappa+1)$ into the previous equation results in

$$\mu < (1+\gamma\kappa)(\gamma h_I + \mu),$$

which is always true. Thus, we have $\rho_1 < \rho_0$, completing the base case.

Induction Step: Suppose for $n \le N-1$, we have $\rho_0 > \rho_1 > \cdots > \rho_n$. Suppose for contradiction that $\rho_{n+1} \ge \rho_n$. By Lemma 4.4, we have $y_n(i) \le y_{n+1}(i)$ for $i = 1, \ldots, n+1$. By the induction hypothesis and Lemma 4.4, we know that $y_{n-1}(i) > y_n(i)$ for $i = 1, \ldots, n$. Since $y_{n-1}(n) = y_n(n+1) = \gamma\kappa + 1$ by (2.22), we have that $y_n(n) < \gamma\kappa + 1$ and $y_{n+1}(n+1) \ge \gamma\kappa + 1$. Applying equation (2.15) with n and $n+1$ results in

$$\gamma\rho_n = \gamma h(n) - \mu n + \mu n\frac{1}{y_n(n)} + \lambda\gamma\kappa,$$
$$\gamma\rho_{n+1} = \gamma h(n+1) - \mu(n+1) + \mu(n+1)\frac{1}{y_{n+1}(n+1)} + \lambda\gamma\kappa\,.$$

Using $y_n(n) < 1 + \gamma\kappa$, and $y_{n+1}(n+1) \ge 1 + \gamma\kappa$, we have

$$\gamma\rho_n > \gamma h(n) - \mu n + \mu n\frac{1}{1+\gamma\kappa} + \lambda\gamma\kappa,$$
$$\gamma\rho_{n+1} \le \gamma h(n+1) - \mu(n+1) + \mu(n+1)\frac{1}{1+\gamma\kappa} + \lambda\gamma\kappa\,.$$

Subtracting these two inequalities, we obtain

$$\gamma(\rho_n - \rho_{n+1}) > \gamma h_I + \mu - \frac{\mu}{1+\gamma\kappa} > 0,$$

as $h(n) - h(n+1) = h_I$ for $n \leq N-1$.

However, we assumed that $\rho_n \leq \rho_{n+1}$, and so we arrive at a contradiction. Thus, we have $\rho_n > \rho_{n+1}$. This completes the verification of the inductive step and also the proof. □

Proof of Theorem 2.2. Our approach is inspired and closely related to the proof of Theorem 3.2 in [16], with the primary differences being that we work with the multiplicative Bellman equation and a more general condition involving the cost and sensitivity parameters. At the end of the proof, we explicitly define a solution to the multiplicative Bellman equation in Lemma 2.2, which results in characterizing an optimal threshold policy. First, we suppose for contradiction that ρ_m is not a global minimum of the sequence $\{\rho_k\}_{k\geq 0}$, so there exists a state $n \geq m+2$ such that $\rho_n < \rho_m$. Second, we construct a modified finite state space admission control problem, and define $\tilde{\rho}_k$ as the cost corresponding to a threshold k in the modified problem. Third, we argue that $\rho_m = \tilde{\rho}_m \leq \tilde{\rho}_n$ by using the Bellman equation for the modified problem. Finally, we show that $\tilde{\rho}_n \leq \rho_n$, which creates a contradiction since we then have $\rho_n < \rho_m \leq \tilde{\rho}_n \leq \rho_n$. This proves that ρ_m is a global minimum of $\{\rho_k\}_{k\geq 0}$. This allows us to easily construct a solution to the Bellman equation in Lemma 2.2, proving that controlling with a threshold level of m is optimal over all admissible control policies.

Following the approach described above, suppose that ρ_m, $m \geq N$ is the first local minimum of $\{\rho_0, \rho_1, \rho_2, \dots\}$ (the fact that $m \geq N$ is a consequence of Lemma 4.5), and suppose for contradiction that ρ_m is not a global minimum of this sequence. This implies that there exists $n \geq m+2$ such that $\rho_n < \rho_m$. Without loss of generality, we also have that $\rho_n \leq \rho_i \ \forall i \leq n$. This is immediately true if $n = m+2$. If $n > m+2$ and $\rho_n > \rho_k$ for some $m+2 < k < n$, then we can simply work with k instead of n. Continuing this chain of reasoning, we can always select an n such that $\rho_n \leq \rho_k$ for all $k \leq n$.

Now, let $\delta = \rho_{m+1} - \rho_m$. By assumption of ρ_m being a local minimum, it is immediate that $\delta \geq 0$. We now construct a modified admission control problem with a finite state space $\{0, \dots, n+1\}$. We define the modified cost and death rates as follows:

$$\tilde{h}(i) = \begin{cases} h(i), & i \leq m, \\ h(m+1) - \delta, & m+1 \leq i \leq n+1, \end{cases} \tag{4.33}$$

$$\tilde{\mu}(i) = \begin{cases} \overline{\mu}(i), & i \leq m, \\ \overline{\mu}(m+1), & m+1 \leq i \leq n+1\,. \end{cases} \tag{4.34}$$

Since the state space is finite, we also set the birth rate at $n+1$ to be 0, which we specify incurs a cost of rejecting an admission. Now, we consider the ERSC problem with the modified cost and death rates given in (4.33)-(4.34). Since the state space is finite, it is clear by using Dynkin's formula that having a solution to the ERSC Bellman equation characterizes optimality (see Theorem 4.1 of [9] and the paragraph immediately following it). Additionally, the same analysis leading up to Lemma 2.2 shows that we can characterize a solution pair of the ERSC Bellman equation as a pair $(\tilde{\rho}, \tilde{y})$ which satisfies

$$\gamma\tilde{\rho} = \gamma\tilde{h}(i) - \tilde{\mu}(i) + \frac{\tilde{\mu}(i)}{\tilde{y}^*(i)} + \lambda \min\{\tilde{y}(i+1) - 1, \gamma\kappa\}, \quad i = 1, \dots, n, \tag{4.35}$$

$$\gamma\tilde{\rho} = \gamma\tilde{h}(0) + \lambda\tilde{y}(1) - \lambda, \tag{4.36}$$

$$\gamma\tilde{\rho} = \gamma\tilde{h}(n+1) - \tilde{\mu}(n+1) + \frac{\tilde{\mu}(n+1)}{\tilde{y}(n+1)} + \lambda\gamma\kappa\,. \tag{4.37}$$

