

ASYMPTOTIC OPTIMALITY OF THE BINOMIAL-EXHAUSTIVE POLICY FOR POLLING SYSTEMS WITH LARGE SWITCHOVER TIMES

BY YUE HU*, JING DONG[†] AND OHAD PERRY[‡]

**Decision, Risk, and Operations, Columbia Business School, yhu22@gsb.columbia.edu*

[†]*Decision, Risk, and Operations, Columbia Business School, jing.dong@gsb.columbia.edu*

[‡]*Department of Industrial Engineering and Management Sciences, Northwestern University, ohad.perry@northwestern.edu*

We study an optimal-control problem of polling systems with large switchover times, when a holding cost is incurred on the queues. In particular, we consider a stochastic network with a single server that switches between several buffers (queues) according to a pre-specified order, assuming that the switchover times between the queues are large relative to the processing times of individual jobs. Due to its complexity, computing an optimal control for such a system is prohibitive, and so we instead search for an asymptotically optimal control. To this end, we first solve an optimal control problem for a deterministic relaxation (namely, for a fluid model), that is represented as a hybrid dynamical system. We then “translate” the solution to that fluid problem to a binomial-exhaustive policy for the underlying stochastic system, and prove that this policy is asymptotically optimal in a large-switchover-time scaling regime, provided a certain uniform integrability (UI) condition holds. Finally, we demonstrate that the aforementioned UI condition holds in the following cases: (i) the holding cost has (at most) linear growth, and all service times have finite second moments; (ii) the holding cost grows at most at a polynomial rate (of any degree), and the service-time distributions possess finite moment generating functions.

1. Introduction. A polling system is a queueing network in which a single server attends multiple queues according to a pre-specified routing mechanism. This class of models has been extensively studied since the 1950’s, starting with [Mack et al. \(1957\)](#), and have been employed in numerous application settings, such as computer-communication ([Bux, 1981](#)), production ([Federgruen and Katalan, 1996](#); [Olsen, 1999](#)), inventory-control ([Winands et al., 2011](#)), transportation ([Altman et al., 1992](#); [Van den Broek et al., 2006](#)), and health-care ([Cicin-Sain et al., 2001](#); [Vlasiou et al., 2009](#)). We refer to [Levy and Sidi \(1990\)](#); [Takagi \(1997\)](#); [Vishnevskii and Semenova \(2006\)](#); [Boon et al. \(2011a\)](#); [Borst and Boxma \(2018\)](#) for comprehensive reviews of the relevant literature.

Exact analysis of polling systems is in general prohibitively hard; [Resing \(1993\)](#) argues that, unless the switching policy has a certain branching property, the system is not amenable to exact analysis. However, even for those “branching-type” policies, results are typically expressed via multi-dimensional transforms that can be hard to analyze. Thus, despite being among the most extensively studied class of stochastic networks ([Boon et al., 2011b](#)), little is known about how to optimally control polling systems, except for special cases, such as the two-queue system in [Hofri and Ross \(1987\)](#), a symmetric cost structure in [Levy et al. \(1990\)](#); [Liu et al. \(1992\)](#), or a limited control problem that is solved for a subset of the queues in [Duenyas and Van Oyen \(1996\)](#); [Van der Mei and Levy \(1997\)](#); [Matveev et al. \(2016\)](#). As will be seen below, due to the dimensionality and the switching dynamics of the queue process, finding an optimal control is a difficult problem even for deterministic polling systems.

MSC 2010 subject classifications: Primary 60K25, 60F17, 90B15; secondary 37A50

Keywords and phrases: Polling system, hybrid dynamical system, fluid-control problem, asymptotic optimality

Scalings of the Switchover Times. To achieve analytical simplification, it is sometimes assumed that the server’s switchover times are instantaneous. This assumption is reasonable to make when those switchover times are sufficiently small relative to the service times, *and* the total traffic intensity is not too close to 1. (If the system is nearly critical, then even small deviations from the “ideal” modeling assumptions can have substantial negative impacts on the performance; see the discussion in (Perry and Whitt, 2016, Section 9).) However, switchover times are often quite large, and sometimes can even be considered to be of a *larger order of magnitude* than the service durations, see, e.g., Nahmias and Rothkopf (1984); Federgruen and Katalan (1994); Olsen (2001); Winands et al. (2011). In such cases, one can turn the analytical disadvantage of having switchover times into an advantage by taking limits as those switchover times increase without bound. (This approach is analogous to the one in which the switchover times are assumed to be instantaneous, which is in turn tantamount to taking limits as the switchover times decrease to 0.) This large-switchover-times asymptotic approach was taken in Olsen (2001); Van der Mei (1999); Winands (2007, 2011) to approximate stationary performance measures, and was identified as an important future direction in Boon et al. (2011a); Borst and Boxma (2018). The same approach is taken here to solve an optimal-control problem in an appropriate asymptotic sense that will be explained below.

1.1. *Optimal Control of Stochastic Networks.* A standard approach in the stochastic-network literature to solving optimal-control problems follows an asymptotic scheme that was first proposed by Harrison (1988). This approach can be roughly summarized as follows: (I) formulate and solve a Brownian control problem (BCP) inspired by a heavy-traffic approximation for the system; (II) “translate” the resulting optimal Brownian control to a control for the stochastic system; and (III) prove that the control for the system is asymptotically optimal in an appropriate sense.

In this paper, we follow the main line or reasoning of the above scheme, but with important differences. First, instead of a BCP, we solve a fluid-control problem (FCP) related to the stochastic control problem. Moreover, an important step in solving the BCP in Harrison’s scheme involves solving an *equivalent workload formulation* that is rigorously achieved for the controlled stochastic system by showing that *state-space collapse* (SSC) holds asymptotically, namely, that the limit diffusion process is confined to a subspace having a lower dimension than that of the prelimit; see, e.g., Harrison et al. (1997). Under fluid scaling, SSC corresponds to sliding motion of the fluid limit on a lower-dimensional manifold, as explained in (Perry and Whitt, 2016, Section 1). For a specific example, see Atar et al. (2011), which considers the problem of asymptotically minimizing long-run average costs in an overloaded many-server fluid regime. (There is customer abandonment with rate θ_i in queue i , keeping the system stable despite being overloaded.) The proposed $c\mu/\theta$ priority rule induces asymptotic SSC in the stationary fluid model, because all the queues that receive service, except the ones with the smallest $c\mu_i/\theta_i$ parameters, are asymptotically null in stationarity; see Equation (18) in this reference. However, SSC does not occur in our setting, because all the queues increase at fixed rates (their respective arrival rates) during the switchover times, which are non-negligible in the fluid time scale. In particular, under the large-switchover-times asymptotic, there is no reduction in the dimensionality of the limiting process, implying that the dynamics of the fluid limits are necessarily discontinuous. (In fact, the fluid limits may not even exhibit continuous dependence on their initial condition, and so do not adhere to the classical definition of well-posed dynamical systems.)

We remark that asymptotic SSC can occur in polling systems if the switchover times are small relative to the time scaling used to derive the limiting process. Such SSC following from the averaging principles are proved in Coffman Jr et al. (1995) for polling systems with zero switchover times, and in Coffman Jr et al. (1998) for systems with positive switchover times that do not scale in the limit (i.e., are negligible asymptotically).

The Optimal-Control Problem. We consider the (asymptotically) optimal-control problem for a polling system in which the server moves among the different queues in a fixed order that is specified by a table, with the objective of minimizing a long-run average holding cost on the queue process. In this setting, a dynamic control is a state-dependent server-routing policy which determines when the server should switch away from its current queue to the next queue in the table.

Since solving the optimal-control problem is prohibitively hard, we seek a control that is optimal in an appropriate asymptotic sense. To this end, we consider a sequence of systems under a functional weak law of large numbers (FWLLN) scaling, and analyze the resulting fluid limits as solutions to a hybrid dynamical system (HDS).¹ We then identify an optimal fluid control for the HDS, which we “translate” to a control for the underlying polling system. In particular, the control we propose is the well-known binomial-exhaustive policy whose specific control parameters are taken directly from the optimal fluid control; see Sections 2 and 5.2 for more details. Finally, we prove that the binomial-exhaustive policy (with the fluid-optimal control parameters) is asymptotically optimal under the fluid scaling, in that it asymptotically achieves a lower bound on the long-run average cost.

The proof of asymptotic optimality requires that the sequence of cumulative holding costs incurred over a table cycle in stationarity is uniformly integrable (UI). We demonstrate that the UI condition holds in two important cases: (i) when the holding cost is linear and the service times have finite second moments; and (ii) when the holding cost grows at most at a polynomial rate (of any degree), and all service times have finite moment generating functions (m.g.f.).

1.2. Conventions About Notation. All the random variables and processes are defined on a single probability space (Ω, \mathcal{F}, P) . We write \mathbb{E} to denote the expectation operator, and $\mathbb{E}_{\mathbb{P}}$ when we want to emphasize that the expectation is with respect to a specific probability measure \mathbb{P} . We let \mathbb{R} , \mathbb{Z} and \mathbb{N} denote the sets of real numbers, integers and strictly positive integers, respectively, $\mathbb{Z}_+ := \mathbb{N} \cup \{0\}$, and $\mathbb{R}_+ := [0, \infty)$. For $k \in \mathbb{N}$, we let \mathbb{R}^k denote the space of k -dimensional vectors with real components, and denote these vectors with bold letters and numbers; in particular, we write $\mathbf{1} := (1, \dots, 1)$ for the vector of 1’s. We let D^k denote the space of right-continuous \mathbb{R}^k -valued functions (on arbitrary finite time intervals) with limits everywhere, endowed with the usual Skorokhod J_1 topology; see Chapter 11 of Whitt (2002). We let $D := D^1$. We use C^k (and $C := C^1$) to denote the subspace of D^k of continuous functions. It is well-known that the J_1 topology relativized to C^k coincides with the uniform topology on C^k , which is induced by the norm

$$\|x\|_t := \sup_{0 \leq u \leq t} \|x(u)\|,$$

where $\|x\|$ denotes the usual Euclidean norm of $x \in \mathbb{R}^k$. We use “ \Rightarrow ” to denote weak convergence of random variables in \mathbb{R}^k , and of stochastic processes over compact time intervals.

For $f : \mathbb{R}^k \rightarrow [0, \infty)$, $g : \mathbb{R}^k \rightarrow [0, \infty)$ and $a \in \mathbb{R}_+^k \cup \{\infty\}$ we write $f(x) = O(g(x))$ as $x \rightarrow a$ if $\limsup_{x \rightarrow a} f(x)/g(x) < \infty$, and $f(x) = o(g(x))$ if $\lim_{x \rightarrow a} f(x)/g(x) = 0$.

Given a sequence of random variables $\{X^n : n \geq 1\}$ and a sequence of non-negative real numbers $\{a^n : n \geq 1\}$, we write $X^n = O_p(a^n)$ if $\|X^n\|/a^n$ is stochastically bounded, i.e., for any $\epsilon > 0$, there exist finite $M, N \in \mathbb{N}$ such that $\mathbb{P}(\|X^n\|/a^n > M) < \epsilon$ for all $n \geq N$. We write $X^n = o_p(a^n)$ if $\|X^n\|/a^n$ converges to zero in probability, and $X^n = \Theta_p(a^n)$ if X^n is $O_p(a^n)$ but not $o_p(a^n)$. We write that a sequence of stochastic processes $\{X^n : n \geq 1\}$ is $O_p(a^n)$, $o_p(a^n)$, and $\Theta_p(a^n)$ if the corresponding property holds for $\|X^n\|_t$ for any $t \in (0, \infty)$.

¹We use the acronym HDS for both singular and plural forms (system and systems).

For $x, y \in \mathbb{R}$, we write $x \wedge y := \max\{x, y\}$, $x \vee y := \min\{x, y\}$, and $x^+ := \max\{x, 0\}$. For a function $x \in D$, $x(a-)$ denotes the left-hand limit at a , i.e., $x(a-) := \lim_{t \uparrow a} x(t)$ is the left-hand limit at the point a . For a vector $\mathbf{v} \in \mathbb{R}^\ell$, $\ell \in \mathbb{N}$, we use $\dim(\mathbf{v})$ to denote the dimension of \mathbf{v} ; namely, $\dim(\mathbf{v}) = \ell$.

We use a “bar” to denote fluid-scale quantities: $\bar{X}^n := X^n/n$ for a sequence of random variables $\{X^n : n \geq 1\}$, and $\bar{X}^n(t) := X^n(nt)/n$, $t \geq 0$, if the $\{X^n : n \geq 1\}$ is a sequence of processes.

1.3. Organization of the Paper. The rest of the paper is organized as follows. In Section 2 we introduce the model, the main results, and a roadmap for our approach to proving those result. In Section 3 we consider a deterministic relaxation (a fluid model) to the optimal control problem, which is characterized as the set of solutions to an HDS. It is also shown that the fluid model is related to the sequence of stochastic systems via functional weak laws. In Section 4 we analyze the FCP, propose an optimal fluid control, and establish important qualitative properties of the fluid model under this control. In Section 5 we relate the proposed fluid control to the binomial-exhaustive policy. In Section 6 we prove that the binomial-exhaustive policy is asymptotically optimal in our setting. Section 7 is dedicated to the proofs of the main theorems. We conclude in Section 8. Complementary proofs appear in the appendix.

2. Problem Formulation and Main Results. We consider a polling system with K queues numbered $1, \dots, K$. Customers (or jobs) arrive at queue $k \in \mathcal{K} := \{1, \dots, K\}$ according to a Poisson process with rate $\lambda_k > 0$, and wait for their turn to be served in a buffer with infinite capacity (so that no customers are blocked). We refer to customers who arrive to queue k as class- k customers. The service times for class- k customers are independent and identically distributed (i.i.d.) random variables with mean $1/\mu_k < \infty$. At this point, we do not impose any other assumptions on the service time distributions (other than assuming that they all have finite means), but further assumptions on the existence of higher moments will be needed to prove the aforementioned UI condition. We denote by S_k a generic random variable that has the service time distribution of class- k customers.

A single server visits the queues periodically according to a fixed order specified by a *table*. In particular, the table consists of I stages, $I \geq K$, and the queue to be served at each stage is defined by a polling function $p : \mathcal{I} \rightarrow \mathcal{K}$, for $\mathcal{I} := \{1, \dots, I\}$, where $p(i)$ is the queue attended (polled) by the server at stage i , and $(p(i), i \in \mathcal{I})$ is the table. Note that a queue may appear more than once in a table, in which case that queue is attended by the server in two or more nonconsecutive stages. We refer to each such attendance as a *visit* (of the server to the queue). We refer to the starting time of a visit as a polling epoch and the ending time of a visit as a departure epoch (of the server from the queue). A *table cycle* is the time elapsed between two consecutive polling epochs of stage 1 (the first visit to queue $p(1)$) in \mathcal{I} . The table is said to be *cyclic* if all the queues appear in the table exactly once (so that the server visits each queue exactly once in a table cycle), in which case $\mathcal{K} = \mathcal{I}$.

For $i \in \mathcal{I}$, we assume that the switchover time of the server from stage i to stage $i + 1$ is a random variable V_i with mean $s_i := \mathbb{E}[V_i] < \infty$, that is independent of all other random variables and processes in the system. We let $s := \sum_{i \in \mathcal{I}} s_i$ denote the total expected switchover time incurred within a table cycle, and assume that $s > 0$.

For a given table, the *switching policy* (the control) is the set of rules specifying when the server should switch from each stage to the next. Note that if a queue is visited more than once in the table, then the control may prescribe a different switching rule for each visit. In addition, we allow the switching policy to induce an *augmented table* in which the queues appear in a periodic pattern that is an L multiple of the pattern of the basic table, for some

integer $L \geq 2$. We refer to a switching policy inducing an L -cycle augmented table as an L -cycle control, and denote the set of stages in that augmented table by $\mathcal{I}^L := \{1, \dots, IL\}$. We refer to the original table $(p(i), i \in \mathcal{I})$ as the *basic table* and to a corresponding control, whose switching rules are repeated after I stages, as a one-cycle control.

A *server cycle* is the time elapsed between two consecutive polling epochs of stage 1 in \mathcal{I}^L . Thus, a server cycle in an L -cycle control consists of L table cycles (and the server cycle is equal to the table cycle under a one-cycle control). Under an L -cycle control, an ℓ th table cycle is time elapsed between stage $1 + (\ell - 1)I$ and stage $1 + \ell I$ in the augmented table, for $1 \leq \ell \leq L$. (That is, the time it takes the server to complete the ℓ th basic table within the augmented table.) Further, we let the polling function $p : \mathcal{I}^L \rightarrow \mathcal{K}$ map a stage in the augmented table to the queue being visited at that stage. The corresponding augmented table is given by $(p(i), i \in \mathcal{I}^L)$. The expected total switchover time in a corresponding server cycle is then sL .