Now, define $\tilde{y} : \{0, \dots, n+1\} \to \mathbb{R}_+$ by

$$\tilde{y}(i) = \begin{cases} y_m(i), & i \le m+1, \\ y_{m+1}(m+1), & m+1 < i \le n+1\,. \end{cases}$$

We verify that $(\rho_m, \tilde{y})$ satisfies (4.35)-(4.37). Since $\rho_0 > \rho_1 > \cdots > \rho_m$, it follows from Lemma 4.4 that $y_m(i) < \gamma\kappa + 1$ for $i \le m$. From this as well as the definition of (ρ_m, y_m) in (2.20), we have that $(\rho_m, \tilde{y})$ satisfies (4.35)-(4.36) for $i \le m$, as shown below

$$\begin{aligned} \gamma\rho_m &= \gamma h(i) - \mu(i) + \frac{\mu(i)}{y_m(i)} + \lambda y_m(i+1) - \lambda \\ &= \gamma\tilde{h}(i) - \tilde{\mu}(i) + \frac{\tilde{\mu}(i)}{\tilde{y}(i)} + \lambda \min\{\tilde{y}(i+1) - 1, \gamma\kappa\}\,. \end{aligned}$$

Next, since $\rho_m \le \rho_{m+1}$, it follows by Lemma 4.4 that $y_m(i) \le y_{m+1}(i)$ for $i \le m+1$. Furthermore, $y_m(m+1) = \gamma\kappa + 1$ by its definition in (2.22), which also means that $y_{m+1}(m+1) \ge \gamma\kappa + 1$. And so substituting $(\rho_m, \tilde{y})$ into (4.35) for $i \ge m+1$, we have

$$\begin{aligned} &\gamma\rho_m = \gamma\tilde{h}(i) - \tilde{\mu}(i) + \frac{\tilde{\mu}(i)}{\tilde{y}(i)} + \lambda \min\{\tilde{y}(i+1), \gamma\kappa\} \\ \iff &\gamma\rho_m = \gamma(h(m+1) - \delta) - \overline{\mu}(m+1) + \frac{\overline{\mu}(m+1)}{y_{m+1}(m+1)} + \lambda \min\{y_{m+1}(m+1) - 1, \gamma\kappa\} \\ \iff &\gamma\rho_{m+1} = \gamma h(m+1) - \overline{\mu}(m+1) + \frac{\overline{\mu}(m+1)}{y_{m+1}(m+1)} + \lambda y_{m+1}(m+1) - \lambda, \end{aligned}$$

which is true by the definition of (ρ_{m+1}, y_{m+1}) in (2.20). Hence, we have shown that $(\rho_m, \tilde{y})$ solve the Bellman equation for the modified finite state space problem, which proves that $\rho_m = \tilde{\rho}^*$ (the optimal cost of the modified admission control problem), and so a threshold policy regulating at m is optimal for the modified problem. In particular, it follows that $\tilde{\rho}_n \ge \rho_m$.

From here, we argue that $\tilde{\rho}_n \le \rho_n$. Observe that

$$\gamma(\tilde{\rho}_n - \rho_n) = \gamma(\tilde{h}(n) - h(n)) - \gamma\delta + \overline{\mu}(n) - \tilde{\mu}(n) + \frac{\tilde{\mu}(n)}{\tilde{y}_n(n)} - \frac{\overline{\mu}(n)}{y_n(n)}$$

by their definitions in (2.20) and (4.35). By adding and subtracting $\frac{\tilde{\mu}(n)}{y_n(n)}$ on the right hand side of the above equation, we have

$$\begin{aligned} \gamma(\tilde{\rho}_n - \rho_n) &= \gamma(\tilde{h}(n) - h(n)) - \gamma\delta + (\tilde{\mu}(n) - \overline{\mu}(n))\left(\frac{1}{y_n(n)} - 1\right) - \tilde{\mu}(n)\left(\frac{1}{y_n(n)} - \frac{1}{\tilde{y}_n(n)}\right) \\ &\le a\theta\gamma(m+1-n) + \theta(m+1-n)\left(\frac{1}{y_n(n)} - 1\right) - \tilde{\mu}(n)\left(\frac{1}{y_n(n)} - \frac{1}{\tilde{y}_n(n)}\right). \end{aligned} \tag{4.38}$$

Now, recall that n was chosen so that $\rho_n \le \rho_{n-1}$. By Lemma 4.4, $y_n(n) \le y_{n-1}(n) = \gamma\kappa + 1$ from its definition in (2.22). Using this and the fact that $m + 1 - n < 0$, from (4.38) we have

$$\begin{aligned} \gamma(\tilde{\rho}_n - \rho_n) &\le a\theta\gamma(m+1-n) + \theta(m+1-n)\left(-\frac{\gamma\kappa}{\gamma\kappa+1}\right) - \tilde{\mu}(n)\left(\frac{1}{y_n(n)} - \frac{1}{\tilde{y}_n(n)}\right) \\ &\le \gamma\theta(m+1-n)\left(a - \frac{\kappa}{\gamma\kappa+1}\right) - \tilde{\mu}(n)\left(\frac{1}{y_n(n)} - \frac{1}{\tilde{y}_n(n)}\right). \end{aligned} \tag{4.39}$$

Because we assumed $a \ge \frac{\kappa}{\gamma\kappa+1}$, from (4.39) it follows that $\tilde{\rho}_n \le \rho_n$ if $y_n(n) \le \tilde{y}_n(n)$ (and similarly, $\tilde{\rho}_n < \rho_n$ if $y_n(n) < \tilde{y}_n(n)$). We now argue that $\tilde{\rho}_n \le \rho_n$ by showing that the alternative implies $y_n(n) < \tilde{y}_n(n)$, which leads to a contradiction since then $\tilde{\rho}_n \le \rho_n$ as per (4.39) and the discussion at the start of the paragraph. As such, we suppose that $\tilde{\rho}_n > \rho_n$. We can apply Lemma 4.4 to the

modified problem for the first m states since the modified Bellman equation is the same up to m and obtain $y_n(i) < \tilde{y}_n(i)$ for $i \leq m+1$. To show that $y_n(i) < \tilde{y}_n(i)$ for $i \geq m+2$, we proceed by induction. For the base case, plugging $i = m+1$ into the optimality equations for the original and modified problem results in

$$\gamma\rho_n = \gamma h(m+1) - \overline{\mu}(m+1) + \overline{\mu}(m+1)\frac{1}{y_n(m+1)} + \lambda y_n(m+2) - \lambda,$$
$$\gamma\tilde{\rho}_n = \gamma\tilde{h}(m+1) - \tilde{\mu}(m+1) + \tilde{\mu}(m+1)\frac{1}{\tilde{y}_n(m+1)} + \lambda\tilde{y}_n(m+2) - \lambda\,.$$

Taking the difference of the above equations and using $\tilde{\rho}_n > \rho_n$, $\tilde{y}_n(m+1) > y_n(m+1)$, and $\overline{\mu}(m+1) - \tilde{\mu}(m+1) = 0$, $h(m+1) - \tilde{h}(m+1) = \delta$, it follows that $\tilde{y}_n(m+2) > y_n(m+2)$. For the induction step, assume $\tilde{y}_n(i) > y_n(i)$ for $m+2 \leq i \leq n-1$. Applying a similar argument to (4.38)-(4.39) in addition to $\tilde{y}_n(i) > y_n(i)$ and $\tilde{\rho}_n > \rho_n$ leads to $\tilde{y}_n(i+1) > y_n(i+1)$. Hence, by induction we have shown $\tilde{y}_n(n) > y_n(n)$. By the discussion in the previous paragraph, this is a contradiction. This proves that ρ_n must be a global minimum of the cost iterates.

From here, using the fact that $\rho_n \leq \rho_k$ for all $k \in S$, we construct a global solution to Lemma 2.2. Define y^* by

$$y^*(i) = \begin{cases} y_n(i), & i \leq n, \\ \overline{y}(i), & i > n, \end{cases}$$

where $\overline{y}(i)$ is defined as the solution to

$$\gamma\rho_n = \gamma h(i) - \overline{\mu}(i) + \frac{\overline{\mu}(i)}{\overline{y}(i)} + \lambda\gamma\kappa, \quad i \geq n+1\,.$$

From here, since $\rho_n \leq \rho_k$ for all $k \geq n+1$, it is straightforward to verify that $\overline{y}(i) \geq \gamma\kappa + 1$ for all $i \geq n+1$, and so (ρ_n, y^*) solve (2.15) for all $i \in S$. Furthermore, by construction it follows that $\limsup_{i\to\infty} y^*(i) = \frac{1}{1-\gamma a}$, and so y^* is uniformly bounded on S. As a result, the pair (ρ_n, y^*) must satisfy Lemma 2.2, and so controlling with a threshold n is optimal among all admissible control policies. □

Remark 4.2. *Note that in the last steps of the above proof where we defined $\overline{y}(i)$, we showed that $\overline{y}(i) \geq \gamma\kappa + 1$ and $\lim_{i\to\infty} \overline{y}(i) = \frac{1}{1-a\gamma}$. Since $a\gamma < \frac{1}{2}$ in Assumption 2.1, this does not contradict the positivity of $\overline{y}$. Furthermore, it is easy to check that $\gamma\kappa + 1 \leq \frac{1}{1-\gamma a}$ is equivalent to $a \geq \frac{\kappa}{\gamma\kappa+1}$, which is the assumed condition in Theorem 2.2.*

Moreover, this highlights the criticality of the assumption $a \geq \frac{\kappa}{\gamma\kappa+1}$ in guaranteeing that we can construct a suitable pair (ρ, y) which satisfies Lemma 2.2. Suppose that there is a level n^* such that $\rho_{n^*} \leq \rho_k$ for all k, and suppose for contradiction that $a < \frac{\kappa}{\gamma\kappa+1}$. Then, if we try to construct a suitable value function $\overline{y}$ in the same manner as above, then we will again end up with $\gamma\kappa + 1 \leq \lim_{i\to\infty} \overline{y}(i) = \frac{1}{1-\gamma a}$, which leads to a contradiction since we assumed $a < \frac{\kappa}{\gamma\kappa+1}$. As a result, we do NOT *necessarily know if the threshold policy controlling with a threshold n^* is optimal among all admissible control policies.*

4.4. **Proof of Theorem 2.3.** For this proof, we study the limiting behavior of the pair (ρ_n, y_n) as $n \to \infty$ to derive a solution pair to (2.15) on S. The following lemma shows that y_n does not vanish at any point in S.

Lemma 4.6. *Suppose $\{\rho_n\}_{n\geq 0}$ is strictly decreasing. Then we have that* $\inf_n y_n(i) > 0 \quad \forall i \geq 1$.

Proof. Clearly ρ_n is bounded from below by the optimal cost in (2.8). Since $\{\rho_n\}_{n\geq 0}$ is strictly decreasing, it is clear that $\rho_n \downarrow \rho$ for some $\rho \geq 0$. To show that y_n does not vanish at any point in S as we take $n \to \infty$, we will first analyze V_n (recall $y_n(i) = \frac{V_n(i)}{V_n(i-1)}$).

Since $V_n > 0$, we can normalize each V_n so that $V_n(0) = 1$ for all n. Additionally, the fact that $\{\rho_n\}_{n\geq 0}$ is strictly decreasing implies that $y_n(i) < \gamma\kappa + 1$ for all $i \leq n$ and for all n. As such, following Remark 4.1, there is a $\delta \in (0,1)$ (independent of n) such that

$$V_n(i) \leq (\gamma\kappa + 1)^i \leq \mathcal{V}(i)^\delta, \quad \forall i \in S\,.$$

As a result, using a standard diagonalization argument, there exists a subsequence of n and a function $V : S \to [0,\infty)$ of order $\mathcal{V}^\delta$ such that $V_n(i) \to V(i)$ for each $i \in S$. Since $V_n(0) = 1$ for every n, it follows that $V(0) = 1$. As a consequence of (ρ_n, V_n) solving (2.18)-(2.19), we have

$$\gamma\rho V(0) = \gamma h(0) + \lambda V(1) - \lambda V(0), \tag{4.40}$$

$$\gamma\rho V(i) = \gamma h(i)V(i) + \lambda V(i+1) + \overline{\mu}(i)V(i-1) - (\overline{\mu}(i) + \lambda)V(i), \quad \forall i \geq 1\,. \tag{4.41}$$