For concreteness, consider a polling system with three queues ($K = 3$) visited according to the basic table $(1, 2, 3, 2, 3)$. The basic table contains five stages ($I = 5$): queue 2 is visited in stages 2 and 4, and queue 3 is visited in stages 3 and 5. Hence, $p(1) = 1$, $p(2) = p(4) = 2$ and $p(3) = p(5) = 3$. Under a one-cycle control ($L=1$), the basic and L -cycle augmented table, as well as table and server cycles, are all equivalent notions. In contrast, under a two-cycle control ($L = 2$), the augmented table is given by $(1, 2, 3, 2, 3; 1, 2, 3, 2, 3)$, so that queue 1 is visited twice, and queues 2 and 3 are each visited four times during a server cycle, which now consists of two table cycles. Of course, by a two-cycle control we mean that the switching rule of queue $p(i)$ is different than the rule of queue $p(i + 5)$, for at least one $i \in \{1, \dots, 5\}$.

REMARK 2.1. The term ‘‘basic table’’ typically suggests that the order at which the server visits the queues has no repeated pattern. It may therefore seem artificial to consider augmented tables with $L \geq 2$ consecutive repetitions of the basic table. However, one cannot rule out at the outset the possibility that an L -cycle control, for some $L > 1$, is better (reduces the cost) than a one-cycle control. Further, by considering L -cycle controls we can prove in some important special cases that the asymptotically optimal one-cycle control is the overall asymptotically optimal control.

Let $\rho_k := \lambda_k / \mu_k$ denote the traffic intensity corresponding to queue k , and let

$$\rho := \sum_{k \in \mathcal{K}} \lambda_k / \mu_k.$$

We assume that $\rho < 1$, so that the system can be stabilized in the sense that there exist service policies under which the queue process admits a stationary distribution (Fricker and Jaibi, 1994; Boon et al., 2011b). We note that the system is stable under the binomial-exhaustive policy if and only if $\rho < 1$.

Let $U^{(0)} := 0$, and for $m \geq 1$, let $U^{(m-1)}$ denote the beginning of the m th server cycle (end of the $(m - 1)$ st server cycle), where without loss of generality, we take time 0 to be a polling epoch of stage 1. Let $A_i^{(m)}$ and $D_i^{(m)}$, $i \in \mathcal{I}^L$, denote the polling and departure epochs of stage i during the m th server cycle. Then, $B_i^{(m)} := D_i^{(m)} - A_i^{(m)}$ is the *busy time* at stage i in the m th server cycle. Lastly, let $T^{(m)}$ be the length of the m th server cycle, i.e., $T^{(m)} = U^{(m)} - U^{(m-1)}$ and $T^{(m)} = \sum_{i \in \mathcal{I}^L} (B_i^{(m)} + V_i^{(m)})$, where $V_i^{(m)} \stackrel{d}{=} V_i$. Under a given switching policy π , we denote by $Q_{\pi,k}(t)$ the number of customers in queue k at time t , $k \in \mathcal{K}$, and let $Q_\pi(t) := (Q_{\pi,k}(t), k \in \mathcal{K}), t \geq 0$.

The Optimal Control Problem. Let $\psi : \mathbb{R}_+^K \rightarrow \mathbb{R}_+$ denote the holding cost, so that $\psi(Q(t))$ is the cost incurred at time t when the state of the queue is $Q(t)$. We assume that ψ is non-negative, non-decreasing and continuous. Our goal is to find an asymptotically optimal control π within a family Π of *admissible controls* (see Definition 2.1 below and Section 5), that minimizes the following expected long-run average costs

$$(2.1) \quad \liminf_{t \rightarrow \infty} \frac{1}{t} \mathbb{E} \left[\int_0^t \psi(Q_\pi(u)) du \right] \quad \text{and} \quad \limsup_{t \rightarrow \infty} \frac{1}{t} \mathbb{E} \left[\int_0^t \psi(Q_\pi(u)) du \right].$$

Let

$$(2.2) \quad \tilde{Q}_\pi(m) := Q_\pi(U^{(m)}), \quad m \geq 0.$$

The service policies we consider are state-dependent controls that may depend on the value of $\tilde{Q}_\pi(m)$, such that the process $\tilde{Q}_\pi := \{\tilde{Q}_\pi(m) : m \geq 0\}$ is a discrete-time Markov chain (DTMC) (see Lemma 5.1 below), and is therefore regenerative whenever it is positive recurrent (as must be the case under an optimal policy). We refer to this DTMC as *the embedded DTMC*, and remark that, under the asymptotically optimal control we propose, namely, under the binomial-exhaustive policy, the embedded DTMC is ergodic; see Fricker and Jaibi (1994).

The Family of Admissible Controls. We say that a switching control is non-idling if the server does not idle while attending a non-empty queue, and in addition, it switches immediately to the next queue in the table if it empties the attended buffer. On the other hand, if the server finds a buffer empty upon its polling epoch, we allow it to wait for work to arrive. This latter event is asymptotically null, because the server always finds a queue upon arrival to a buffer in the fluid limits, and thus has no impact on our asymptotic analysis. Let $\{\mathcal{F}_t : t \geq 0\}$ denote the σ -algebra generated by the queue process.

DEFINITION 2.1 (admissible control). *A switching control is admissible if*

- (i) *The policy is non-idling.*
- (ii) *For $i \in \mathcal{I}^L$, $m \geq 1$, the number of customers served during the busy time at the i th stage in the m th server cycle conditional on $Q(A_i^{(m)})$ is independent of $\mathcal{F}_{A_i^{(m)}}$.*
- (iii) *The policy is non-anticipative.*

It is significant that the set of admissible controls contains a wide range of controls studied in the literature. For example, the family of branching-type controls, which includes the exhaustive, gated, binomial-exhaustive, binomial-gated, and Bernoulli-type policies (Resing, 1993; Levy, 1988, 1989), limited-type policies (Boxma, 1986; Szpankowski and Rego, 1987), and base-stock policies (Federguen and Katalan, 1996) are all admissible. More generally, all the policies studied in Fricker and Jaibi (1994) are admissible. In particular, the policies considered in Fricker and Jaibi (1994) satisfy the three requirements in Definition 2.1, in addition to a certain stochastic-monotonicity condition that we do not impose; see Section 2 in this reference.

The Binomial-Exhaustive Policy. Let $Q_{p(i)}(A_i^{(m)})$ denote the number of customers in queue $p(i)$ upon its polling epoch in the m th server cycle, $i \in \mathcal{I}^L$, $m \geq 1$. The binomial-exhaustive policy, which was proposed in Levy (1988), is fully characterized by two parameters: an integer L , that specifies the number of table cycles contained in a server cycle, and a vector $\mathbf{r} = (r_1, \dots, r_{IL}) \in [0, 1]^{IL}$, whose component r_i is the “success probability” of the binomial random variable corresponding to stage $i \in \mathcal{I}^L$. Note that, if $\sum_{\{i \in \mathcal{I}^L : p(i)=k\}} r_i = 0$, then

queue k explodes since it never gets served. We therefore consider \mathbf{r} to be an element in the set

$$(2.3) \quad \mathcal{R} := \left\{ \mathbf{r} \in [0, 1]^{IL} : \sum_{\{i \in \mathcal{I}^L : p(i) = k\}} r_i > 0 \text{ for all } k \in \mathcal{K} \right\}.$$

DEFINITION 2.2 (binomial-exhaustive policy). *For $(L, \mathbf{r}) \in \mathbb{N} \times \mathcal{R}$, $i \in \mathcal{I}^L$ and $m \geq 1$, conditional on the event $\{Q_{p(i)}(A_i^{(m)}) = N\}$, $N \in \mathbb{Z}_+$, the number of customers that the server leaves behind at the departure epoch of queue $p(i)$ is $N - Y_i^{(m)}(N, r_i)$, where $Y_i^{(m)}(N, r_i)$ is a binomial random variable with parameters N and r_i , which is independent of all other random variables and processes.*

The binomial-exhaustive policy can equivalently be described as one in which the server performs an independent Bernoulli trial for each customer in the queue at stage i , having “success probability” r_i . If the outcome of that trial is a “success,” then the server serves that customer as well as all of the new arrivals during the service duration of that customer. Thus, the server attends the queue polled at stage i in the m th server cycle for $Y_i^{(m)}(N, r_i)$ busy periods of an $M/G/1$ queue having arrival rate $\lambda_{p(i)}$ and service rate $\mu_{p(i)}$.

2.1. The Large-Switchover-Time Asymptotic Regime. To carry out our asymptotic analysis, we consider a sequence of systems indexed by $n \geq 1$, and append a superscript n to all random variables and processes that scale with n . Let V_i^n denote the switchover time from stage i in system n . Under the large-switchover-time scaling, we keep λ_k and μ_k fixed (they do not scale with n), and impose the following assumptions on the sequence of switchover times.

ASSUMPTION 1. $\bar{V}_i^n := V_i^n/n \Rightarrow s_i$ as $n \rightarrow \infty$. Further, $\mathbb{E}[V_i^n] = ns_i$ for all $i \in \mathcal{I}$.

We make two remarks: First, we allow $V_i^n = o_p(n)$ for some, but not all, $i \in \mathcal{I}$, so that $s_i = 0$, but $s > 0$. Second, the latter part of Assumption 1 can be easily relaxed to $\mathbb{E}[V_i^n]/n \rightarrow s_i$ as $n \rightarrow \infty$. However, this relaxation comes at the expense of more cumbersome notation in some proofs, and has no practical significance (for the actual stochastic system under consideration).

Under the large-switchover-time scaling, the server spends $\Theta_p(n)$ time switching, so that the queues at polling epochs are also of order $\Theta_p(n)$, namely, the queue process is strictly positive in fluid scale. Recall that fluid-scaled quantities (random variables, processes and parameters) are denoted with a bar, e.g., $\bar{Q}_{\pi^n}^n(t) := Q_{\pi^n}^n(nt)/n$. Let

$$\bar{C}_{\pi^n}^n(t) := \frac{1}{t} \int_0^t \psi(\bar{Q}_{\pi^n}^n(u)) du,$$

where π^n is the control employed in the n th system (and is allowed to depend on n).

We say that a sequence of controls $\tilde{\pi}_* := \{\tilde{\pi}_*^n : n \geq 1\}$ is asymptotically optimal if

$$(2.4) \quad \limsup_{n \rightarrow \infty} \limsup_{t \rightarrow \infty} \mathbb{E}[\bar{C}_{\tilde{\pi}_*^n}^n(t)] \leq \liminf_{n \rightarrow \infty} \liminf_{t \rightarrow \infty} \mathbb{E}[\bar{C}_{\pi^n}^n(t)],$$

for any other sequence of admissible controls $\pi := \{\pi^n : n \geq 1\}$.

REMARK 2.2. Since we seek an effective control for a given stochastic system, the sequence $\tilde{\pi}_*$ of asymptotically optimal controls should be considered to be a single control

whose parameters may depend on n . For example, if a threshold-type control is exercised, then the control parameters (the thresholds) must increase linearly with n in order to appear in the fluid limits. Hence, there is no abuse of terminology in saying that a control (as opposed to a sequence of controls) is asymptotically optimal. On the other hand, the elements of π are allowed to change arbitrarily with n .

2.2. Summary of Main Results. Our main result establishes that the binomial-exhaustive policy, with properly selected parameters (L_*, \mathbf{r}_*) , is asymptotically optimal (among the set of admissible controls) for a large family of cost functions ψ . The specific control parameters (L_*, \mathbf{r}_*) are computed by solving a corresponding FCP, as will be explained below, and are referred to as the *optimal (control) parameters*. We thus denote the sequence of binomial-exhaustive policies by π_* . Note that the same control parameters (L_*, \mathbf{r}_*) are used for all $n \geq 1$; in particular, the same control is considered for all the systems along the sequence. This property of the asymptotically optimal control we propose is attractive, because applying the control in a given system can be done directly, without any engineering considerations which are often needed in order to determine the size of the control parameters for a specific system.

To formally state our main result, let $\bar{T}^n := T^n/n$, where T^n is the length of the *stationary* server cycle in the n th system, which is finite w.p.1 when the embedded DTMC is positive recurrent. For each $n \geq 1$ and control π^n , let

$$(2.5) \quad \bar{\Psi}_{\pi^n}^n := \int_0^{\bar{T}^n} \psi(\bar{Q}_{\pi^n}^n(u)) du$$

denote the cumulative fluid-scaled cost under π^n over a stationary server cycle, namely, when $\bar{Q}_{\pi^n}^n(0)$ is distributed according to a stationary distribution of the embedded DTMC. Let (L_*, \mathbf{r}_*) be the optimal FCP parameters, and c_* be the optimal objective value of the FCP. In addition, let π_*^n be the binomial-exhaustive policy with these parameters (which are fixed along the sequence). The following theorem is the main result of the paper.

THEOREM 1. *If $\{\bar{\Psi}_{\pi_*^n}^n : n \geq 1\}$ is UI, then for any sequence of admissible controls π ,*

$$(2.6) \quad \liminf_{n \rightarrow \infty} \liminf_{t \rightarrow \infty} \bar{C}_{\pi^n}^n(t) \geq \lim_{n \rightarrow \infty} \lim_{t \rightarrow \infty} \bar{C}_{\pi_*^n}^n(t) = c_* \quad w.p.1.$$

PROOF. The result follows from Theorems 4 and 5 in Section 6. □

The next corollary is a simple consequence of Theorem 1.

COROLLARY 2.1. *If $\{\bar{\Psi}_{\pi_*^n}^n : n \geq 1\}$ is UI, then π_* satisfies (2.4), i.e., it is asymptotically optimal.*

PROOF. It follows from (2.6) by applying Fatou's Lemma twice that

$$\liminf_{n \rightarrow \infty} \liminf_{t \rightarrow \infty} \mathbb{E} [\bar{C}_{\pi^n}^n(t)] \geq \mathbb{E} \left[\lim_{n \rightarrow \infty} \lim_{t \rightarrow \infty} \bar{C}_{\pi_*^n}^n(t) \right] = c_*.$$

Moreover, for each $n \geq 1$, the embedded DTMC $\bar{Q}_{\pi_*^n}^n$ is ergodic by (Fricker and Jaibi, 1994, Proposition 1), so that $\bar{C}_{\pi_*^n}^n(t)$ converges to a deterministic finite value w.p.1 as $t \rightarrow \infty$ (see (7.28) in Section 7.2 for a characterization of this constant). Since $\bar{C}_{\pi_*^n}^n(t)$ is continuous in t , it holds that $\mathbb{E} \left[\sup_{t \geq 0} \bar{C}_{\pi_*^n}^n(t) \right] < \infty$. Thus,

$$\lim_{n \rightarrow \infty} \lim_{t \rightarrow \infty} \mathbb{E} [\bar{C}_{\pi_*^n}^n(t)] = \lim_{n \rightarrow \infty} \mathbb{E} \left[\lim_{t \rightarrow \infty} \bar{C}_{\pi_*^n}^n(t) \right] = c_*,$$

where the first equality follows from the dominated convergence theorem, and the second equality follows from (2.6) and the fact that $\lim_{t \rightarrow \infty} \bar{C}_{\pi_*}^n(t)$ is a constant w.p.1. \square

To apply Theorem 1 (and Corollary 2.1) we must (i) compute the fluid-optimal control parameters (L_*, \mathbf{r}_*) , which are also the parameters of the binomial-exhaustive policy π_*^n for all $n \geq 1$, and (ii) establish that the UI condition holds. We now discuss these two conditions, starting with the latter.

The UI Condition in Theorem 1. Theorem 2 below provides sufficient conditions for $\{\bar{\Psi}_{\pi_*^n}^n : n \geq 1\}$ to be UI, whenever the next assumption holds.

ASSUMPTION 2. *The following two conditions hold for all $i \in \mathcal{I}$.*

- (i) $\mathbb{E} [e^{tV_i^n}] < \infty$ for all $t \geq 0$ and $n \geq 1$.
- (ii) $\mathbb{E} [(\bar{V}_i^n)^\ell] \rightarrow s_i^\ell$ as $n \rightarrow \infty$ for all $\ell \geq 2$.

Recall that S_k denotes a generic random variable having the service time distribution of the class- k customers, $k \in \mathcal{K}$.

THEOREM 2. *For $p \geq 1$, let $\psi(x) = O(\|x\|^p)$. Under Assumption 2, $\{\bar{\Psi}_{\pi_*^n}^n : n \geq 1\}$ is UI if either of the following two conditions holds for all $k \in \mathcal{K}$.*

- (i) $p > 1$, and for some $\epsilon > 0$, $\mathbb{E} [e^{tS_k}] < \infty$ for all $t \in (-\epsilon, \epsilon)$.
- (ii) $p = 1$, and $\mathbb{E} [S_k^2] < \infty$.

PROOF. See Section 7.3. \square

As an immediate corollary to Theorems 1 and 2, we obtain that π_* is asymptotically optimal under either one of the assertions in Theorem 2.

We remark that the condition that the second moments of the service times are finite when $p = 1$ is also necessary in order for the desired UI to hold; see Theorem 3 in Hu et al. (2020). Thus, for cost functions that grow at most at a linear rate, π_* is asymptotically optimal if and only if $\mathbb{E} [S_k^2] < \infty$ for all $k \in \mathcal{K}$.