Now, suppose for contradiction that $\inf_n y_n(i_0) = 0$ for some $i_0 \geq 1$. Then there exists some subsequence of n such that $y_n(i_0) \to 0$ as $n \to \infty$. Since $V_n(i_0) = y_n(i_0)V_n(i_0 - 1)$, and $V_n(i_0 - 1)$ is bounded from above for all n, this implies that $V_n(i_0) \to 0$ as $n \to \infty$ (along a further subsequence of n if necessary). Thus, it follows that $V(i_0) = 0$. Writing out (4.41) for $i = i_0$, we have

$$\begin{aligned} &\gamma\rho V(i_0) = \gamma h(i_0)V(i_0) + \lambda V(i_0+1) + \overline{\mu}(i_0)V(i_0-1) - (\overline{\mu}(i_0) + \lambda)V(i_0) \\ &\implies 0 = \lambda V(i_0+1) + \overline{\mu}(i_0)V(i_0-1)\,. \end{aligned}$$

Since $V \geq 0$, this is only possible if $V(i_0+1) = V(i_0-1) = 0$. By proceeding in an inductive manner, we eventually have $V(0) = 0$, which is a contradiction. This gives the desired result. □

Proof of Theorem 2.3. Since $\{\rho_n\}_{n\geq 0}$ is strictly decreasing and bounded from below, there exists $\rho \geq 0$ such that $\rho_n \downarrow \rho$. Additionally, we know from Lemma 4.4 that $y_n(i) < \gamma\kappa + 1$ for all $i \leq n$ and for all n. This, in conjunction with Lemma 4.6, implies that

$$0 < \liminf_{n\to\infty} y_n(i) \leq \gamma\kappa + 1\,.$$

Now, define $y(i) = \liminf_{n\to\infty} y_n(i)$ for each $i \geq 1$. Additionally, recall that (ρ_n, y_n) satisfy the following equation for all $i \leq n$ and $n \geq 1$:

$$\gamma\rho_n = \gamma h(i) - \overline{\mu}(i) + \frac{\overline{\mu}(i)}{y_n(i)} + \lambda(y_n(i+1) - 1)\,.$$

Since $y(i) > 0$ for each i, taking $n \to \infty$ (along an appropriate subsequence) gives us

$$\gamma\rho = \gamma h(i) - \overline{\mu}(i) + \frac{\overline{\mu}(i)}{y(i)} + \lambda(y(i+1) - 1), \quad \forall i \geq 1\,. \tag{4.42}$$

Since $\rho < \rho_n$ for all n, by Lemma 4.4 we must have that $y(i) < y_n(i) < \gamma\kappa + 1$ for all $i \geq 0$. This together with (4.42) implies that the pair (ρ, y) satisfies

$$\gamma\rho = \gamma h(i) - \overline{\mu}(i) + \frac{\overline{\mu}(i)}{y(i)} + \lambda \min(y(i+1) - 1, \gamma\kappa), \quad \forall i \geq 1,$$

with the minimum always being attained at $y(i) - 1$. By Lemma 2.2, we have ρ equals the optimal cost, and the policy which never exercises admission control is optimal. This completes the proof. □

Remark 4.3. *In the case that the sequence $\{\rho_n\}$ is not decreasing, one may still apply an argument similar to the one used in this section coupled with Theorem 2.1 to show that $\rho_n \to \rho_\infty$, the ERSC cost corresponding to the policy which never exercises admission control. First, note that the sequence $\{\rho_n\}_{n\geq 0}$ is bounded (see the last sentence in Remark 2.2). Furthermore, one can apply a similar argument used in the proof of* [10, Lemma 3.11] *to rescale the value functions V_n so that $V_n \leq \mathcal{V}^\delta$ on S. From here, one can use a diagonalization approach to obtain a limiting pair $(\overline{\rho}, \overline{V}) \in \mathbb{R}_+ \times \mathcal{O}^+(\mathcal{V}^\delta)$ (where the strict positivity of $\overline{V}$ follows from the same argument argument in the proof of Lemma 4.6). Defining $\overline{y}(i) \doteq \frac{\overline{V}(i)}{\overline{V}(i-1)}$, it follows that $(\overline{\rho}, \overline{y})$ satisfies (4.42), and the fact that $\overline{\rho} = \rho_\infty$ can be seen from an adaptation of the proof of Theorem 2.1.*

4.5. **Proof of Proposition 2.1.** To begin with, for any $\gamma > 0$ and $n \geq 0$, we let ρ_n^γ represent the ERSC cost corresponding to the sensitivity γ when controlling with a threshold of n, and $y_n^\gamma : \{0, 1, \ldots, n+1\} \to (0, \infty)$ represents the corresponding value function from (2.20)-(2.22). The following characterizes a useful monotonicity result on y_n^γ.

Lemma 4.7. *Let γ_1, γ_2 be two sensitivity parameters such that $\gamma_1 > \gamma_2$, and let n_1, n_2 denote any two non-negative integers. If $\gamma_1 > \gamma_2$ and $\rho_{n_1}^{\gamma_1} > \rho_{n_2}^{\gamma_2}$, then $\frac{1}{\gamma_1}(y_{n_1}^{\gamma_1}(i) - 1) > \frac{1}{\gamma_2}(y_{n_2}^{\gamma_2}(i) - 1)$ for $i = 1, \ldots, \min(n_1, n_2) + 1$. Moreover, the above result holds if " $>$ " is replaced by " $\geq$ ".*

Proof. This proof is nearly identical to the one given in Lemma 4.4 except that we carefully account for the parameters γ_1 and γ_2. We first work under the case that $n_1 = n_2 = n$. In this case, to simplify the notation, we write $\rho(\gamma_k) = \rho_n^{\gamma_k}$ and $y^{\gamma_k}(\cdot) = y_n^{\gamma_k}(\cdot)$ for $k = 1, 2$. Now, as in the statement of the lemma, suppose that $\rho(\gamma_1) > \rho(\gamma_2)$. To show that $y^{\gamma_1}(\cdot) > y^{\gamma_2}(\cdot)$, we proceed by induction.