The Optimal Control Parameters. Solving the FCP in order to compute the optimal control parameters (L_*, \mathbf{r}_*) is not always feasible, because it requires optimizing over the table structure (within the infinite set of all possible augmented tables) simultaneously with optimizing the parameters. Nevertheless, in addition to solving the FCP on a case-by-case basis, it can also be solved in certain general settings. The most important case for which we can solve the FCP is that of cyclic basic tables, when the cost function ψ is separable convex, namely, is of the form $\psi(x) = \sum_{k=1}^K \psi_k(x_k)$, $x = (x_1, \dots, x_K)$, where $\psi_k : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is convex for each $k \in \mathcal{K}$. In this setting, we prove that the exhaustive policy, under which the server empties the queues in all visits, is fluid optimal; see Proposition 4.1. See also Corollary 6.1 for the corresponding asymptotic-optimality result.

We also consider a restricted optimal-control problem for cases in which the FCP cannot be solved. In the restricted problem, we optimize the control parameters for a finite set of values of L (including the case $L = 1$). Unlike the FCP, the *restricted FCP* (RFCP) can always be solved, and the corresponding binomial-exhaustive policy is then asymptotically optimal under the same conditions in Theorem 2, although among a smaller family of admissible controls; see Definition 5.2 in Section 5.1.1 and Theorem 6 in Section 6.2.

2.3. *Roadmap to the Proof of Theorem 1.* We now describe the main steps in the proof of Theorem 1. We emphasize that the description here is provided for overview, and is not meant to be fully rigorous. In addition, the proof scheme outlined below is for the general case, in which we optimize over all possible augmented tables. The aforementioned restricted problem follows a similar procedure, except that the family of admissible controls is smaller for that latter problem.

(I) **Formalizing an FCP (Sections 3 and 4.1).** To formalize an FCP corresponding to the control problem for the stochastic system, we consider a fluid model of the original stochastic system. Specifically, we consider a deterministic polling system that has the same basic table as the stochastic system, in which the arrival and service processes are replaced by deterministic continuous processes with the same rates λ_k and μ_k , $k \in \mathcal{K}$. In that deterministic counterpart, the queue process Q is replaced by a *fluid model* $q := \{q(t) : t \geq 0\}$, whose dynamics are determined by its initial condition and the switching policy.

We then seek a fluid control that minimizes the long-run average cost. Let $q_\phi(t)$ denote the value of the fluid queue at time t under control ϕ , for ϕ in some appropriate family of admissible fluid controls. To apply the asymptotic-optimization framework, we want q_ϕ to be related to Q_π via a FWLLN. Next, as will be shown in the proof of Theorem 4, all possible fluid limits for Q_π in stationarity are “almost periodic” in the sense that each such limit is arbitrarily close to a *periodic equilibrium* (PE).² (A fluid model q_e is a PE if $q_e(t + \tau) = q_e(t)$ for all $t \geq 0$, for some $\tau > 0$; see Definition 3.1.) Thus, optimizing the long-run average cost is equivalent to first optimizing over all possible PE, and then finding a control that guarantees convergence of the fluid model to the optimal PE. In particular, we can take the set of admissible fluid controls, denoted by Φ , to be the set of all controls under which the fluid model converges to a PE, so that the FCP reduces to

$$\inf_{\phi \in \Phi} \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \psi(q_\phi(u)) du = \frac{1}{\tau_*} \int_0^{\tau_*} \psi(q_*(u)) du,$$

where q_* is the optimal PE and τ_* is its period.

(II) **Solving the FCP (Section 4.2).** From the description of step (I), solving the FCP consists of two components: first, we need to identify an optimal PE q_* , and second, we need to design a control $\phi_* \in \Phi$ such that q_{ϕ_*} converges in an appropriate sense (see Definition 3.3) to q_* . These two components are interconnected, because we have quite some flexibility in how we characterize the optimal PE. In particular, the orbit of any PE q_e is a loop (a closed curve) in \mathbb{R}_+^K , which is fully characterized by specifying the server’s departure epochs during that server cycle, together with a single point on the PE, because the dynamics of the fluid model are deterministic between switching epochs. A fluid control ϕ is then a switching rule that produces the desired trajectory q_e whenever the initial point is on that PE’s trajectory, and is in Φ if it guarantees the desired convergence. (The main difficulty in establishing that a control ϕ_* is optimal is in establishing that it is an element of Φ .)

The fluid control ϕ_* we propose prescribes reducing queue $p(i)$ by a fixed proportion r_i of its size at the polling epoch. Specifically, letting the value of the fluid queue polled at stage i be $q_{p(i)}(a_i^{(m)})$ at the polling epoch, the server will switch away from that queue when its value reaches $(1 - r_i)q_{p(i)}(a_i^{(m)})$, $i \in \mathcal{I}^L$, $m \geq 1$. We refer to this control as stage-based proportion reduction (SB-PR), and to the SB-PR control with the optimal parameters (L_*, \mathbf{r}_*) as the *optimal SB-PR*.

²We use the acronym PE for both singular and plural forms, i.e., periodic equilibrium and periodic equilibria.

(III) **Proving asymptotic optimality (Sections 5 and 6).** The translation step of the optimal SB-PR to the binomial-exhaustive policy is straightforward, and was discussed above and further detailed in Section 5. To prove that the binomial-exhaustive policy with parameters (L_*, \mathbf{r}_*) is asymptotically optimal, we first show (in Theorem 4) that the limiting holding cost of any sequence of admissible controls is lower bounded by the optimal fluid cost. We then prove that under the conditions in Theorem 2, the binomial-exhaustive policy with the optimal SB-PR parameters achieves the lower bound asymptotically; see Theorem 5.

3. The Fluid Model. To formulate the FCP, we start by constructing a fluid model for the polling system. To this end, we consider a deterministic polling system having the exact same system's topology and basic table as the stochastic system, but in which arrivals and service completions occur continuously and deterministically at rates λ_k and μ_k , $k \in \mathcal{K}$, respectively. Let $q(t)$ denote the fluid content at time t , and for $i \in \mathcal{I}^L$ and $m \geq 1$, let $a_i^{(m)}$, $d_i^{(m)}$ and $b_i^{(m)}$ denote the polling epoch, departure epoch, and the busy time of stage i during the m th server cycle. Let $u^{(m-1)}$ be the time at which the m th server cycle begins, and $\tau^{(m)}$ be the length of the m th server cycle. For the following, we write q instead of q_ϕ to simplify the notations whenever the control is fixed, and refer to q as the “queue” or “fluid content” interchangeably.

Let $\mathbf{k} = (i \in \mathcal{I}^L : p(i) = k)$ denote the vector of ordered stages at which queue k is visited in a server cycle, so that queue k is visited a total of $\dim(\mathbf{k})$ times over a server cycle. Then the fluid queue over the first server cycle satisfies

$$q_k(t) = q_k(0) + \lambda_k t - \mu_k \sum_{j=1}^{\dim(\mathbf{k})} \int_0^t \mathbf{1}_{[a_{k_j}^{(1)}, d_{k_j}^{(1)}]}(s) ds, \quad k \in \mathcal{K}, t \in [u^{(0)}, u^{(1)}].$$

Since the fluid model is time-invariant, it can be described inductively via its dynamics over one server cycle; in particular, the dynamics of q_k over the time interval $[u^{(m-1)}, u^{(m)}]$, namely, during the m th server cycle, can be described by

$$(3.1) \quad q_k(t) = q_k(u^{(m-1)}) + \lambda_k(t - u^{(m-1)}) - \mu_k \sum_{j=1}^{\dim(\mathbf{k})} \int_{u^{(m-1)}}^t \mathbf{1}_{[a_{k_j}^{(m)}, d_{k_j}^{(m)}]}(s) ds,$$

for $t \in [u^{(m-1)}, u^{(m)}]$, $k \in \mathcal{K}$, $m \geq 1$.

3.1. *The Fluid Model as a Hybrid Dynamical System.* Note that the values of $b_i^{(m)}$, $a_i^{(m)}$ and $d_i^{(m)}$ are determined by the state of q and the control, and are therefore not available a-priori (those values must be computed on the fly). It is therefore more useful to represent q as a solution to a differential equation. To achieve such a representation, let $z(t)$ denote the location of the server at time t : we write $z(t) = i$ if the server is actively serving queue $p(i)$ at time t , and $z(t) = \ominus_i$ if the server is switching from stage i to stage $i + 1$ at time t (with $i + 1 := 1$ for $i = IL$). We let

$$(3.2) \quad \mathcal{Z} := \{1, \dots, IL, \ominus_1, \dots, \ominus_{IL}\}$$

denote the state space of the server-location process z .

If a control depends only on the state of the queue process q and the location of the server, then we should keep track of the state of the process (q, z) in order to determine the values of the switching times. However, since q is a “surrogate” for the stochastic process Q , and since we consider controls under which $\{\tilde{Q}(m) : m \geq 0\}$ in (2.2) is a DTMC, we also allow

the control to depend on the value of the fluid queue at the last polling epoch prior to t , i.e., on $q(a(t))$, where

$$a(t) := \max \{a_i^{(m)} \leq t : i \in \mathcal{I}^L, m \geq 1\}.$$

Thus, we consider the process

$$(3.3) \quad x(t) := (q(t), q(a(t)), z(t)), \quad t \geq 0,$$

taking values in $\mathbb{R}_+^K \times \mathbb{R}_+^K \times \mathcal{Z}$. Note that x in (3.3) is a hybrid of the fluid-content process q , which has a continuous state space, and the server-location process z , which has a finite state space, and is therefore an HDS. (In fact, x is a slight generalization of standard HDS due to the additional processes $q(a(t))$.) Then x is a solution to the following state equations.

$$(3.4) \quad \begin{aligned} \dot{q}(t) &= f(z(t)) \\ z(t) &= g(q(t), q(a(t-)), z(t-)), \\ a(t) &= h(q(t), q(a(t-)), z(t-)), \end{aligned}$$

where $f : \mathcal{Z} \rightarrow \mathbb{R}^K$, $g : \mathbb{R}_+^K \times \mathbb{R}_+^K \times \mathcal{Z} \rightarrow \mathcal{Z}$, and $h : \mathbb{R}_+^K \times \mathbb{R}_+^K \times \mathcal{Z} \rightarrow \mathbb{R}_+^K$ are the functions specified below.

First, the function f determines the dynamics of the queues, which change at the polling and departure epochs of each stage. Thus, for each $k \in \mathcal{K}$, f is defined via

$$f_k(q(t), z(t)) = \begin{cases} \lambda_k - \mu_k & \text{if } z(t) = i \text{ and } p(i) = k \\ \lambda_k & \text{otherwise.} \end{cases}$$

The functions g and h are determined by the control; to characterize these function, we define a *service function* $\phi_i : \mathbb{R}_+^K \rightarrow \mathbb{R}_+$ mapping the queue length at the polling epoch of stage i to the immediate busy time of the server;

$$(3.5) \quad \phi_i(q(a_i^{(m)})) := b_i^{(m)}, \quad i \in \mathcal{I}^L, m \geq 1.$$

The non-idling property we impose implies that

$$b_i^{(m)} \leq q(a_i^{(m)}) / (\mu_{p(i)} - \lambda_{p(i)}), \quad i \in \mathcal{I}^L, m \geq 1.$$

Indeed, the expression on the right-hand side of the above inequality is the time at which the fluid queue that is attended by the server hits state 0 if the server keeps processing work continuously.

Now, the function g characterizing the location of the server as follows:

(i) If $z(t-) = i$ and $q_{p(i)}(t) = q_{p(i)}(a(t-)) - (\mu_{p(i)} - \lambda_{p(i)})\phi_i(q(a(t-)))$, define

$$\begin{aligned} j_s &:= \min\{j \geq i : s_{(j \bmod IL)} > 0\} \\ j_\phi &:= \min\{j > i : \phi_{(j \bmod IL)}(q(a(t-))) > 0\}. \end{aligned}$$

(a) If $j_s < j_\phi$, then $g(q(t), q(a(t-)), z(t-)) = \ominus_{(j_s \bmod IL)}$.

(b) Otherwise, $g(q(t), q(a(t-)), z(t-)) = (j_\phi \bmod IL)$.

(ii) If $z(t-) = \ominus_i$ and $q_{p(i+1)}(t) = q_{p(i+1)}(a(t-)) + \lambda_{p(i+1)}(\phi_i(q(a(t-))) + s_i)$, define

$$\begin{aligned} j_s &:= \min\{j > i : s_{(j \bmod IL)} > 0\} \\ j_\phi &:= \min\{j > i : \phi_{(j \bmod IL)}(q(a(t-))) > 0\}. \end{aligned}$$

(a) If $j_s < j_\phi$, then $g(q(t), q(a(t-)), z(t-)) = \ominus_{(j_s \bmod IL)}$.

(b) Otherwise, $g(q(t), q(a(t-)), z(t-)) = (j_\phi \bmod IL)$.

(iii) Otherwise, $g(q(t), q(a(t-)), z(t-)) = z(t-)$.

Lastly, the function h updates the most recent polling epoch according to

$$h(q(t), q(a(t-)), z(t-)) = \begin{cases} t & \text{if } z(t-) = \ominus_i \text{ and } q_{p(i+1)}(t) = q_{p(i+1)}(a(t-)) + \lambda_{p(i+1)}(\phi_i(q(a(t-))) + s_i) \\ a(t-) & \text{otherwise.} \end{cases}$$

3.2. *Qualitative Behavior of the HDS.* Our qualitative analysis of the HDS relies on fundamental concepts defined in this section.

DEFINITION 3.1 (PE). *A solution x_e to the HDS (3.4) is a PE if there exists $\tau > 0$ such that $x_e(t + \tau) = x_e(t)$ for all $t \geq 0$. The smallest such τ is called the period.*

Note that a solution x_e is a PE if and only if the orbit of q_e , namely, the image of q_e in \mathbb{R}_+^K , is a loop. Thus, we will henceforth refer to the queue component q_e as a PE.

DEFINITION 3.2 (L -cycle PE). *A solution x_e to the HDS (3.4) is an L -cycle PE if $x_e(t + \tau_L) = x_e(t)$ for all $t \geq 0$, where τ_L is its cycle length spanning L table cycles.*

Clearly, the cycle length τ_L of a PE is an integer product of the period of that PE. It follows from basic flow-balance equations that the cycle length of L -cycle PE satisfies

$$(3.6) \quad \tau_L = sL/(1 - \rho).$$

To see this, observe that the server must be working a fraction ρ of the time, and is therefore switching between stages for a fraction $1 - \rho$ of the time. Since the total switchover time over L table cycles is sL , it holds that $\tau_L(1 - \rho) = sL$, from which (3.6) follows.

Stable PE. The purpose of the fluid-optimal control is to steer every possible trajectory q to a desired PE q_* . It is significant that convergence of trajectories to a PE cannot occur in the Lyapunov sense, i.e., it does not hold that $\|q(t) - q_e(t)\| \rightarrow 0$ as $t \rightarrow \infty$ for a trajectory q that converges to the PE q_e . Instead, convergence of q to the PE q_e is said to hold if the orbit of q in \mathbb{R}_+^K “spirals” towards the closed orbit of q_e . Recall that, without loss of generality, $u^{(0)} = 0$, namely, the beginning of the first server cycle is time 0. Similarly, we take $u_e^{(0)} = 0$ for a PE x_e .

DEFINITION 3.3 (convergence to a PE). *A solution x to the HDS (3.4) is said to converge to a PE x_e if $\|q(u^{(m)} + \cdot) - q_e(\cdot)\|_t \rightarrow 0$ as $m \rightarrow \infty$, for all $t > 0$.*

A PE q_e may be of several types; if any other trajectory in some neighborhood of q_e converges to it, then q_e is called a *stable limit cycle*. (It is unstable if the trajectories in its neighborhood are “spiraling” away from it, and semi-stable if some trajectories in its neighborhood converge, while other are repelled.) For our optimality result, we require a stronger stability property to hold.

DEFINITION 3.4 (global limit cycle). *A PE q_e of the HDS is said to be a global limit cycle if all the trajectories of the HDS converge to q_e .*

In ending we remark that determining the number of limit cycles of a dynamical system is in general a hard problem, even in the classical setting of dynamical systems with continuous vector fields. (For planar systems with a polynomial vector field of degree greater than 1, this

is part of Hilbert's 16th open problem, which is still unsolved.) Further, HDS of the form (3.1) can exhibit chaotic behavior, and in particular, possess infinitely many PE, *none of which is a limit cycle*, even when the continuous-state process q is of a dimension as low as 3; see Chase et al. (1993). In contrast, the fluid model (and limit) under our proposed SB-PR control will be shown to possess a global limit cycle (which is necessarily unique).

3.3. Fluid Limits and Their Relation to the Fluid Model. Whereas the fluid model is derived for deterministic polling systems, the *fluid limits*, namely, the subsequential limits of the sequence of fluid-scaled queue processes, may not be deterministic under an arbitrary sequence of controls. A FWLLN holds, and the resulting fluid limit is deterministic, under an extra regularity condition; see Proposition 3.1 below. Since the deterministic fluid model is the basis for solving the FCP and deriving the asymptotically optimal control, it is significant that the FWLLN holds for the binomial-exhaustive policy.

Consider the stochastic polling system, and let $Z(t)$ denote the location of the server at time t , defined on the same state space \mathcal{Z} in (3.2); that is, Z is the stochastic counterpart of the server-location process z in the fluid model. For $t \geq 0$, let

$$(3.7) \quad A(t) := \max\{A_i^{(m)} \leq t : i \in \mathcal{I}^L, m \geq 1\}.$$

We define the state-process (of the stochastic system)

$$X(t) := (Q(t), Q(A(t)), Z(t)), \quad t \geq 0,$$

where we removed π from the notation to simply it.

Let $\mathcal{P}_k := \{\mathcal{P}_k(t) : t \geq 0\}$ denote the Poisson arrival process to buffer k , and let $\mathcal{S}_k := \{\mathcal{S}_k(t) : t \geq 0\}$ denote the *potential* service process in buffer k , namely, $\mathcal{S}_k(t)$ would be the number of class- k service completions by time t if the server were to process work from queue k continuously during $[0, t)$. In particular,

$$\mathcal{S}_k(t) := \sup \left\{ m \geq 1 : \sum_{j=1}^m S_k^{(j)} \leq t \right\},$$

where $\{S_k^{(j)} : j \geq 1\}$ is a sequence of i.i.d. random variables distributed like S_k . Then for $k \in \mathcal{K}$,

$$(3.8) \quad Q_k(t) = Q_k(0) + \mathcal{P}_k(t) - \mathcal{S}_k \left(\sum_{m=1}^{\infty} \sum_{\ell=1}^{\dim(\mathbf{k})} \int_0^t \mathbf{1}_{[A_{k_\ell}^{(m)}, D_{k_\ell}^{(m)}]}(u) du \right), \quad t \geq 0.$$

Now, consider the sequence of stochastic systems under the large-switchover-time scaling. For the n th system, let $A_i^{(m),n}$ and $D_i^{(m),n}$ denote, respectively, the polling and departure epoch of stage i in the m th server cycle, $i \in \mathcal{I}^L$, $m \geq 1$. The corresponding fluid-scaled server-switching epochs (arrival and departure epochs to and from the queues) are given by $\bar{A}_i^{(m),n} := A_i^{(m),n}/n$ and $\bar{D}_i^{(m),n} := D_i^{(m),n}/n$. Analogously to (3.7), we denote the most recent polling epoch prior to time t in system n via

$$A^n(t) := \max\{A_i^{(m),n} \leq nt : i \in \mathcal{I}^L, m \geq 1\}, \quad t \geq 0.$$

The fluid-scaled state-process is given by

$$\bar{X}^n(t) := (\bar{Q}^n(t), \bar{Q}^n(\bar{A}^n(t)), Z(nt)), \quad t \geq 0,$$

where $\bar{Q}^n(t) := Q(nt)/n$ and $\bar{A}^n(t) := A^n(t)/n$ (there is no spacial scaling of the process $Z(nt)$).

For $k \in \mathcal{K}$ and $n \geq 1$, define $\mathcal{S}_k^n(t) := \mathcal{S}(nt)$, $\mathcal{P}_k^n(t) := \mathcal{P}_k(nt)$, $\bar{\mathcal{S}}_k^n(t) := \mathcal{S}_k^n(nt)/n$ and $\bar{\mathcal{P}}_k^n(t) := \mathcal{P}_k^n(nt)/n$. Then the representation (3.8) for the queue in the n th system becomes

$$(3.9) \quad Q_k^n(t) = Q_k^n(0) + \mathcal{P}_k^n(t) - \mathcal{S}_k^n \left(\sum_{m=1}^{\infty} \sum_{\ell=1}^{\dim(\mathbf{k})} \int_0^t \mathbf{1}_{[\bar{A}_{k_\ell}^{(m),n}, \bar{D}_{k_\ell}^{(m),n})}(u) du \right), \quad t \geq 0.$$

LEMMA 3.1 (tightness). *If $\{\bar{Q}^n(0) : n \geq 1\}$ is tight in \mathbb{R}_+^K , then $\{\bar{Q}^n : n \geq 1\}$ is C -tight in D^K , and the sample paths of its subsequential limits are of the form (3.1).*

It is significant that, for $i \in \mathcal{I}$ and $m \geq 1$, the time epochs $u^{(m)}$, $a_i^{(m)}$ and $d_i^{(m)}$ of a subsequential limit of \bar{Q}^n may be random variables, in which case that limit q is stochastic. However, Lemma 3.1 states that, even in this case, the evolution of a stochastic limit q between any two consecutive server-switching epochs is deterministic, and is characterized in (3.1).

PROOF OF LEMMA 3.1. Fix $T > 0$. Due to the scaling of the switchover times in Assumption 1, the number of server switchings in system n over the time interval $[0, nT]$ is finite w.p.1 as $n \rightarrow \infty$. Hence, the sequence of fluid-scaled server-switching epochs is tight in $[0, T)$. In particular, any subsequence of the sequences $\{\bar{A}_{k_\ell}^{(m),n} : n \geq 1\}$ and $\{\bar{D}_{k_\ell}^{(m),n} : n \geq 1\}$ in (3.9) has a further converging sub-subsequence (for all m and k_ℓ for which there are infinitely many elements of these sequences in $[0, T)$). Now, the indicator functions in the time-changed service process in (3.9) are fixed at the value 0 or at 1 between any two consecutive server-switching epochs, so that \bar{Q}_k^n is a continuous mapping of its primitives between any two such switching epochs. It follows from (Whitt, 2002, Theorem 13.6.4) that any subsequence of $\{\bar{Q}_k^n : n \geq 1\}$ for which all the server-switching epochs in $[0, T)$ converge, converges in D^K to q_k in (3.1) as $n \rightarrow \infty$. \square

It follows immediately from the proof of Lemma 3.1 that if the sequences of fluid-scaled server-switching epochs converge in $[0, T)$ for all $T > 0$, then $\bar{Q}_k^n \Rightarrow q_k$ in D^K as $n \rightarrow \infty$, for q_k in (3.1). In fact, since the dynamics of the queues are deterministic between any two server-switching epochs, convergence of the server departure times implies that the server arrival times also converge. We therefore have the following FWLLN.

PROPOSITION 3.1 (FWLLN). *Assume that $\bar{Q}^n(0) \Rightarrow q(0)$ in \mathbb{R}_+^K as $n \rightarrow \infty$. If $\bar{D}_i^{(m),n} \Rightarrow d_i^{(m)}$ in \mathbb{R}_+^K for all $m \geq 1$ and $i \in \mathcal{I}^L$, then $\bar{Q}^n \Rightarrow q$ in D^K as $n \rightarrow \infty$, where each element q_k , $k \in \mathcal{K}$, of the vector process q satisfies (3.1).*

Note that if $q(0)$ and $d_i^{(m)}$ are deterministic for all $m \geq 1$ and $i \in \mathcal{I}^L$, then the fluid limit q is the unique solution to an HDS of the form (3.1).

4. The Fluid Control Problems. In this section we formally define the FCP, whose solution is an optimal fluid control for the family of all augmented tables, and the restricted problem, namely the RFCP, whose solution is an optimal fluid control for a finite set of augmented tables.

The FCP. For the FCP, we consider the set Φ of controls for which the following holds for each control $\phi \in \Phi$:

(i) There exists a unique solution $q_\phi^\gamma := \{q_\phi^\gamma(t) : t \geq 0\}$ to the HDS (3.4) under ϕ for any initial condition $\gamma \in \mathbb{R}_+^K$.

(ii) Any solution q_ϕ^γ converges to a limit cycle as $t \rightarrow \infty$.

For $\gamma \in \mathbb{R}_+^K$, let

$$C_\phi(\gamma) := \inf_{\phi \in \Phi} \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \psi(q_\phi^\gamma(u)) du.$$

DEFINITION 4.1 (fluid optimal control). *We say that ϕ_* is fluid-optimal if $C_{\phi_*}(\gamma) \leq C_\phi(\gamma)$ for all $\phi \in \Phi$ uniformly in γ .*

The following lemma, whose proof appears in Section 4.3, motivates searching for an “optimal PE”, namely, a PE that achieves the lowest possible time-average cost over its cycle length among all possible PE, and then devising a control ensuring that that PE is a global limit cycle.

LEMMA 4.1. *For $\phi \in \Phi$ and $\gamma \in \mathbb{R}_+^K$, let q_ϕ^γ denote the unique solution to the HDS when control ϕ is exercised and when $q_\phi^\gamma(0) = \gamma$. Let q_e^γ denote the limit cycle to which q_ϕ^γ converges, and τ_e^γ denote its cycle length. Then*

$$\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \psi(q_\phi^\gamma(u)) du = \frac{1}{\tau_e^\gamma} \int_0^{\tau_e^\gamma} \psi(q_e^\gamma(u)) du.$$

Due to Lemma 4.1, the FCP is concerned with finding a control ϕ_* that achieves the optimal long-run average c_* , where

$$(4.1) \quad c_* := \inf_{\phi \in \Phi} C_\phi(\gamma) := \inf_{\phi \in \Phi} \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \psi(q_\phi^\gamma(u)) du, \quad \text{for all } \gamma \in \mathbb{R}_+^K.$$

In turn, to solve the FCP, we seek a control $\phi_* \in \Phi$ under which there exists a *global limit cycle* q_* , such that

$$(4.2) \quad \frac{1}{\tau_*} \int_0^{\tau_*} \psi(q_*(u)) du \leq \frac{1}{\tau_e} \int_0^{\tau_e} \psi(q_e(u)) du$$

holds for any other PE q_e (whose cycle length is τ_e). Note that both τ_* and τ_e in (4.2) are allowed to have any possible value of τ_L in (3.6), so that we are effectively optimizing the PE over all possible augmented tables.

Solving the FCP. We start by identifying closed curves in \mathbb{R}_+^K which are possible solution to the HDS (namely, they can be obtained as a PE under some control). We refer to each such closed curve q_e as a *PE-candidate*, and treat it as a mapping from $[0, \tau_e]$ to \mathbb{R}_+^K (where $q_e(0) = q_e(\tau_e)$). We then optimize over all possible PE-candidates in order to find an *optimal PE-candidate* q_* for which (4.2) holds. Finally, we design an optimal control $\phi_* \in \Phi$ under which the optimal PE-candidate q_* is a global limit cycle for the HDS, so that (4.1) holds for any solution q_ϕ^γ to (3.4) with initial condition $\gamma \in \mathbb{R}_+^K$.

We emphasize two points: (i) We do not rule out the possibility that, in general, the infimum c_* is not achievable via a PE-candidate, namely, that there exists no PE-candidate whose time-average cost over the cycle length is c_* . (However, we are unaware of such pathological examples; we do not study this problem due to its impracticability, as explained in the next point.) (ii) Computing a PE-candidate for which c_* is attained is not always practically feasible, due to the need to optimize the table structure among all the possible augmented tables. (Hence, proving that a given problem is well-posed may also be impractical.)

As was mentioned in Section 2.2, solving the FCP is possible for specific systems or in specific settings. The most important case for which the FCP can be solved is when the cost function is separable convex (including linear), and the basic table is cyclic; see Proposition 4.1 and Corollary 6.1 for the corresponding asymptotic-optimality result.

REMARK 4.1 (On the set Φ). It is significant that the set of fluid limits is larger than the set of possible fluid models under Φ . In particular, fluid limits under a sequence of admissible controls can be non-stable, in the sense that they do not converge to a limit cycle, and can also be stochastic. Thus, Φ is smaller than the set of possible controls for the fluid limits. However, Theorem 4 in Section 6 proves that c_* in (4.1) is a lower bound on the achievable costs asymptotically (as $n \rightarrow \infty$), so that, it is sufficient to search for control in Φ .

The RFCP. When solving the FCP in (4.1) is not feasible, one can instead optimize among all L -cycle PE for L in some finite subset $\mathcal{N} \subset \mathbb{N}$, e.g., $L \in \mathcal{N} = \{1, \dots, M\}$, where $M \geq 1$ is a finite integer. To this end, we consider the RFCP, whose goal is to find $c_{\mathcal{N}}$, where

$$(4.3) \quad c_{\mathcal{N}} := \min_{L \in \mathcal{N}} \inf_{\phi \in \Phi^L} C_{\phi}(\gamma) := \min_{L \in \mathcal{N}} \inf_{\phi \in \Phi^L} \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \psi(q_{\phi}^{\gamma}(u)) du, \quad \text{for all } \gamma \in \mathbb{R}_+^K,$$

where $\Phi^L \subset \Phi$ is the set of all the controls under which any solution to the HDS (3.4) converges to an L -cycle limit cycle. Correspondingly, for each $L \in \mathcal{N}$ we seek an optimal L -cycle PE-candidate q_*^L such that the inequality

$$\frac{1}{\tau_L} \int_0^{\tau_L} \psi(q_*^L(u)) du \leq \frac{1}{\tau_L} \int_0^{\tau_L} \psi(q_e^L(u)) du,$$

holds for any other L -cycle PE q_e^L . The solution to the RFCP is then

$$q_{\mathcal{N}} := \min_{L \in \mathcal{N}} q_*^L.$$

The procedure for solving the RFCP is similar to that of solving the FCP: We start by computing an optimal PE-candidate for each $L \in \mathcal{N}$, and take the one with the lowest time-average cost over the cycle length to be the optimal PE-candidate for the RFCP. Letting $L_{\mathcal{N}}$ denote the number of table cycles contained in the cycle length of $q_{\mathcal{N}}$, we then design a control $\phi_{\mathcal{N}}$ under which $q_{\mathcal{N}}$ is a global limit cycle for the HDS. Unlike the FCP (4.1), solving the RFCP is always feasible, because computing an optimal PE-candidate q_*^L for any fixed L , and therefore computing $q_{\mathcal{N}}$, is straightforward.

In ending we remark that $L = 1$ should always be an element of \mathcal{N} , not only because it corresponds to the basic table, but also because the period of an L -cycle PE can be smaller than τ_L , i.e., the period might be $\tau_{L_2} < \tau_L$, with L being divisible by L_2 . In particular, an optimal L -cycle PE with $L > 1$ may have period τ_1 .

4.1. *Computing an Optimal PE-Candidate.* We now discuss the first step in solving the FCP and RFCP, namely, characterizing an optimal PE-candidate.

4.1.1. *Optimal PE-Candidates for the FCP.* Let \mathcal{Q} denote the set of all PE-candidates (of all possible cycle lengths τ_L , $L \geq 1$). When a solution to the FCP (4.1) exists, an optimal PE-candidate for this FCP solves the optimization problem

$$(4.4) \quad \min_{q_e \in \mathcal{Q}} \frac{1}{\tau_L} \int_0^{\tau_L} \psi(q_e(u)) du.$$

Let q_{exh} denote the one-cycle PE under the exhaustive policy in which the server empties the queue it attends and then switches to the next queue in the table (the existence of such a PE is established in Lemma 4.3 below). Recall that ψ is separable convex if $\psi(x) = \sum_{k \in \mathcal{K}} \psi_k(x_k)$ for $x \in \mathbb{R}_+^K$, and ψ_k is convex for each $k \in \mathcal{K}$.

PROPOSITION 4.1. *If the basic table is cyclic and ψ is separable convex, then q_{exh} is a solution to (4.4).*

PROOF. See Appendix A. □

Whereas q_{exh} is not a solution to (4.4) in general, as we show below, it is easy to see that each queue must be emptied at least once in a PE-candidate that solves (4.4). In particular, for $q_{*,k}$ denoting the k th component process of a solution q_* to (4.4), and τ_* denoting the period (or cycle length) of q_* , it must hold that $q_{*,k}(t_k) = 0$, for some $t_k \in [0, \tau_*)$, $k \in \mathcal{K}$. To see this, observe that a PE-candidate is completely determined by its initial condition and the busy times $(b_i, i \in \mathcal{I}^L)$. Consider a PE-candidate $q_e^{(1)}$ in which the k th queue, denoted by $q_{e,k}^{(1)}$, is such that $q_{e,k}^{(1)}(t) > 0$ for all $t \in [0, \tau_e)$, where τ_e is the period of $q_{e,k}^{(1)}$. We can construct a PE-candidate $q_e^{(2)}$ that has lower cost than $q_e^{(1)}$ by taking

$$q_{e,k}^{(2)}(0) := q_{e,k}^{(1)}(0) - \min_{t \in [0, \tau_*)} q_{e,k}^{(1)}(t), \quad q_{e,\ell}^{(2)}(0) := q_{e,\ell}^{(1)}(0) \text{ for } \ell \neq k, \ell \in \mathcal{K},$$

and giving $q_e^{(2)}$ the same busy times $(b_i, i \in \mathcal{I}^L)$ of $q_e^{(1)}$. Then $q_e^{(2)}(t) < q_e^{(1)}(t)$ for all $t \in [0, \tau_e)$, and the same inequality holds for the corresponding costs, because ψ is non-decreasing.

4.1.2. Optimal PE-Candidates for the RFCP. Let \mathcal{Q}^L denote the set of all PE-candidates having cycle length τ_L . To solve the RFCP in (4.3), we solve for the optimal L -cycle PE-candidate for each $L \in \mathcal{N}$, taking the one that gives the overall minimal cost as the solution. To this end, we consider the following optimization problem.

$$(4.5) \quad \min_{q_e^L \in \mathcal{Q}^L} \frac{1}{\tau_L} \int_0^{\tau_L} \psi(q_e^L(u)) du, \quad \text{for some (fixed) } L \geq 1.$$

Unlike (4.4), problem (4.5) always admits solution.