For the base case, writing out the Bellman equations for $i = 0$ results in

$$\gamma_1 \rho(\gamma_1) = \gamma_1 h(0) + \lambda(y^{\gamma_1}(1) - 1),$$
$$\gamma_2 \rho(\gamma_2) = \gamma_2 h(0) + \lambda(y^{\gamma_2}(1) - 1).$$

Dividing each equation by γ_1 and γ_2 (respectively), we observe that

$$0 < \rho(\gamma_1) - \rho(\gamma_2) = \lambda\left(\frac{1}{\gamma_1}(y^{\gamma_1}(1) - 1) - \frac{1}{\gamma_2}(y^{\gamma_2}(1) - 1)\right),$$

implying that $\frac{1}{\gamma_1}(y^{\gamma_1}(1) - 1) > \frac{1}{\gamma_2}(y^{\gamma_2}(1) - 1)$ as desired. Next, writing out the Bellman equations for $i = 1$ gives us

$$\gamma_1 \rho(\gamma_1) = \gamma_1 h(1) - \overline{\mu}(1) + \frac{\overline{\mu}(1)}{y^{\gamma_1}(1)} + \lambda(y^{\gamma_1}(2) - 1),$$
$$\gamma_2 \rho(\gamma_2) = \gamma_2 h(1) - \overline{\mu}(1) + \frac{\overline{\mu}(1)}{y^{\gamma_2}(1)} + \lambda(y^{\gamma_2}(2) - 1).$$

Dividing each equation by γ_1 and γ_2 (respectively) and subtracting the two equations gives us

$$0 < \overline{\mu}(1)\left[\frac{1}{\gamma_1}\left(\frac{1}{y^{\gamma_1}(1)} - 1\right) - \frac{1}{\gamma_2}\left(\frac{1}{y^{\gamma_2}(1)} - 1\right)\right] + \lambda\left(\frac{1}{\gamma_1}(y^{\gamma_1}(2) - 1) - \frac{1}{\gamma_2}(y^{\gamma_2}(2) - 1)\right). \tag{4.43}$$

Note that

$$\frac{1}{\gamma_1}\left(\frac{1}{y^{\gamma_1}(1)} - 1\right) - \frac{1}{\gamma_2}\left(\frac{1}{y^{\gamma_2}(1)} - 1\right) < 0 \iff \frac{1}{y^{\gamma_1}(1)} - 1 < \frac{\gamma_1}{\gamma_2}\left(\frac{1}{y^{\gamma_2}(1)} - 1\right). \tag{4.44}$$

Because $\frac{\gamma_1}{\gamma_2} > 1$, (4.44) holds if we show that $\frac{1}{y^{\gamma_1}(1)} - 1 < \frac{1}{y^{\gamma_2}(1)} - 1$. However, this is implied by $\frac{1}{\gamma_1} y^{\gamma_1}(1) > \frac{1}{\gamma_2} y^{\gamma_2}(1)$, which holds as a consequence of the base case. So (4.44) holds. But (4.43) implies that $\frac{1}{\gamma_1}(y^{\gamma_1}(2) - 1) > \frac{1}{\gamma_2}(y^{\gamma_2}(2) - 1)$ as desired. The rest of the induction follows in an identical procedure, and so we have $\frac{1}{\gamma_1} y^{\gamma_1}(i) > \frac{1}{\gamma_2} y^{\gamma_2}(i)$ for $i = 1, \ldots, n+1$.

To prove the desired result in the case of two different n_1, n_2 when $\rho_{n_1}^{\gamma_1} > \rho_{n_2}^{\gamma_2}$, note that the cases for $i = 0$ and $i = 1$ are the same, and for the induction step, one would repeat the same procedure but on $i = 1, \ldots, \min(n_1, n_2) + 1$. □

We are now in a position to prove Proposition 2.1.

Proof of Proposition 2.1. For convenience, let $n_1 = n^{\gamma_1}$, and $n_2 = n^{\gamma_2}$. To see why $\rho^{\gamma_1} \geq \rho^{\gamma_2}$ for $\gamma_1 > \gamma_2$, we make use of the following observation: for a random variable Z such that $\mathbb{E}\left[e^{\gamma_1 Z}\right]$ is

finite, we have from Jensen's inequality

$$\frac{1}{\gamma_1 T}\log\mathbb{E}[e^{\gamma_1 Z}] = \frac{1}{\gamma_1 T}\log\mathbb{E}[\big(e^{\gamma_2 Z}\big)^{\frac{\gamma_1}{\gamma_2}}] > \frac{1}{\gamma_1 T}\log\mathbb{E}[\big(e^{\gamma_2 Z}\big)]^{\frac{\gamma_1}{\gamma_2}}$$
$$= \frac{1}{\gamma_1 T}\times\frac{\gamma_1}{\gamma_2}\log\mathbb{E}[e^{\gamma_2 Z}] = \frac{1}{\gamma_2 T}\log\mathbb{E}[e^{\gamma_2 Z}]\,.$$

Applying this to the logarithm of the finite horizon accrued cost, taking $T\to\infty$, and then taking the infimum over control policies on each side of the inequality, we arrive at $\rho^{\gamma_1}\geq\rho^{\gamma_2}$. This verifies the first part of the conclusion of Proposition 2.1.

For the second part, note that it holds automatically if $n_2=\infty$, and so we consider the case when n_2 is finite. Suppose for contradiction that $n_1>n_2$. By assumption, we know that $\rho^{\gamma_1}_{n_1}=\rho^{\gamma_1}\geq\rho^{\gamma_2}=\rho^{\gamma_2}_{n_2}$ since $\gamma_1>\gamma_2$. Also, Lemma 4.7 tells us that $\frac{1}{\gamma_1}(y^{\gamma_1}_{n_1}(i)-1)\geq\frac{1}{\gamma_2}(y^{\gamma_2}_{n_2}(i)-1)$ for $i=1,\dots,n_2+1$. Since $y^{\gamma_2}_{n_2}(n_2+1)=\gamma_2\kappa+1$ from (2.22), it follows that

$$\frac{1}{\gamma_1}(y^{\gamma_1}_{n_1}(n_2+1)-1)\geq\frac{1}{\gamma_2}\kappa\gamma_2=\kappa$$
$$\implies\ y^{\gamma_1}_{n_1}(n_2+1)\geq\gamma_1\kappa+1\,. \tag{4.45}$$

However, this is a contradiction as $n_2<n_1$ and $y^{\gamma_1}_{n_1}(i)<\gamma_1\kappa+1$ for all $i\leq n_1$, as n_1 is the minimal optimal threshold corresponding to γ_1. Hence, this proves $n_1\leq n_2$ as desired. □

4.6. **Proof of Theorem 2.4.**

Proof of Theorem 2.4(i). From the proof of Proposition 2.1 in Section 4.5, we know that $\{\rho^\gamma\}$ is monotonically decreasing as $\gamma\to 0$ and always larger than the risk-neutral cost. Furthermore, since each γ satisfies Assumption 2.1, there exists a $\mathcal{M}>0$ such that $\rho_{avg}\leq\rho^\gamma\leq\mathcal{M}$ for $\gamma<\overline{\gamma}$, and so $\inf_{\gamma>0}\rho^\gamma\geq\rho_{avg}$. We also know by Theorem 2.1 that the ERSC cost is independent of the initial condition $X_0=i$, so for the rest of this proof we fix $X_0=0$ and can remove the infimum over i in the definition of ρ^γ, and similarly for the definition of ρ_{avg} (see [16, Section 3.1]). Furthermore, ρ_{avg} is attained at some threshold policy v_*. For the rest of this section, we let $\rho^\gamma(v)$ denote the ERSC cost corresponding to sensitivity γ and a threshold policy v, and similarly let $\rho_{avg}(v)$ denote the average cost corresponding to a policy v. Showing that $\rho^\gamma(v_*)\to=\rho_{avg}(v^*)=\rho_{avg}$ will complete the proof because $\inf_{\gamma>0}\rho^\gamma(v_*)\geq\inf_{\gamma>0}\rho^\gamma$.

There are two cases: v^* regulates at a finite threshold, or at an infinite threshold. In the case of a finite threshold, the controlled process $\{X(t)\}$ corresponding to v^* takes values in a finite state space. In this case, it follows from [28, Theorems 3.1 and 3.2] that $\rho^\gamma(v_*)$ uniquely solves

$$\mathbb{E}^{v_*}_0\left[e^{\gamma\int_0^{\tau_0}\left(c(X(t),v_*(X(t)))-\rho^\gamma(v_*)\right)dt}\right]=1, \tag{4.46}$$

where τ_0 denotes the first strictly positive return time of the process to state 0. From here, it is straightforward to see the moment generating function of $\int_0^{\tau_0}\big(c(X(t),v_*(X(t)))-\rho^\gamma(v_*)\big)$ exists, and as a consequence we can apply the power series expansion for the exponential, subtract 1 from both sides, and collect the higher order terms in $o(\gamma)$ to obtain

$$\mathbb{E}^{v_*}_0\left[\gamma\int_0^{\tau_0}\big(c(X(t),v_*(X(t)))-\rho^\gamma(v_*)\big)dt\right]+o(\gamma)=0,\ .$$

Dividing by γ on both sides results in

$$\mathbb{E}^{v_*}_0\left[\int_0^{\tau_0}\big(c(X(t),v_*(X(t)))-\rho^\gamma(v_*)\big)dt\right]+\frac{o(\gamma)}{\gamma}=0.$$

Taking $\gamma \to 0$ and pulling it into the expectation and integral (which is possible since $\mathbb{E}_0^{v_*}[\tau_0] < \infty$ by positive recurrence together with the fact that $\rho^\gamma(v_*)$ is bounded), we have

$$\mathbb{E}_0^{v_*}\left[\int_0^{\tau_0} \left(c(X(t), v_*(X(t))) - \lim_{\gamma\to 0} \rho^\gamma(v_*)\right) dt\right] = 0\,. \tag{4.47}$$

However, (4.47) equals 0 uniquely when ρ_{avg} replaces $\lim_{\gamma\to 0}\rho^\gamma(v_*)$. This follows quickly from the renewal reward theorem together with the fact that v^* is assumed to be optimal for the risk-neutral setting. Hence, it follows that $\lim_{\gamma\to 0}\rho^\gamma(v_*) = \rho_{avg}$ as desired. Note that if v^* regulates with an infinite threshold, then the corresponding process is irreducible on S, hence, by [29, Theorem 4.1] we have the following result for all $\gamma < \overline{\gamma}$ due to the uniform stability of our process:

$$\mathbb{E}_0^{v_*}\left[e^{\gamma\int_0^{\tau_0}\left(c(X(t),v_*(X(t)))-\rho^\gamma(v_*)\right)dt}\right] \le 1\,.$$

Since the quantity in the expectation is finite, an application of Dynkin's formula to $\mathbb{E}_0^{v_*}\left[e^{\gamma\int_0^t\left(c(X(s),v_*(X(s)))-\rho^\gamma(v_*)\right)ds}V_\infty(X_t)\right]$ up to the stopping time τ_0 together with the fact that $(\rho^\gamma(v_*), V_\infty)$ satisfy the Bellman equation for V^* results in

$$\mathbb{E}_0^{v_*}\left[e^{\gamma\int_0^{\tau_0}\left(c(X(t),v_*(X(t)))-\rho^\gamma(v_*)\right)dt}V_\infty(0)\right] = V_\infty(0)\,.$$

From here, the same arguments from the case that v^* regulates with a finite threshold apply to give us the desired result. □

To prove Theorem 2.4(ii), we need the following two auxiliary results.

Lemma 4.8. *Let $\mathcal{X} = \{0,1\}^S$, where $S = \{0,1,2,\dots\}$,* i.e., *the set of binary sequences. For any two binary sequences $w_1, w_2 \in \mathcal{X}$, let $n(w_1,w_2) = \max\{n : w_1(i) = w_2(i), \quad \forall i \le n\}$, which we allow to be ∞ if $w_1 = w_2$. Define the metric d on $\mathcal{X}$ by*

$$d(w_1,w_2) = \frac{1}{2^{n(w_1,w_2)}}.$$

Then $\mathcal{X}$ is compact with respect to the metric d.

Proof. This is a well-known result for Cantor spaces; see [14, Section 2.B]. □

Lemma 4.9. *Let $\mathcal{X} = \{0,1\}^S$ be equipped with the metric d from Lemma 4.8. Let $\mathcal{A}$ denote the subset of binary sequences such that there is at most a single 1 that appears. Then $\mathcal{A}$ is closed with respect to the metric d.*

Proof. Let $\{w_k\}_{k=1}^\infty$ be any convergent sequence in $\mathcal{A}$, and suppose that it converges to $w^* \in \mathcal{X}$. Furthermore, let $m(k)$ denote the index in the binary sequence w_k such that $w_k(m(k)) = 1$ (if $w_k(m) \ne 1 \quad \forall m$, then we set this value to be infinity), and similarly let m^* denote the index in w^* where w^* first equals 1. Note there are two cases: ⓐ $m^* = \infty$, or ⓑ $m^* < \infty$.