LEMMA 4.2. *For any fixed $L \in \mathbb{N}$, the optimization problem (4.5) admits a solution q_e^L .*

The proof of the lemma builds on Lemmas 4.3 and 4.4 which are stated below, and is therefore relegated to Section 4.3.

An analogous result to Proposition 4.1 holds for general cost functions when $L = 1$, due to the aforementioned fact that, in an optimal PE-candidate, each queue must be exhausted at least once. We therefore have:

PROPOSITION 4.2. *If the basic table is cyclic, then q_{exh} is a solution to (4.5) with $L = 1$.*

Finally, to demonstrate that q_{exh} is not an optimal PE-candidate in general, consider a system with three queues and basic (non-cyclic) table $(1, 2, 3, 2, 3)$. We take $\lambda_k = 2$, $\mu_k = 8$, $s_k = 2$ for $k = 1, 2, 3$, and ψ to be linear with $c_1 = c_2 = 1$, and consider $\mathcal{N} = \{1\}$. (We remark that the solution remains unchanged when we optimize over larger values of L ; we conjecture that the optimal one-cycle PE-candidate also solves (4.4).) If $c_3 > 4$, then it is optimal to *not exhaust* q_2 at stage 2. Moreover, the proportion of fluid processed at stage 2 is decreasing to 0 as c_3 increases. It is easy to explain why q_2 is not exhausted in one of its visits. Specifically, as the holding cost of q_3 increases, it becomes more and more advantageous to keep this queue smaller at the expense of making q_2 larger. This can be achieved while keeping q_2 (and its corresponding holding cost) bounded, because q_2 is visited twice, so the server has an opportunity to exhaust it in a server cycle.

4.2. *The SB-PR Control.* Consider an L -cycle PE-candidate q_e^L , and let r_i denote the proportion by which the queue polled in stage $i \in \mathcal{I}^L$ is reduced. In particular, with a_i and d_i denoting, respectively, the polling epoch and departure epoch of stage i ,

$$(4.6) \quad r_i := \begin{cases} \frac{q_{e,p(i)}^L(a_i) - q_{e,p(i)}^L(d_i)}{q_{e,p(i)}^L(a_i)} & \text{if } q_{e,p(i)}^L(a_i) > 0, \\ 0 & \text{otherwise} \end{cases}, \quad i \in \mathcal{I}^L.$$

Clearly, one can always represent a PE-candidate via parameters (L, \mathbf{r}) , where $\mathbf{r} := (r_i, i \in \mathcal{I}^L)$ is a vector whose component r_i is defined in (4.6). The following lemma shows that the reverse is also true; its proof is deferred to Section 4.3. For a given system, recall \mathcal{R} in (2.3) and that \mathcal{Q} is the set of all PE-candidates.

LEMMA 4.3. *For any $(L, \mathbf{r}) \in \mathbb{N} \times \mathcal{R}$, there exists a unique L -cycle PE-candidate q_e^L such that (4.6) is satisfied. In particular, the function $q_e^L \mapsto (L, \mathbf{r})$ is a bijection between \mathcal{Q} and $\mathbb{N} \times \mathcal{R}$.*

Lemma 4.3 motivates our proposed SB-PR control, which will be shown to be fluid optimal in Theorem 3 below.

DEFINITION 4.2 (SB-PR control). *Let $(L, \mathbf{r}) \in \mathbb{N} \times \mathcal{R}$. The SB-PR control with parameters (L, \mathbf{r}) has the service function*

$$\phi_i(q) = r_i q_{p(i)} / (\mu_{p(i)} - \lambda_{p(i)}), \quad i \in \mathcal{I}^L,$$

for ϕ_i in (3.5). In particular, at each stage i , the server reduces the polled queue to a proportion $1 - r_i$ of its value at the polling epoch of this stage.

Let (L_*, \mathbf{r}_*) denote the SB-PR control parameters corresponding to a solution to (4.4) (and optimal PE-candidate for the FCP), and for $\mathcal{N} \subset \mathbb{N}$, let $(L_{\mathcal{N}}, \mathbf{r}_{\mathcal{N}})$ denote the SB-PR control parameters corresponding to a solution to (4.5) (an optimal PE-candidate for the RFCP).

THEOREM 3 (optimality of SB-PR). *SB-PR with parameters (L_*, \mathbf{r}_*) is a solution to the FCP (4.1). Similarly, SB-PR with parameters $(L_{\mathcal{N}}, \mathbf{r}_{\mathcal{N}})$ is a solution to the RFCP (4.3).*

The proof of Theorem 3 relies on the following lemma, which establishes, in particular, that any PE-candidate is a bona-fide PE under the corresponding SB-PR control, and that this PE is a global limit cycle.

LEMMA 4.4 (global stability of SB-PR). *Let q_e^L be an L -cycle PE-candidate, and let \mathbf{r} be the corresponding vector of ratios defined for q_e^L via (4.6). Then q_e^L is a global limit cycle for the HDS (3.4) under SB-PR with parameters (L, \mathbf{r}) .*

PROOF OF LEMMA 4.4. For the HDS under SB-PR with parameters $(L, \mathbf{r}) \in \mathbb{N} \times \mathcal{R}$, define the operator $\Gamma_i : \mathbb{R}_+^K \rightarrow \mathbb{R}_+^K$, $i \in \mathcal{I}^L$, mapping the queue length at the polling epoch of stage i to that at the polling epoch of stage $i + 1$. Note that during the busy time of stage i , queue $p(i)$ decreases at rate $\mu_{p(i)} - \lambda_{p(i)}$, and any other queue $k \neq p(i)$ increases at rate λ_k . If q is the queue length at the polling epoch of stage i , then the busy time at stage i lasts for $r_i q_{p(i)} / (\mu_{p(i)} - \lambda_{p(i)})$ units of time, which is the time it takes to reduce queue $p(i)$ to $(1 - r_i)q_{p(i)}$. During the switchover time from stage i to stage $i + 1$, each queue $k \in \mathcal{K}$ increases at rate λ_k , and the switching takes s_i unit of time.

For $i \in \mathcal{I}^L$, let

$$\Gamma_i(q) := \mathcal{A}_i q + \mathcal{B}_i,$$

where \mathcal{A}_i is the $K \times K$ square matrix and $\mathcal{B}_i \in \mathbb{R}^K$ are given by

$$\mathcal{A}_i := \begin{array}{ccccccc} & & & \text{\scriptsize } p(i)\text{th column} & & & \\ \begin{array}{c} 1 \\ 0 \\ \vdots \\ 0 \\ \vdots \\ 0 \\ 0 \\ 0 \end{array} & \begin{array}{c} 0 \\ 1 \\ \vdots \\ 0 \\ \vdots \\ 0 \\ 0 \\ 0 \end{array} & \begin{array}{c} \cdots \\ \cdots \\ \ddots \\ \cdots \\ \cdots \\ \cdots \\ \cdots \\ \cdots \end{array} & \begin{array}{c} \lambda_1 \frac{r_i}{\mu_{p(i)} - \lambda_{p(i)}} \\ \lambda_2 \frac{r_i}{\mu_{p(i)} - \lambda_{p(i)}} \\ \vdots \\ 1 - r_i \\ \vdots \\ \lambda_{K-1} \frac{r_i}{\mu_{p(i)} - \lambda_{p(i)}} \\ \lambda_K \frac{r_i}{\mu_{p(i)} - \lambda_{p(i)}} \end{array} & \begin{array}{c} \cdots \\ \cdots \\ \ddots \\ \cdots \\ \cdots \\ \cdots \\ \cdots \\ \cdots \end{array} & \begin{array}{c} 0 \\ 0 \\ \vdots \\ 0 \\ \vdots \\ 1 \\ 0 \\ 1 \end{array} & \begin{array}{c} 0 \\ 0 \\ \vdots \\ 0 \\ \vdots \\ 0 \\ 0 \end{array} \\ \text{\scriptsize } p(i)\text{th row,} & & & & & & \end{array} \quad \mathcal{B}_i := \begin{bmatrix} \lambda_1 s_i \\ \vdots \\ \lambda_{p(i)-1} s_i \\ \lambda_{p(i)} s_i \\ \lambda_{p(i)+1} s_i \\ \vdots \\ \lambda_K s_i \end{bmatrix},$$

so that

$$\mathcal{A}_i q := \begin{bmatrix} q_1 + \lambda_1 \frac{r_i q_{p(i)}}{\mu_{p(i)} - \lambda_{p(i)}} \\ \vdots \\ q_{p(i)-1} + \lambda_{p(i)-1} \frac{r_i q_{p(i)}}{\mu_{p(i)} - \lambda_{p(i)}} \\ (1 - r_i) q_{p(i)} \\ q_{p(i)+1} + \lambda_{p(i)+1} \frac{r_i q_{p(i)}}{\mu_{p(i)} - \lambda_{p(i)}} \\ \vdots \\ q_K + \lambda_K \frac{r_i q_{p(i)}}{\mu_{p(i)} - \lambda_{p(i)}} \end{bmatrix}.$$

Let $\Gamma' := \Gamma_{IL} \circ \dots \circ \Gamma_1$ be the composition operator over one server cycle, namely, the operator mapping the value of the queue at the beginning of a server cycle to its value at the beginning of the subsequent server cycle. Then

$$\Gamma'(q) = \mathcal{A}' q + \mathcal{B}'$$

for

$$\mathcal{A}' := \mathcal{A}_{IL} \cdots \mathcal{A}_1 \quad \text{and} \quad \mathcal{B}' := \sum_{i=1}^{IL-1} \left(\prod_{j=i+1}^{IL} \mathcal{A}_j \right) \mathcal{B}_i + \mathcal{B}_{IL}.$$

Since each of the operators Γ_i , $i \in \mathcal{I}^L$, is affine and *positively invariant*, the same is true for Γ' . (An affine operator is positively invariant if it maps \mathbb{R}_+^K into itself; see (Matveev et al., 2016, p.10).) By Lemma 5.1 in Matveev et al. (2016), if $\varrho(\mathcal{A}') < 1$, where $\varrho(\mathcal{A}')$ denotes the spectral radius of the matrix \mathcal{A}' , then the positively invariant affine operator $\Gamma'(q)$ is a contraction mapping in \mathbb{R}_+^K .

Hence, we next show that $\varrho(\mathcal{A}') < 1$. To this end, observe that \mathcal{A}' does not depend on the switchover times, so that if the switchover times in the system are changed, but the arrival and service rates are kept fixed, then the matrix \mathcal{A}' remains unchanged. In particular, the matrix \mathcal{A}' does not change if the switchover times in the system under consideration are modified to $s_i = 0$ for all $i \in \mathcal{I}^L$, with all other parameters remaining unchanged.

Consider an auxiliary system that has the same parameters as the system under consideration, except that $(s_i, i \in \mathcal{I}^L) = \mathbf{0}$, and denote its queue process by $q^a := \{q^a(t) : t \geq 0\}$. Let $W(t) := \sum_{k \in \mathcal{K}} q_k^a(t) / \mu_k$ denote the total workload in this auxiliary system at time t . Since all

the switchover times are null, the server is busy at all times in the set $W_+ := \{t : W(t) > 0\}$, so that

$$\dot{W}(t) = \sum_{\ell \in \mathcal{K}, \ell \neq k} \frac{\lambda_\ell}{\mu_\ell} + \frac{\lambda_k - \mu_k}{\mu_k} = \rho - 1 < 0, \quad t \in W_+.$$

Now, \mathcal{A}' is a non-negative square matrix, and so by the Perron-Frobenius theorem (e.g., (Meyer, 2000, Chapter 8.3)), it has a maximal eigenvalue which is strictly positive. This implies that $\varrho(\mathcal{A}') > 0$, and that the eigenvector v associated with $\varrho(\mathcal{A}')$ has strictly positive components. Hence, the eigenvector v is a legitimate state for the queue process q^a .

Take $q^a(0) = v$. Then, at the end of the first server cycle, we have $q^a(u^{(1)}) = \mathcal{A}'v = \varrho(\mathcal{A}')v$, with the second equality holding because v and $\varrho(\mathcal{A}')$ are the associated eigenvector and eigenvalue of \mathcal{A}' . In addition, the workload in the system changes from $W(0) = \sum_{k \in \mathcal{K}} v_k / \mu_k$ to $W(u^{(1)}) = \sum_{k \in \mathcal{K}} \varrho(\mathcal{A}')v_k / \mu_k$. Since the workload process W is strictly decreasing, it holds that

$$\varrho(\mathcal{A}') \sum_{k \in \mathcal{K}} \frac{v_k}{\mu_k} < \sum_{k \in \mathcal{K}} \frac{v_k}{\mu_k},$$

so that $\varrho(\mathcal{A}') < 1$, from which it follows that Γ' is a contraction mapping in \mathbb{R}^K . In turn, under SB-PR (with any control parameters $(L, \mathbf{r}) \in \mathbb{N} \times \mathcal{R}$), there exists a global limit cycle for the HDS³ (3.4) if (and only if) $\rho < 1$. \square

As a consequence of Lemmas 4.1 and 4.4, we also have the following corollary, which in turn, implies the statement of Theorem 3.

COROLLARY 4.1. *Let q_e^L be an L -cycle PE-candidate with ratios \mathbf{r} in (4.6). Then, under SB-PR with parameters (L, \mathbf{r}) , it holds that*

$$\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \psi(q(u)) du = \frac{1}{\tau_L} \int_0^{\tau_L} \psi(q_e^L(u)) du,$$

for any solution q to the HDS (3.4).

PROOF OF THEOREM 3. The proof follows immediately from Corollary 4.1 by taking the SB-PR control parameters to be (L_*, \mathbf{r}_*) for the FCP, or $(L_{\mathcal{N}}, \mathbf{r}_{\mathcal{N}})$ for the RFCP. \square

4.3. Proofs of Lemmas 4.1, 4.2 and 4.3.

PROOF OF LEMMA 4.1. Let L_e denote the number of table cycles contained in the period of the limit cycle q_e^γ . Let $v^{(m-1)}$ denote the beginning epoch of the $((m-1)L_e + 1)$ th table cycle, $m \geq 1$. Define $\tilde{T}^{(m)} := v^{(m)} - v^{(m-1)}$. By construction, $\tilde{T}^{(m)}$ contains exactly L_e table cycles. Since q_e^γ is the limit cycle for q_ϕ^γ , it follows that for any fixed $\epsilon > 0$, there exists $N_\epsilon \geq 1$, such that, for all $m \geq N_\epsilon$,

$$(4.7) \quad \|q_\phi^\gamma(v^{(m-1)} + \cdot) - q_e^\gamma(\cdot)\|_t < \epsilon \quad \text{for all } t > 0, \quad |\tilde{T}^{(m)} - \tau_e^\gamma| < \epsilon,$$

and

$$\left| \int_{v^{(m-1)}}^{v^{(m)}} q_\phi^\gamma(s) ds - \int_0^{\tau_e^\gamma} q_e^\gamma(s) ds \right| < \epsilon.$$

³It is easily seen that the global limit cycle under SB-PR for a system with zero switchover times is trivial, namely, a fixed point; in particular, the limit cycle for this system is the origin.

Since the PE q_e^γ is bounded (componentwise), (4.7) implies that q_ϕ^γ is also bounded. Due to the continuity of ψ , q_e^γ and q_ϕ^γ , the composites $\psi \circ q_e^\gamma$ and $\psi \circ q_\phi^\gamma$ are uniformly continuous over any compact time interval. It follows that for any $\epsilon > 0$, there exists $M_\epsilon \geq N_\epsilon$, such that

$$(4.8) \quad \left| \int_{v^{(m-1)}}^{v^{(m)}} \psi(q_\phi^\gamma(s)) ds - \int_0^{\tau_e^\gamma} \psi(q_e^\gamma(s)) ds \right| < \epsilon, \quad \text{for all } m \geq M_\epsilon.$$

Let $M(t) := \max\{m \geq 1 : v^{(m)} \leq t\}$. Then

$$\frac{1}{t} \int_0^t \psi(q_\phi^\gamma(s)) ds = \frac{1}{t} \sum_{m=1}^{M(t)} \int_{v^{(m-1)}}^{v^{(m)}} \psi(q_\phi^\gamma(s)) ds + \frac{1}{t} \int_{v^{(M(t))}}^t \psi(q_\phi^\gamma(s)) ds.$$

Since $0 \leq t - v^{(M(t))} \leq \tilde{T}^{(M(t)+1)}$ and $\tilde{T}^{(M(t)+1)}$ is bounded by virtue of (4.7), the second term on the right-hand side of the equality above converges to 0 as $t \rightarrow \infty$. Now, for all t large enough, it holds that $M(t) > M_\epsilon$, so that

$$\frac{1}{t} \sum_{m=1}^{M(t)} \int_{v^{(m-1)}}^{v^{(m)}} \psi(q_\phi^\gamma(s)) ds = \frac{1}{t} \sum_{m=1}^{M_\epsilon-1} \int_{v^{(m-1)}}^{v^{(m)}} \psi(q_\phi^\gamma(s)) ds + \frac{1}{t} \sum_{m=M_\epsilon}^{M(t)} \int_{v^{(m-1)}}^{v^{(m)}} \psi(q_\phi^\gamma(s)) ds.$$

For fixed $\epsilon > 0$, M_ϵ is fixed, so that the first term in the right-hand side of the equality converges to 0 as $t \rightarrow \infty$. Applying (4.8) for the second term gives that for t large enough, we get that

$$(4.9) \quad \begin{aligned} \frac{1}{t} \sum_{m=M_\epsilon}^{M(t)} \int_{v^{(m-1)}}^{v^{(m)}} \psi(q_\phi^\gamma(s)) ds &\leq \frac{M(t)}{t} \frac{1}{M(t)} \sum_{m=M_\epsilon}^{M(t)} \left(\int_0^{\tau_e^\gamma} \psi(q_e^\gamma(s)) ds + \epsilon \right) \\ &= \frac{M(t)}{t} \frac{M(t) - M_\epsilon}{M(t)} \left(\int_0^{\tau_e^\gamma} \psi(q_e^\gamma(s)) ds + \epsilon \right) \\ &\leq \left(\frac{1}{\tau_e^\gamma - \epsilon} + o(1) \right) \left(\int_0^{\tau_e^\gamma} \psi(q_e^\gamma(s)) ds + \epsilon \right) \\ &\rightarrow \frac{1}{\tau_e^\gamma - \epsilon} \left(\int_0^{\tau_e^\gamma} \psi(q_e^\gamma(s)) ds + \epsilon \right) \quad \text{as } t \rightarrow \infty. \end{aligned}$$

In the second inequality above, we have used the fact that

$$\begin{aligned} \frac{t}{M(t)} &= \frac{1}{M(t)} \left(\sum_{m=1}^{M_\epsilon-1} \tilde{T}^{(m)} + \sum_{m=M_\epsilon}^{M(t)} \tilde{T}^{(m)} + (t - v^{(M(t))}) \right) \\ &= \frac{1}{M(t)} \sum_{m=M_\epsilon}^{M(t)} \tilde{T}^{(m)} + o(1) \\ &\geq \frac{M(t) - M_\epsilon}{M(t)} (\tau_e^\gamma - \epsilon) + o(1) \rightarrow \tau_e^\gamma - \epsilon \quad \text{as } t \rightarrow \infty. \end{aligned}$$

It follows from (4.9) that

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{m=M_\epsilon}^{M(t)} \int_{v^{(m-1)}}^{v^{(m)}} \psi(q_\phi^\gamma(s)) ds \leq \frac{1}{\tau_e^\gamma - \epsilon} \left(\int_0^{\tau_e^\gamma} \psi(q_e^\gamma(s)) ds + \epsilon \right).$$

We can similarly show that

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \sum_{m=M_\epsilon}^{M(t)} \int_{v^{(m-1)}}^{v^{(m)}} \psi(q_\phi^\gamma(s)) ds \geq \frac{1}{\tau_e^\gamma + \epsilon} \left(\int_0^{\tau_e^\gamma} \psi(q_e^\gamma(s)) ds - \epsilon \right),$$

and so the statement follows by taking $\epsilon \rightarrow 0$. \square

PROOF OF LEMMA 4.2. Let $e_{i,j}$ denote the time elapsed between the departure epoch of stage i and the polling epoch of stage j in q_e^L , $i, j \in \mathcal{I}^L$, namely,

$$e_{i,j} := \begin{cases} s_i + \sum_{\ell=i+1}^{j-1} (s_\ell + b_\ell) & \text{if } i < j \\ s_i + \sum_{\ell=i+1}^{IL} (s_\ell + b_\ell) + \sum_{\ell=1}^{i-1} (s_\ell + b_\ell) & \text{if } i \geq j, \end{cases}$$

with $\sum_{\ell=\ell_1}^{\ell_2} (s_\ell + b_\ell) := 0$ for $\ell_1 > \ell_2$. Then an L -cycle PE-candidate q_e^L necessarily satisfies the following systems of equations

$$(4.10) \quad q_k^L(a_{k_j})(1 - r_{k_j}) + \lambda_k e_{k_j, k_{j+1}} = q_k^L(a_{k_{j+1}}), \quad j = 1, \dots, \dim(\mathbf{k}), k \in \mathcal{K},$$

where $k_{\dim(\mathbf{k})+1} := k_1$. Since the L -cycle PE-candidate parameterized by \mathbf{r} is unique by virtue of Lemma 4.3, the linear system (4.10) admits a unique solution. Hence, solving (4.10) at all possible value of $\mathbf{r} \in \mathcal{R}$ for the corresponding PE gives the entire constraint set of (4.5), because for given (L, \mathbf{r}) , q_e^L is determined by the solution to (4.10), $(q_{e,k}^L(a_{k_j}), j = 1, \dots, \dim(\mathbf{k}), k \in \mathcal{K})$. Thus, (4.5) can be reformulated equivalently as follows.

$$\begin{aligned} \min_{\mathbf{r} \in \mathcal{R}} \quad & \frac{1}{\tau_L} \int_0^{\tau_L} \psi(q_e^L(u)) du \\ \text{s.t.} \quad & q_{e,k}^L(a_{k_j})(1 - r_{k_j}) + \lambda_k e_{k_j, k_{j+1}} = q_{e,k}^L(a_{k_{j+1}}), \quad j = 1, \dots, \dim(\mathbf{k}), k \in \mathcal{K} \\ & q_e^L \text{ is determined by } (q_{e,k}^L(a_{k_j}), j = 1, \dots, \dim(\mathbf{k}), k \in \mathcal{K}). \end{aligned}$$

Now, as was explained in Section 4.1.1, each queue in an optimal PE-candidate must be emptied at least once within a server cycle, and so the vector \mathbf{r} corresponding to an optimal PE-candidate is an element of the set

$$\mathcal{R}' := \left\{ \mathbf{r} \in [0, 1]^{IL} : \sum_{\{i \in \mathcal{I}^L : p(i) = k\}} r_i \geq 1 \quad \text{for all } k \in \mathcal{K} \right\}.$$

Note that \mathcal{R}' is a compact subset of the (non-compact) set \mathcal{R} in (2.3).

Finally, since (4.10) is a system of linear equations for a given \mathbf{r} , its unique solution $(q_{e,k}^L(a_{k_j}), j = 1, \dots, \dim(\mathbf{k}), k \in \mathcal{K})$, is continuous in \mathbf{r} . It follows that, for a given $\epsilon > 0$, there exists a $\delta > 0$, such that for all $\mathbf{r}_1, \mathbf{r}_2 \in \mathcal{R}'$ and their corresponding PE-candidates $q_e^{L,(1)}, q_e^{L,(2)}$, if $\|\mathbf{r}_1 - \mathbf{r}_2\| < \delta$, then $|\psi(q_e^{L,(1)}(u)) - \psi(q_e^{L,(2)}(u))| < \epsilon$ for all $u \in [0, \tau_L)$, so that

$$\left| \frac{1}{\tau_L} \int_0^{\tau_L} \psi(q_e^{L,(1)}(u)) du - \frac{1}{\tau_L} \int_0^{\tau_L} \psi(q_e^{L,(2)}(u)) du \right| < \epsilon.$$

Thus, we established an equivalent formulation for problem (4.5), in which the objective function is continuous over the compact constraint set \mathcal{R}' . It follows from Weierstrass theorem that a global minimum exists. \square

PROOF OF LEMMA 4.3. It follows from the proof of Lemma 4.4 that under SB-PR with parameters $(L, \mathbf{r}) \in \mathbb{N} \times \mathcal{R}$, the HDS converges to a global limit cycle. Thus, an L -cycle PE q_e^L that satisfies (2.3) exists. Moreover, this PE is a global limit cycle, and is therefore the unique PE characterized via (L, \mathbf{r}) . The statement of the lemma follows, because a PE is a PE-candidate by definition. \square

5. Translating SB-PR to the Stochastic System. As discussed in Section 1, we translate SB-PR with control parameters (L, \mathbf{r}) in the deterministic system to the binomial-exhaustive policy with the same control parameters in the stochastic system. To show that the binomial-exhaustive policy with the optimal fluid-control parameters is asymptotically optimal, we first establish general results for admissible policies. Recall that, for a control π , \tilde{Q}_π is the embedded process defined via (2.2). The proof of the following lemma follows from the proof of (Fricker and Jaibi, 1994, Proposition 1), and is thus omitted.

LEMMA 5.1. \tilde{Q}_π is a homogeneous, aperiodic DTMC for any admissible control π .

We remark that the controls considered in Fricker and Jaibi (1994) are assumed to satisfy a certain stochastic monotonicity property, in addition to the conditions in our definition of admissible controls. Thus, the set of controls in this reference is smaller than ours. However, that extra stochastic-monotonicity property does not determine the Markov property of the embedded process \tilde{Q}_π ; see the proof of (Fricker and Jaibi, 1994, Proposition 1).

DEFINITION 5.1. We say that a control π is stable if \tilde{Q}_π is absorbed in a positive recurrent class, regardless of its initial distribution.

It follows from Lemma 5.1 that, for a stable control π ,

$$Q_\pi(m) \Rightarrow \tilde{Q}_\pi(\infty) \quad \text{as } m \rightarrow \infty,$$

where $\tilde{Q}_\pi(\infty)$ is a random variable distributed according to a stationary distribution of the DTMC \tilde{Q}_π . By flow-balance arguments, see, e.g., Boon et al. (2011b), the length of a stationary server-cycle over an L -cycle augmented table T_L has mean

$$\mathbb{E}[T_L] = Ls/(1 - \rho), \quad \text{for } L \geq 1.$$

Clearly, only stable controls are relevant for our (asymptotic) control-optimization problem. However, we note that the stability region of a given control, namely, the set of values of the service and arrival rates for which the system is stable, can be hard to characterize; see Takagi (1988). The most general characterization of the stability region we are aware of was developed in Fricker and Jaibi (1994) (under the aforementioned stochastic-monotonicity property).

5.1. *Sequences of Admissible Controls.* We say that a sequence of controls $\pi = \{\pi^n : n \geq 1\}$ is admissible if π^n is an admissible policy for each $n \geq 1$, and denote the family of all such sequences by Π . For $n \geq 1$ and $U^{(0),n} := 0$, let $U^{(m),n}$ denote the beginning of the $(m+1)$ st server cycle of the n th system, $m \geq 0$. Then, for $\pi \in \Pi$,

$$\tilde{Q}_{\pi^n}^n(m) := \tilde{Q}_{\pi^n}^n(\bar{U}^{(m),n}), \quad m \geq 0,$$

is a DTMC for all $n \geq 1$ by Lemma 5.1. If, in addition, the control is stable for each $n \geq 1$, then there exists a stationary distribution for each of the DTMCs in the sequence, and we say that π is stable.

For the queue process in stationarity, the server-cycle length T_L^n (when the control is designed for an L -cycle augmented table) has mean $nsL/(1 - \rho)$, and for $\bar{T}_L^n := T_L^n/n$,

$$(5.1) \quad \mathbb{E}[\bar{T}_L^n] = sL/(1 - \rho),$$

which is equal to the equilibrium cycle length τ_L in any L -cycle PE of the fluid model.

In order for a sequence of controls $\pi \in \Pi$ to be asymptotically optimal, it must be stable and the sequence of corresponding stationary distributions $\{\tilde{Q}_{\pi^n}^n(\infty) : n \geq 1\}$ must be tight in \mathbb{R}_+^K . However, we remark at the outset that, even if π is stable and $\{\tilde{Q}_{\pi^n}^n(\infty) : n \geq 1\}$ is tight, there is no guarantee that there exists a global limit cycle for any of the resulting fluid limits, because the limits as $n \rightarrow \infty$ and as $t \rightarrow \infty$ need not commute.

5.1.1. *L-Cyclic Controls for the Restricted Problem.* As is the case for the unrestricted problem, for a sequence of controls to be asymptotically optimal with respect to the restricted optimal-control problem, that sequence must be stable, and the corresponding sequence of stationary distributions must be tight. The difference between the two versions of the optimal-control problem is that, in the restricted problem, we have fixed values of table cycles L which we target.

Let q_e denote a PE for a fluid limit when the sequence of controls is a stable sequence π , and when $\tilde{Q}_{\pi^n}^n(0) \stackrel{d}{=} \tilde{Q}_{\pi^n}^n(\infty)$. The fact that a fluid limit for such a sequence exists follows from Lemma 3.1 because the sequence of initial distribution is stationary, and is assumed to be tight, for the reason described above. From the asymptotic perspective, there is clearly no point in considering L -cycle controls which give rise in the limit to PE that have a period that does not divide τ_L . (We always allow the period of an L -cycle PE to be smaller than the cycle length.) Thus, when solving the restricted problem over a set $\mathcal{N} \subset \mathbb{N}$, we should only consider sequences of admissible controls that give rise to L -cycle PE for $L \in \mathcal{N}$, which motivates considering the following family of controls.

DEFINITION 5.2. *A sequence of admissible controls $\pi \in \Pi$ is said to be L -cyclic if any fluid limit of $\{\tilde{Q}_{\pi^n}^n : n \geq 1\}$ with initial condition $\tilde{Q}_{\pi^n}^n(0) \stackrel{d}{=} \tilde{Q}_{\pi^n}^n(\infty)$, $n \geq 1$, is an L -cycle PE.*

We denote the subset of L -cyclic controls by Π^L .

5.2. *SB-PR and the Corresponding Binomial-Exhaustive Policy.* A FWLLN for the binomial-exhaustive policy follows easily from Proposition 3.1, as the next corollary shows.

COROLLARY 5.1 (FWLLN under binomial-exhaustive). *Let $\{Q^n : n \geq 1\}$ denote a sequence of queues where, for each $n \geq 1$, the system operates under the binomial-exhaustive policy with the same parameters $(L, \mathbf{r}) \in \mathbb{N} \times \mathcal{R}$. If $\tilde{Q}^n(0) \Rightarrow q(0)$ in \mathbb{R}_+^K , then $\tilde{Q}^n \Rightarrow q$ in D^K , where q is the fluid queue process under SB-PR with parameters (L, \mathbf{r}) and initial condition $q(0)$.*

PROOF. We verify that the condition in Proposition 3.1 holds under SB-PR, namely, $\bar{D}_i^{(m),n} \Rightarrow d_i^{(m)}$ in \mathbb{R}_+ as $n \rightarrow \infty$, for all $m \geq 1$ and $i \in \mathcal{I}^L$. To this end, consider the first departure epoch, which is also the first busy time of the server at stage 1 in the first server cycle, i.e., time $\bar{D}_1^{(1),n} = \bar{B}_1^{(1),n}$, for $n \geq 1$. Under the binomial-exhaustive policy, all the arrivals to queue $k \in \mathcal{K}$ during the service time of a customer from that same queue are served as well, and so the *total service time* of each served customer and all the arrivals during his service time is distributed like a busy period in an $M/G/1$ queue that has arrival rate λ_k and service rate μ_k .

For each stage $i \in \mathcal{I}^L$ and the corresponding queue $p(i)$, denote by $\Theta_{p(i)}^{(\ell)}$ the busy period “generated” by the service of the ℓ th served customer in this queue. Let $\{Y_i^{(\ell)} : \ell \geq 1\}$ be a sequence of i.i.d. Bernoulli r.v.’s with success probability r_i . We use Y_i and $\Theta_{p(i)}$ to denote corresponding generic random variables. Then

$$\bar{B}_1^{(1),n} = \frac{1}{n} \sum_{\ell=1}^{Q_{p(1)}^n(0)} \Theta_{p(1)}^{(\ell)} Y_1^{(\ell)} \Rightarrow q_{p(1)}(0) \mathbb{E} [\Theta_{p(1)} Y_1] \quad \text{as } n \rightarrow \infty,$$

and due to the independence of $\Theta_{p(1)}$ and Y_1 ,

$$(5.2) \quad \bar{B}_1^{(1),n} = \bar{D}_1^{(1),n} \Rightarrow q_{p(1)}(0) \mathbb{E} [\Theta_{p(1)}] \mathbb{E} [Y_1] = r_1 q_{p(1)}(0) / (\mu_{p(1)} - \lambda_{p(1)}) = d_1^{(1)},$$

where the weak convergence holds as $n \rightarrow \infty$. Furthermore, the length of queue $p(1)$ at the end of the busy time is given by

$$\bar{Q}_{p(1)}^n(\bar{D}_1^{(1),n}) = \bar{Q}_{p(1)}^n(0) - \frac{1}{n} \sum_{\ell=1}^{Q_{p(1)}^n(0)} Y_1^{(\ell)} \Rightarrow q_{p(1)}(0) - q_{p(1)}(0)r_1 = q_{p(1)}(d_1^{(1)}) \quad \text{as } n \rightarrow \infty.$$

It follows from the FWLLN for the Poisson process and (5.2) that, for all $k \neq p(1)$,

$$\bar{Q}_k^n(\bar{D}_1^{(1),n}) \Rightarrow q_k(d_1^{(1)}) = q_k(0) + \lambda_k b_1^{(1)} \quad \text{as } n \rightarrow \infty,$$

and that

$$\bar{Q}(\bar{D}_1^{(1),n} + \bar{V}_1^{(1),n}) \Rightarrow q(d_1^{(1)} + s_1) \quad \text{as } n \rightarrow \infty.$$

Continuing with the same line of arguments gives $\bar{D}_i^{(m),n} \Rightarrow d_i^{(m)}$ as $n \rightarrow \infty$, for all $m \geq 1$, $i \in \mathcal{I}^L$, as required. \square

The FWLLN under the binomial-exhaustive policy remains to hold if the condition that the initial queue converges is replaced with the condition that the initial distribution of the queue is equal to its stationary distribution at the beginning of a server cycle. In this case, the resulting fluid limit is the global limit cycle (the unique PE) under the corresponding SB-PR control. This result, stated formally in the following lemma, will be employed in the proofs of our main theorems.