For ⓐ, this means that $w^* = \{0,0,\dots\}$, which automatically implies that $w^* \in \mathcal{A}$. For ⓑ, since $w_k \to w^*$, we know that $n(w_k, w^*)$ converges to infinity, where $n(w_1,w_2)$ is given in Lemma 4.8. This implies that $m(k) = m^*$ for large enough k. Furthermore, if there was another index $j > m^*$ such that $w^*(j) = 1$, this would mean that $m(k) = j$ for large enough k, as w_k can only be equal to 1 at a single index. This is clearly a contradiction, and thus $w^*(m^*) = 1$ and $w^*(j) = 0 \quad \forall j > m^*$, which implies that $w^* \in \mathcal{A}$. This proves that $\mathcal{A}$ is closed. □

Proof of Theorem 2.4(ii). Let $\{\gamma_k\}_{k\ge 0}$ denote a sequence of sensitivity parameters such that $\gamma_k \downarrow 0$ as $k \to \infty$. Furthermore, let $v_k : S \to \{0,1\}$ denote the control policy that controls at a threshold n^{γ_k}. It is straightforward that we can identify each control policy as an element of $\{0,1\}^S$. Furthermore, since each threshold policy is unique up to the first state where admission control is implemented, then the controls after that state we can represent as simply 0. In other words, if $v \in \mathfrak{U}_{thresh}$

regulates at a threshold of ℓ, then we can set $v = \{0, \dots, 0, 1, 0, \dots\}$, where the 1 occurs at the ℓ-th position in the sequence. Likewise, for a policy $v \in \mathfrak{U}_{thresh}$, recall that we defined $\rho^\gamma(v)$ as the risk-sensitive cost value for sensitivity γ under the policy v, and similarly $\rho_{avg}(v)$ as the risk-neutral cost value under v.

From Lemma 4.8, we know that $\{0,1\}^S$ is a compact space under the topology induced by the metric d that is defined in the statement of the lemma. Furthermore, Lemma 4.9 tells us that $\mathfrak{U}_{thresh}$ is closed as a subset of $\{0,1\}^S$ under this metric, which implies it is also compact. Since $\{v_k\}$ is a sequence in $\mathfrak{U}_{thresh}$, this implies that there exists some $\overline{v} \in \mathfrak{U}_{thresh}$ such that $v_k \to \overline{v}$ along a subsequence. Next, we consider two cases: ① $n_{avg} < \infty$, or ② $n_{avg} = \infty$. We focus on the first case since the second case can be argued similarly. We first suppose for contradiction that $\overline{v}$ regulates at an infinite threshold. Observe that upon showing that $\rho^{\gamma_k}(v_k) \to \rho_{avg}(\overline{v})$ as $k \to \infty$, a contradiction occurs. This is because $\overline{v}$ regulates at an infinite threshold, which is sub-optimal for the risk-neutral setting, and $\rho^{\gamma_k}(v_k) \to \rho_{avg}$. Hence, this would mean that $\rho_{avg}(\overline{v}) = \rho_{avg}$, which is a contradiction.

To show that $\rho^{\gamma_k}(v_k) \to \rho_{avg}(\overline{v})$ as $k \to \infty$, we use a diagonalization argument. We first use the fact that for any fixed γ, $\rho^\gamma(v_k) \to \rho^\gamma(\overline{v})$ as $k \to \infty$. This follows by the fact that $v_k \to \overline{v}$ together with an argument analogous to the policy iteration proof in Theorem 2.3. Next, we use the fact that $\rho^\gamma(\overline{v}) \to \rho_{avg}(\overline{v})$ as $\gamma \to 0$ by part (i).

For any j, k, observe that

$$\rho^{\gamma_j}(v_k) - \rho_{avg}(\overline{v}) \leq |\rho^{\gamma_j}(v_k) - \rho^{\gamma_j}(\overline{v})| + |\rho^{\gamma_j}(\overline{v}) - \rho_{avg}(\overline{v})|.$$

Fix $\varepsilon > 0$ and choose K_1 so that $j \geq K_1$ implies $|\rho^{\gamma_j}(\overline{v}) - \rho_{avg}(\overline{v})| < \frac{\varepsilon}{2}$. Next, for such a j, choose K_2 so that $k \geq K_2$ implies that $|\rho^{\gamma_j}(v_k) - \rho^{\gamma_j}(\overline{v})| < \frac{\varepsilon}{2}$ (which is possible by Remark 4.3). Hence, letting $j = k \geq \max(K_1, K_2)$, we have $\rho^{\gamma_k}(v_k) - \rho_{avg}(\overline{v}) < \varepsilon$ as desired, which gives us the desired convergence. By the discussion in the previous paragraph, we obtain the desired contradiction.

As a result, we must have that $\overline{v}$ regulates at a finite threshold. From the policy convergence, this implies that $v_k = \overline{v}$ for large enough k, so $\rho^{\gamma_k}(v_k) = \rho^{\gamma_k}(\overline{v})$ when k is large enough. Thus, $\rho^{\gamma_k}(\overline{v}) \to \rho_{avg}$, so we have $\rho_{avg}(\overline{v}) = \rho_{avg}$. By the same reasoning as before, this implies that $\overline{v}$ is the optimal policy in the risk-neutral setting as desired.

We now show $n^\gamma \uparrow n_{avg}$. We first consider the case when n_{avg} is finite, and argue that $n^\gamma \leq n_{avg}$ for all $\gamma > 0$. Suppose for contradiction that there exists a γ such that $n^\gamma > n_{avg}$. By the arguments from above, we know that for any sequence $\gamma_k \to 0$, we must have n^{γ_k} converges to n_{avg} as $k \to \infty$. As a result, finiteness of n_{avg} implies $n^{\gamma_k} = n_{avg}$ for all sufficiently small γ_k. But this means that there exists a γ_k smaller than γ where $n^{\gamma_k} = n_{avg} < n^\gamma$. This is a contradiction, as Lemma 4.7 tells us that $n^\gamma \leq n^{\gamma_k}$ for $\gamma_k < \gamma$. Hence, this shows that $n^\gamma \leq n_{avg}$ for all γ, and it follows that $n^\gamma \uparrow n_{avg}$ as $\gamma \to 0$.