LEMMA 5.2 (interchange of limits). *Let $\{Q^n : n \geq 1\}$ denote a sequence of queues where, for each $n \geq 1$, the system operates under the binomial-exhaustive policy with the same parameters $(L, \mathbf{r}) \in \mathbb{N} \times \mathcal{R}$. Then for any real-valued, continuous, and bounded function f on \mathbb{R}_+^K ,*

$$(5.3) \quad \lim_{m \rightarrow \infty} \lim_{n \rightarrow \infty} \mathbb{E} \left[f \left(\tilde{Q}^n(m) \right) \right] = \lim_{n \rightarrow \infty} \lim_{m \rightarrow \infty} \mathbb{E} \left[f \left(\tilde{Q}^n(m) \right) \right] = f(q_e(a_1)),$$

where q_e is the PE under SB-PR with parameters (L, \mathbf{r}) . In particular, if $\bar{Q}^n(0) \stackrel{d}{=} \bar{Q}^n(\infty)$ for all $n \geq 1$, then $\bar{Q}^n \Rightarrow q_e$ in D^K as $n \rightarrow \infty$.

PROOF. The key to the proof is the fact that $\mathbb{E}[\bar{Q}^n(\bar{A}_1^n)] = q_e(a_1)$ for all $n \geq 1$. This fact, which is established in Lemma 7.3, implies that $\sup_n \mathbb{E}[\bar{Q}^n(\bar{A}_1^n)] < \infty$. It follows from Markov's inequality that $\{\bar{Q}^n(\bar{A}_1^n) : n \geq 1\}$ is UI, and thus tight in \mathbb{R}_+^K . Since $\bar{Q}^n(0) = \bar{Q}^n(\bar{A}_1^n)$ by definition, $\{\bar{Q}^n(0) : n \geq 1\}$ is tight. Further, $\{\bar{Q}^n(U^{(m),n}) : m \geq 0\}$ is a stationary sequence, so that, since $A_1^n = U^{(0),n}$, we have convergence along subsequences $\bar{Q}^n(U^{(m),n_k}) \Rightarrow \bar{Q}(0)$ as $k \rightarrow \infty$, for all $m \geq 0$. Note that, conditional on $\bar{Q}(0)$, the fluid limit \bar{Q} is deterministic, and converges to the global limit cycle q_e as $t \rightarrow \infty$, regardless of the realized value of $\bar{Q}(0)$.

Assume, in order to arrive at a contradiction, that there exists a set $E \subsetneq \mathbb{R}^K$, with $q_e(a_1) \notin E$, such that $\mathbb{P}(E) > 0$, where \mathbb{P} denotes the probability distribution of $\bar{Q}(0)$. Due to the convergence of $\bar{Q}(t)$ to q_e as $t \rightarrow \infty$, there exists an m_0 , such that $\|\bar{Q}(\bar{U}^{(m)}) - q_e(u^{(1)})\| < \epsilon$ w.p.1 for all $m \geq m_0$ and for any $\epsilon > 0$. It follows that, for all m large enough, $\bar{Q}(\bar{U}^{(m)}) \notin E$. Since this holds for all the trajectories \bar{Q} with $\bar{Q}(0) \in E$, it follows that E is a set of transient states, contradicting the stationarity of $\{\bar{Q}(\bar{U}^{(m)}) : m \geq 0\}$. Thus, $\mathbb{P}(E) = 0$, and in turn, $\bar{Q}(0) = q_e(a_1)$ w.p.1. This latter equality holds for all converging subsequences of $\{\bar{Q}^n(0) : n \geq 1\}$, and so it holds for the sequence itself, namely, $\bar{Q}^n(0) \Rightarrow q_e(a_1)$ as $n \rightarrow \infty$, implying (5.3). This, together with the FWLLN in Corollary 5.1 when $\bar{Q}^n(0) \stackrel{d}{=} \bar{Q}^n(\infty)$, $n \geq 1$, implies that $\bar{Q}^n \Rightarrow q_e$ in D^K as $n \rightarrow \infty$. \square

6. Asymptotic Optimality of Binomial-Exhaustive. In this section we consider the global optimal-control problem, which is the subject of Theorem 1 and the corresponding FCP, and the restricted optimal-control problem.

6.1. *Asymptotic Optimality for the Global Problem.* Theorems 4 and 5 below imply Theorem 1. Recall that c_* is the optimal objective value of the FCP.

THEOREM 4 (asymptotic lower bound). $\liminf_{n \rightarrow \infty} \liminf_{t \rightarrow \infty} \bar{C}_{\pi^n}^n(t) \geq c_*$ w.p.1, for any $\pi \in \Pi$.

PROOF. See Section 7.1. □

Recall $\Psi_{\pi^n}^n$ in (2.5), and that π_*^n is the binomial-exhaustive policy with the same parameters (L_*, \mathbf{r}_*) for all $n \geq 1$, where (L_*, \mathbf{r}_*) are the optimal FCP parameters.

THEOREM 5 (asymptotic optimality). *If $\{\bar{\Psi}_{\pi_*^n}^n : n \geq 1\}$ is UI, then*

$$\lim_{n \rightarrow \infty} \lim_{t \rightarrow \infty} \bar{C}_{\pi_*^n}^n(t) = c_* \quad \text{w.p.1.}$$

PROOF. See Section 7.2. □

The following is an immediate corollary to Theorems 4 and 5 (alternatively, to Theorem 1), Corollary 2.1, Proposition 4.1, and Theorem 2.

COROLLARY 6.1. *Assume that Assumption 2 holds and that the basic table is cyclic. Then the exhaustive policy is asymptotically optimal under either of the following:*

- (i) *For some $p \geq 1$, $\psi(x) = O(\|x\|^p)$ and is separable convex, and in addition, there exists an $\epsilon > 0$ such that $\mathbb{E}[e^{tS_k}] < \infty$ for all $t \in (-\epsilon, \epsilon)$ and for all $k \in \mathcal{K}$.*
- (ii) *$\psi(x) = O(\|x\|)$, and in addition, $\mathbb{E}[S_k^2] < \infty$ for all $k \in \mathcal{K}$.*

6.2. *Asymptotic Optimality for the Restricted Problem.* Recall that $L_{\mathcal{N}}$ is the number of table cycles contained in one server cycle of $q_{\mathcal{N}}$, and that $\mathbf{r}_{\mathcal{N}}$ is the vector of proportion reductions at each stage in $q_{\mathcal{N}}$. Let $\pi_{\mathcal{N}} := \{\pi_{\mathcal{N}}^n : n \geq 1\}$ denote the sequence of binomial-exhaustive policies with parameters $(L_{\mathcal{N}}, \mathbf{r}_{\mathcal{N}})$. We then have the following asymptotic optimality result for the restricted class of admissible controls. The proof of this result follows similar lines of arguments to those in the proofs of Theorems 4 and 5, and is therefore omitted.

THEOREM 6 (asymptotic optimality for the restricted problem). *For all $\pi \in \bigcup_{L \in \mathcal{N}} \Pi^L$ it holds that*

$$\liminf_{n \rightarrow \infty} \liminf_{t \rightarrow \infty} \bar{C}_{\pi^n}^n(t) \geq c_{\mathcal{N}} \quad \text{w.p.1,}$$

for $c_{\mathcal{N}}$ in (4.3). *If, in addition, $\{\bar{\Psi}_{\pi_{\mathcal{N}}^n}^n : n \geq 1\}$ is UI, then*

$$\lim_{n \rightarrow \infty} \lim_{t \rightarrow \infty} \bar{C}_{\pi_{\mathcal{N}}^n}^n(t) = c_{\mathcal{N}} \quad \text{w.p.1.}$$

Following the same lines of arguments as in Corollary 2.1, Theorem 6 implies that if $\{\bar{\Psi}_{\pi_{\mathcal{N}}^n}^n : n \geq 1\}$ is UI, then $\pi_{\mathcal{N}}$ is asymptotically optimal among the restricted class of admissible controls. In particular, (2.4) holds. This, together with Theorem 2, implies the following corollary.

COROLLARY 6.2. *Suppose that $\psi(x) = O(\|x\|^p)$, for some $p \geq 1$, and that Assumption 2 holds. If, for some $\epsilon > 0$, $\mathbb{E}[e^{tS_k}] < \infty$ for all $t \in (-\epsilon, \epsilon)$ and for all $k \in \mathcal{K}$, then the binomial-exhaustive policy with parameters $(L_{\mathcal{N}}, \mathbf{r}_{\mathcal{N}})$ is asymptotically optimal among $\bigcup_{L \in \mathcal{N}} \Pi^L$.*

6.3. Summary of Established Asymptotic Optimality Results. We summarize the conditions and results of Corollaries 6.1 and 6.2 in Table 1.

	Corollary 6.1 (i)	Corollary 6.1 (ii)	Corollary 6.2
Admissible controls	Π	Π	$\bigcup_{L \in \mathcal{N}} \Pi^L$
Cost function	polynomial growth and separable convex	linear growth	polynomial growth
Basic table	cyclic	cyclic	general
Service time distributions	finite m.g.f.'s	second moments	finite m.g.f.'s
Optimal control	$(L, \mathbf{r}) = (1, \mathbf{1})$ (exhaustive)	$(L, \mathbf{r}) = (1, \mathbf{1})$ (exhaustive)	binomial-exhaustive with parameters $(L_{\mathcal{N}}, \mathbf{r}_{\mathcal{N}})$

TABLE 1
Established asymptotic optimality results

7. Proofs of the Main Results. In this section we prove Theorems 4 and 5, from which Theorem 1 follows, and Theorem 2. Some technical results which are employed in the proofs are proved in the appendix.

7.1. Proof of Theorem 4. To establish Theorem 4, it is sufficient restrict attention to sequences of admissible controls $\pi \in \Pi$ under which the corresponding sequences of embedded stationary DTMC's $\{\tilde{Q}_{\pi^n}^n(\infty) : n \geq 1\}$ are tight; the set of such controls π is not empty due to Lemma 5.2. Take $\tilde{Q}_{\pi^n}^n(0) \stackrel{d}{=} \tilde{Q}_{\pi^n}^n(\infty)$ for each $n \geq 1$. Then $\{\tilde{Q}_{\pi^n}^n(0) : n \geq 1\}$ is tight, so that $\{\tilde{Q}_{\pi^n}^n : n \geq 1\}$ is C -tight in D^K by Lemma 3.1.

To decrease the notational burden, we fix a sequence of admissible controls π and a corresponding converging subsequence of $\{\tilde{Q}_{\pi^n}^n : n \geq 1\}$, but we remove the subscript π^n from the notation, and denote the converging subsequence by a superscript ℓ . For example, $\tilde{Q}^\ell := \tilde{Q}_{\pi^{n_\ell}}^{n_\ell}$ denotes the fluid-scaled queue process in system n_ℓ , $\ell \geq 1$, operating under the control π^{n_ℓ} in the converging subsequence of $\{\tilde{Q}_{\pi^{n_\ell}}^{n_\ell} : \ell \geq 1\}$.

Let \tilde{Q} denote the limit of $\{\tilde{Q}^\ell : \ell \geq 1\}$, and let α^ℓ denote the stationary distribution of the corresponding embedded DTMC $\{\tilde{Q}^\ell(m) : m \geq 0\}$. Since each process in the pre-limit is stationary, the limit $\{\tilde{Q}(m) : m \geq 0\}$ of this subsequence of DTMCs is also stationary; we denote the corresponding stationary distribution by α . For $r \geq 0$, let $\mathcal{B}(r)$ denote a ball in \mathbb{R}_+^K with positive α -measure, namely, $\alpha(\mathcal{B}(r)) \in (0, 1]$, and let $B^o = (B_1^o, \dots, B_K^o)$ denote the center of this ball. Note that we do not rule out the case where $r = 0$, which is tantamount to $\mathcal{B}(r)$ being a point in \mathbb{R}_+^K and the limiting distribution α having a point mass on B^o .

Due to the weak convergence of $\{\tilde{Q}^\ell : \ell \geq 1\}$ to \tilde{Q} , we have

$$(7.1) \quad \lim_{\ell \rightarrow \infty} \mathbb{P}(\tilde{Q}^\ell(0) \in \mathcal{B}(r)) = \mathbb{P}(\tilde{Q}(0) \in \mathcal{B}(r)),$$

so that

$$\lim_{\ell \rightarrow \infty} \alpha^\ell(\mathcal{B}(r)) = \alpha(\mathcal{B}(r)) > 0.$$

It follows from (7.1) that $\alpha^\ell(\mathcal{B}(r)) > 0$ for all ℓ large enough, so that $\{\tilde{Q}^\ell(m) : m \geq 0\}$ must return to $\mathcal{B}(r)$ infinitely often for any such ℓ . Similarly, there are infinitely many m 's for

which $\tilde{Q}(m) \in \mathcal{B}(r)$. Let

$$N_r^\ell := \inf\{m \geq 1 : \tilde{Q}^\ell(m) \in \mathcal{B}(r)\} \quad \text{and} \quad N_r := \inf\{m \geq 1 : \tilde{Q}(m) \in \mathcal{B}(r)\}.$$

Then for

$$(7.2) \quad \alpha_r^\ell(\cdot) := \mathbb{P}\left(\tilde{Q}^\ell(0) \in \cdot \mid \tilde{Q}^\ell(0) \in \mathcal{B}(r)\right) \quad \text{and} \quad \alpha_r(\cdot) := \mathbb{P}\left(\tilde{Q}(0) \in \cdot \mid \tilde{Q}(0) \in \mathcal{B}(r)\right),$$

we have

$$(7.3) \quad \lim_{\ell \rightarrow \infty} \mathbb{E}_{\alpha_r^\ell}[N_r^\ell] = \lim_{\ell \rightarrow \infty} \frac{1}{\alpha^\ell(\mathcal{B}(r))} = \frac{1}{\alpha(\mathcal{B}(r))} = \mathbb{E}_{\alpha_r}[N_r].$$

Define the following *first return times* to $\mathcal{B}(r)$

$$(7.4) \quad \begin{aligned} \bar{R}_r^\ell &:= \inf\{\bar{U}^{(m),\ell} > 0 : \bar{Q}^\ell(\bar{U}^{(m),\ell}) \in \mathcal{B}(r)\} \\ \bar{R}_r &:= \inf\{\bar{U}^{(m)} > 0 : \bar{Q}(\bar{U}^{(m)}) \in \mathcal{B}(r)\}. \end{aligned}$$

The next lemma is proved in Section 7.1.1.

LEMMA 7.1. *The subsequence $\{\bar{R}_r^\ell : \ell \geq 1\}$ in (7.4) is UI and satisfies $\bar{R}_r^\ell \Rightarrow \bar{R}_r$. Hence,*

$$\mathbb{E}_{\alpha_r^\ell}[\bar{R}_r^\ell] \rightarrow \mathbb{E}_{\alpha_r}[\bar{R}_r] \quad \text{as } \ell \rightarrow \infty.$$

Observe that the trajectory of \bar{Q} over one return time (from time 0 to \bar{R}_r) is “nearly periodic” for small r , in the sense that both $\bar{Q}(0)$ and $\bar{Q}(\bar{R}_r)$ are in $\mathcal{B}(r)$, although the return time \bar{R}_r may increase as r decreases.

The next lemma, whose proof is given in Section 7.1.1 below, provides an upper bound on the value of \bar{R}_r , and formalizes the observation that \bar{Q} is “nearly periodic,” by proving that it can be made arbitrarily close to a PE-candidate. To emphasize the fact that that PE-candidate depends on the realization of \bar{Q} , and therefore on the sample point $\omega \in \Omega$ (where Ω is the underlying sample space), we make explicit the dependence on ω by adding it to the notation when needed. For example, we write $\bar{Q}(\omega, \cdot)$ for the sample path $\{\bar{Q}(t) : t \geq 0\}$ and $\bar{R}_r(\omega)$ for the realization of the random variable \bar{R}_r corresponding to ω .

LEMMA 7.2. *There exist constants $d_1, d_2 > 0$ such that the following hold.*

(i) $|\bar{R}_r - \tau_{N_r}| \leq d_1 r$ w.p.1.

(ii) *There exists a set $E \subseteq \Omega$, with $P(E) = 1$, such that, for each $\omega \in E$, there exists an $N_r(\omega)$ -cycle PE-candidate q^ω for which*

$$(7.5) \quad \|\bar{Q}(\omega, \cdot) - q^\omega\|_{\bar{R}_r(\omega) \vee \tau_{N_r(\omega)}} \leq d_2 r.$$

Consider the set E in Lemma 7.2, and fix $\omega \in E$. Assume that $\bar{R}_r(\omega) \geq \tau_{N_r(\omega)}$; similar arguments to those below hold for the case $\bar{R}_r(\omega) < \tau_{N_r(\omega)}$.