For the case that $n_{avg} = \infty$, it immediately follows that $n^\gamma \leq n_{avg}$ for all γ, and the same reasoning as the finite case results in $n^\gamma \uparrow \infty$ as $\gamma \to 0$. □

Acknowledgement

This work is funded by the NSF Grant DMS 2216765.

References

[1] E. Altman, T. Jiménez, and G. Koole. On optimal call admission control in a resource-sharing system. *IEEE Transactions on Communications*, 49(9):1659–1668, 2001.

[2] S. R. Anugu and G. Pang. Ergodic risk sensitive control of large-scale parallel server networks with abandonment. *submitted*, 2025.

[3] S. R. Anugu and G. Pang. Ergodic risk sensitive control of Markovian multiclass many-server queues with abandonment. *The Annals of Applied Probability*, 35(3):1570–1621, 2025.

[4] A. Arapostathis, A. Biswas, and S. Saha. Strict monotonicity of principal eigenvalues of elliptic operators in $\mathbb{R}^d$ and risk-sensitive control. *Journal de Mathématiques Pures et Appliquées*, 124:169–219, 2019.

[5] H. Ayhan. Optimal admission control in queues with abandonments. *Operations Research Letters*, 50(6):712–718, 2022.

[6] N. Bäuerle, M. Pitera, and L. Stettner. Blackwell optimality and policy stability for long-run risk-sensitive stochastic control. *SIAM Journal on Control and Optimization*, 62(6):3172–3194, 2024.

[7] M. Benelli and R. Hassin. Rational joining behavior in a queueing system with abandonments. *Operations Research Letters*, 49(3):426–430, 2021.

[8] S. Benjaafar, J.-P. Gayon, and S. Tepe. Optimal control of a production–inventory system with customer impatience. *Operations Research Letters*, 38(4):267–272, 2010.

[9] A. Biswas and V. S. Borkar. Ergodic risk-sensitive control—a survey. *Annual Reviews in Control*, 55:118–141, 2023.

[10] A. Biswas and S. Pradhan. Ergodic risk-sensitive control of Markov processes on countable state space revisited. *ESAIM: Control, Optimisation and Calculus of Variations*, 28, 2021.

[11] N. Bäuerle and A. Jaśkiewicz. Markov decision processes with risk-sensitive criteria: An overview. *Mathematical Methods of Operations Research*, 99:141–178, 2024.

[12] A. P. Ghosh and A. P. Weerasinghe. Optimal buffer size and dynamic rate control for a queueing system with impatient customers in heavy traffic. *Stochastic Processes and their Applications*, 120(11):2103–2141, 2010.

[13] X. Guo and Y. Huang. Risk-sensitive average continuous-time Markov decision processes with unbounded transition and cost rates. *Journal of Applied Probability*, 58(2):523–550, 2021.

[14] A. S. Kechris. *Classical Descriptive Set Theory*. Springer New York, NY, 1st edition, 1995.

[15] N. C. Knudsen. Individual and social optimization in a multiserver queue with a general cost-benefit structure. *Econometrica*, 40(3):515–528, 1972.

[16] Y. L. Koçağa and A. Ward. Admission control for a multi-server queue with abandonment. *Queueing Systems*, 65:275–323, 2010.

[17] M. Larrañaga, O. J. Boxma, R. Núñez-Queija, and M. S. Squillante. Efficient content delivery in the presence of impatient jobs. In *2015 27th International Teletraffic Congress*, pages 73–81, 2015.

[18] M. E. Lewis, H. Ayhan, and R. D. Foley. Bias optimality in a queue with admission control. *Probability in the Engineering and Informational Sciences*, 13(3):309–327, 1999.

[19] M. E. Lewis, H. Ayhan, and R. D. Foley. Bias optimal admission control policies for a multiclass nonstationary queueing system. *Journal of Applied Probability*, 39:20 – 37, 2002.

[20] X. Liu and A. Weerasinghe. Admission control for double-ended queues. *arXiv preprint*, 2022.

[21] P. Naor. The regulation of queue size by levying tolls. *Econometrica*, 37(1):15–24, 1969.

[22] S. Pradhan. Risk-sensitive ergodic control of reflected diffusion processes in orthant. *Applied Mathematics and Optimization*, 83:1739–1764, 2021.

[23] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., USA, 1st edition, 1994.

[24] R. Singh, X. Guo, and E. Modiano. Risk-sensitive optimal control of queues. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 3563–3568. IEEE, 2017.

[25] S. Stidham. Optimal control of admission to a queueing system. *IEEE Transactions on Automatic Control*, 30(8):705–713, 1985.

[26] A. R. Ward and S. Kumar. Asymptotically optimal admission control of a queue with impatient customers. *Mathematics of Operations Research*, 33(1):167–202, 2008.

[27] A. Weerasinghe and A. Mandelbaum. Abandonment versus blocking in many-server queues: Asymptotic optimality in the qed regime. *Queueing Systems*, 75:279–337, 2013.

[28] Q. Wei and X. Chen. Continuous-time Markov decision processes under the risk-sensitive average cost criterion. *Operations Research Letters*, 44(4):457–462, 2016.

[29] Q. Wei and X. Chen. Risk-sensitive average continuous-time Markov decision processes with unbounded rates. *Optimization*, 68(4):773–800, 2019.

[30] W. Whitt. What you should know about queueing models to set staffing requirements in service systems. *Naval Research Logistics*, 54:476 – 484, 08 2007.

[31] R. Wu and H. Ayhan. Optimal admission control in queues with multiple customer classes and abandonments. *Queueing Systems*, 109(1), 2025.

[32] Z. Wu and R. Xu. Risk-sensitive Markov decision process and learning under general utility functions. *arXiv*, 2311.13589, 2024.

[33] S. Yoon and M. E. Lewis. Optimal pricing and admission control in a queueing system with periodically varying parameters. *Queueing Systems*, 47(3):177–199, 2004.

[34] G. Zayas-Cabán and M. E. Lewis. Admission control in a two-class loss system with periodically varying parameters and abandonments. *Queueing Systems*, 94(1–2):175–210, 2020.