Clearly, (7.5) implies that

$$(7.6) \quad \max_{k \in \mathcal{K}} |\bar{Q}_k(\omega, t) - q_k^\omega(t)| \leq d_2 r \quad \text{for all } t \in [0, \bar{R}_r(\omega)].$$

Now, $\|\bar{Q}(\omega, 0) - B^o\| \leq r$ because $\bar{Q}(\omega, 0) \in \mathcal{B}(r)$, so that

$$\max_{k \in \mathcal{K}} \bar{Q}_k(\omega, 0) \leq \max_{k \in \mathcal{K}} B_k^o + r.$$

Then (7.5) implies that

$$\max_{k \in \mathcal{K}} q_k^\omega(0) \leq \max_{k \in \mathcal{K}} B_k^o + r + d_2 r,$$

and in turn, for all $t \in [0, \bar{R}_r(\omega)]$,

$$(7.7) \quad \begin{aligned} \max_{k \in \mathcal{K}} q_k^\omega(t) &\leq \max_{k \in \mathcal{K}} q_k^\omega(0) + \max_{k \in \mathcal{K}} \{\lambda_k(1 - \rho_k)\} \tau_{N_r(\omega)} \\ &\leq \max_{k \in \mathcal{K}} B_k^o + r + d_2 r + \max_{k \in \mathcal{K}} \{\lambda_k(1 - \rho_k)\} \tau_{N_r(\omega)}. \end{aligned}$$

Together with (7.6), (7.7) implies that for all $t \in [\tau_{N_r(\omega)}, \bar{R}_r(\omega)]$,

$$(7.8) \quad \begin{aligned} \frac{1}{\tau_{N_r}} \max_{k \in \mathcal{K}} \bar{Q}_k(\omega, t) &\leq \frac{1}{\tau_{N_r(\omega)}} \left(\max_{k \in \mathcal{K}} B_k^o + r + 2d_2 r \right) + \max_{k \in \mathcal{K}} \{\lambda_k(1 - \rho_k)\} \\ &\leq \frac{1}{\tau_1} \left(\max_{k \in \mathcal{K}} B_k^o + r + 2d_2 r \right) + \max_{k \in \mathcal{K}} \{\lambda_k(1 - \rho_k)\} \\ &= d_3 r + d_4, \end{aligned}$$

where d_3 and d_4 are the following constants (that do not depend on ω)

$$d_3 := \frac{1}{\tau_1} (1 + 2d_2) \quad \text{and} \quad d_4 := \frac{1}{\tau_1} \max_{k \in \mathcal{K}} B_k^o + \max_{k \in \mathcal{K}} \{\lambda_k(1 - \rho_k)\}.$$

Thus, for any $\omega \in E$, it holds that

$$(7.9) \quad \begin{aligned} &\max_{k \in \mathcal{K}} \left| \frac{1}{\bar{R}_r(\omega)} \int_0^{\bar{R}_r(\omega)} \bar{Q}_k(\omega, u) du - \frac{1}{\tau_{N_r(\omega)}} \int_0^{\tau_{N_r(\omega)}} q_k^\omega(u) du \right| \\ &= \max_{k \in \mathcal{K}} \left| \frac{1}{\bar{R}_r(\omega)} \left(\int_0^{\tau_{N_r(\omega)}} \bar{Q}_k(\omega, u) du + \int_{\tau_{N_r(\omega)}}^{\bar{R}_r(\omega)} \bar{Q}_k(\omega, u) du \right) - \frac{1}{\tau_{N_r(\omega)}} \int_0^{\tau_{N_r(\omega)}} q_k^\omega(u) du \right| \\ &\leq \max_{k \in \mathcal{K}} \left(\frac{1}{\tau_{N_r(\omega)}} \int_0^{\tau_{N_r(\omega)}} |\bar{Q}_k(\omega, u) - q_k^\omega(u)| du + \frac{1}{\tau_{N_r(\omega)}} \int_{\tau_{N_r(\omega)}}^{\bar{R}_r(\omega)} \bar{Q}_k(\omega, u) du \right) \\ &\leq d_2 r + (d_3 r + d_4) d_1 r, \end{aligned}$$

where the last inequality follows from (7.6), statement (i) in Lemma 7.2, and (7.8). It follows from (7.9) that, for any $\epsilon > 0$, there exists $r_\epsilon > 0$ (that does not depend on ω), such that for all $r < r_\epsilon$,

$$(7.10) \quad \max_{k \in \mathcal{K}} \left| \frac{1}{\bar{R}_r(\omega)} \int_0^{\bar{R}_r(\omega)} \bar{Q}_k(\omega, u) du - \frac{1}{\tau_{N_r(\omega)}} \int_0^{\tau_{N_r(\omega)}} q_k^\omega(u) du \right| < \epsilon.$$

Since q^ω is bounded, (7.5) implies that $\bar{Q}(\omega, \cdot)$ is bounded as well. Therefore, due to the continuity of ψ , q^ω , and of the sample path $\bar{Q}(\omega, \cdot)$, the composite functions $\psi \circ q^\omega$ and $\psi \circ \bar{Q}(\omega, \cdot)$ are both uniformly continuous over compact time intervals. It therefore follows from (7.10) that for any $\delta > 0$, there exists an $\epsilon > 0$ and a corresponding $r_\epsilon > 0$, such that

$$\left| \frac{1}{\bar{R}_r(\omega)} \int_0^{\bar{R}_r(\omega)} \psi(\bar{Q}(\omega, u)) du - \frac{1}{\tau_{N_r(\omega)}} \int_0^{\tau_{N_r(\omega)}} \psi(q^\omega(u)) du \right| < \delta, \quad \text{for all } r \in (0, r_\epsilon),$$

so that

$$(7.11) \quad \frac{1}{\bar{R}_r(\omega)} \int_0^{\bar{R}_r(\omega)} \psi(\bar{Q}(\omega, u)) du > c^\omega - \delta \geq c_* - \delta,$$

where

$$c^\omega := \frac{1}{\tau_{N_r(\omega)}} \int_0^{\tau_{N_r(\omega)}} \psi(q^\omega(u)) du$$

is the time-average holding cost of q^ω , and is necessarily no smaller than c_* by the definition of the latter term. Hence, due to the regenerative structure of \bar{Q}^ℓ for all $\ell \geq 1$, we have (considering the random elements, and thus dropping ω from the notation)

$$\begin{aligned} \liminf_{\ell \rightarrow \infty} \lim_{t \rightarrow \infty} \bar{C}^\ell(t) &= \liminf_{\ell \rightarrow \infty} \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \psi(\bar{Q}^\ell(u)) du \\ &= \liminf_{\ell \rightarrow \infty} \frac{\mathbb{E}_{\alpha_r^\ell} \left[\int_0^{\bar{R}_r^\ell} \psi(\bar{Q}^\ell(u)) du \right]}{\mathbb{E}_{\alpha_r^\ell} [\bar{R}_r^\ell]} && \text{w.p.1 by renewal-reward theorem} \\ &\geq \frac{\liminf_{\ell \rightarrow \infty} \mathbb{E}_{\alpha_r^\ell} \left[\int_0^{\bar{R}_r^\ell} \psi(\bar{Q}^\ell(u)) du \right]}{\limsup_{\ell \rightarrow \infty} \mathbb{E}_{\alpha_r^\ell} [\bar{R}_r^\ell]} \\ &= \frac{\liminf_{\ell \rightarrow \infty} \mathbb{E}_{\alpha_r^\ell} \left[\int_0^{\bar{R}_r^\ell} \psi(\bar{Q}^\ell(u)) du \right]}{\mathbb{E}_{\alpha_r} [\bar{R}_r]} && \text{by Lemma 7.1} \\ &\geq \frac{\mathbb{E}_{\alpha_r} \left[\liminf_{\ell \rightarrow \infty} \int_0^{\bar{R}_r^\ell} \psi(\bar{Q}^\ell(u)) du \right]}{\mathbb{E}_{\alpha_r} [\bar{R}_r]} && \text{by Fatou's lemma} \\ &= \frac{\mathbb{E}_{\alpha_r} \left[\left(\frac{1}{\bar{R}_r} \int_0^{\bar{R}_r} \psi(\bar{Q}(u)) du \right) \bar{R}_r \right]}{\mathbb{E}_{\alpha_r} [\bar{R}_r]} \\ &> \frac{\mathbb{E}_{\alpha_r} [(c_* - \delta) \bar{R}_r]}{\mathbb{E}_{\alpha_r} [\bar{R}_r]} && \text{on the event } E \text{ by (7.11)} \\ &= c_* - \delta. \end{aligned}$$

Note that the second equality above holds regardless of whether $\mathbb{E}_{\alpha_r^\ell} \left[\int_0^{\bar{R}_r^\ell} \psi(\bar{Q}^\ell(u)) du \right] < \infty$ because ψ is nonnegative; see, e.g., Theorem 2.2.1 and the corresponding remark on p.42 in Tijms (2003). The result follows because δ is arbitrary \square

7.1.1. Proofs of the Auxiliary Results in the Proof of Theorem 4.

PROOF OF LEMMA 7.1. The weak convergence in the statement follows from the continuous mapping theorem applied to the first passage time (Whitt, 2002, Theorem 13.6.4). To prove the convergence of the means, let $\bar{Q}^\ell(0)$ be distributed according to α^ℓ , and $\bar{Q}(0)$ be distributed according to α , for α^ℓ and α in (7.2).

The length of the return time \bar{R}_r^ℓ consists of the total time the server spends serving each queue k , $k \in \mathcal{K}$, plus the total switchover time in N_r^ℓ table cycles. Let G_k^ℓ denote the number of customers served at queue k over the time interval $[0, \bar{R}_r^\ell]$, and $\bar{G}_k^\ell := G_k^\ell / \ell$. It holds that

$$(7.12) \quad \bar{Q}_k^\ell(0) + \bar{P}_k^\ell \left(\frac{1}{\ell} \sum_{\nu=1}^K \sum_{j=1}^{G_\nu^\ell} S_\nu^{(j)} + \frac{1}{\ell} \sum_{\nu=1}^{N_r^\ell} \sum_{i=1}^I V_i^{(\nu), \ell} \right) - \bar{G}_k^\ell = \bar{Q}_k^\ell(\bar{R}_r^\ell), \quad k \in \mathcal{K},$$

where

$$\bar{R}_r^\ell = \frac{1}{\ell} \sum_{\nu=1}^K \sum_{j=1}^{G_\nu^\ell} S_\nu^{(j)} + \frac{1}{\ell} \sum_{\nu=1}^{N_r^\ell} \sum_{i=1}^I V_i^{(\nu),\ell}.$$

Since \bar{Q} is stationary, both $\bar{Q}^\ell(0)$ and $\bar{Q}^\ell(\bar{R}_r^\ell)$ are distributed according to α_r^ℓ , so that

$$\mathbb{E}_{\alpha_r^\ell} [\bar{Q}_k^\ell(0)] = \mathbb{E}_{\alpha_r^\ell} [\bar{Q}_k^\ell(\bar{R}_r^\ell)], \quad k \in \mathcal{K}.$$

Thus, taking expectations in (7.12) and applying Wald's equation give

$$\mathbb{E}_{\alpha_r^\ell} [\bar{G}_k^\ell] = \lambda_k \left(\sum_{\nu=1}^K \frac{1}{\mu_\nu} \mathbb{E}_{\alpha_r^\ell} [\bar{G}_\nu^\ell] + \mathbb{E}_{\alpha_r^\ell} [N_r^\ell] s \right), \quad k \in \mathcal{K}.$$

It follows that $\mathbb{E}_{\alpha_r^\ell} [\bar{G}_k^\ell] = \frac{\lambda_k s}{1-\rho} \mathbb{E}_{\alpha_r^\ell} [N_r^\ell]$, so that

$$(7.13) \quad \mathbb{E}_{\alpha_r^\ell} [\bar{R}_r^\ell] = \frac{s}{1-\rho} \mathbb{E}_{\alpha_r^\ell} [N_r^\ell], \quad k \in \mathcal{K}.$$

Similar flow equation holds for the subsequential limit process \bar{Q} . Since the sample paths of \bar{Q} are of the form (3.1) by Lemma 3.1, the process \bar{Q} satisfies

$$\bar{Q}_k(t) = \bar{Q}_k(0) + \lambda_k t - \mu_k \bar{\mathcal{B}}_k(t), \quad t \geq 0, \quad k \in \mathcal{K},$$

where $\bar{\mathcal{B}}_k := \{\bar{\mathcal{B}}_k(t) : t \geq 0\}$ is of the form

$$(7.14) \quad \bar{\mathcal{B}}_k(t) = \int_0^t b_k(u) du,$$

for a piecewise-constant function $b_k : \mathbb{R}_+ \rightarrow \{0, 1\}$. Then, by definition of \bar{R}_r , we have

$$(7.15) \quad \bar{Q}_k(\bar{R}_r) = \bar{Q}_k(0) + \lambda_k \bar{R}_r - \mu_k \bar{\mathcal{B}}_k(\bar{R}_r), \quad k \in \mathcal{K},$$

where $\bar{R}_r = \sum_{k=1}^K \bar{\mathcal{B}}_k(\bar{R}_r) + \bar{N}_r s$. As both $\bar{Q}_k(0)$ and $\bar{Q}_k(\bar{R}_r)$ are distributed according to α , it holds that $\mathbb{E}_{\alpha_r} [\bar{Q}_k(0)] = \mathbb{E}_{\alpha_r} [\bar{Q}_k(\bar{R}_r)]$, and therefore

$$\lambda_k \mathbb{E}_{\alpha_r} \left[\sum_{\nu=1}^K \bar{\mathcal{B}}_\nu(\bar{R}_r) + \bar{N}_r s \right] = \mu_k \mathbb{E}_{\alpha_r} [\bar{\mathcal{B}}_k(\bar{R}_r)], \quad k \in \mathcal{K}.$$

In turn, $\mathbb{E}_{\alpha_r} [\bar{\mathcal{B}}_k(\bar{R}_r)] = \frac{\rho_k s}{1-\rho} \mathbb{E}_{\alpha_r} [\bar{N}_r]$, so that

$$(7.16) \quad \mathbb{E}_{\alpha_r} [\bar{R}_r] = \frac{s}{1-\rho} \mathbb{E}_{\alpha_r} [\bar{N}_r], \quad k \in \mathcal{K}.$$

Since $\mathbb{E}_{\alpha_r^\ell} [N_r^\ell] \rightarrow \mathbb{E}_{\alpha_r} [N_r]$ as $\ell \rightarrow \infty$ by (7.3), it follows from (7.13) and (7.16) that $\mathbb{E}_{\alpha_r^\ell} [\bar{R}_r^\ell] \rightarrow \mathbb{E}_{\alpha_r} [\bar{R}_r]$ as $\ell \rightarrow \infty$, and the result follows. Finally, since $\bar{R}_r \geq 0$ and $\bar{R}_r^\ell \geq 0$ for all $\ell \geq 1$ w.p.1, the sequence $\{\bar{R}_r^\ell : \ell \geq 1\}$ is UI by Theorem 5.4 in Billingsley (1968). \square

PROOF OF LEMMA 7.2. We prove the two assertions of the lemma separately.

Proof of (i). By (7.15) and the fact that $\|\bar{Q}(\bar{R}_r) - \bar{Q}(0)\| \leq 2r$, it holds that for each $k \in \mathcal{K}$,

$$(7.17) \quad -2r \leq -\mu_k \bar{B}_k(\bar{R}_r) + \lambda_k \bar{R}_r \leq 2r,$$

so that

$$(7.18) \quad \sum_{k \in \mathcal{K}} (-2r/\mu_k + \rho_k \bar{R}_r) \leq \sum_{k \in \mathcal{K}} \bar{B}_k(\bar{R}_r) \leq \sum_{k \in \mathcal{K}} (2r/\mu_k + \rho_k \bar{R}_r) \quad w.p.1.$$

Since \bar{R}_r is the total length of the N_r table cycles, it equals the total time the server spends switching, which is equal to sN_r , and the total time it spends serving in each of the queues. Hence,

$$\bar{R}_r = sN_r + \sum_{k \in \mathcal{K}} \bar{B}_k(\bar{R}_r).$$

It then follows from (7.18) that

$$\begin{aligned} sN_r + \sum_{k \in \mathcal{K}} (-2r/\mu_k + \rho_k \bar{R}_r) &\leq \bar{R}_r \leq sN_r + \sum_{k \in \mathcal{K}} (2r/\mu_k + \rho_k \bar{R}_r) \quad \text{so that} \\ -2r \sum_{k \in \mathcal{K}} \frac{1}{\mu_k} &\leq (1 - \rho) \bar{R}_r - sN_r \leq 2r \sum_{k \in \mathcal{K}} \frac{1}{\mu_k}, \end{aligned}$$

and employing (3.6) gives

$$(7.19) \quad -\frac{2r}{1 - \rho} \sum_{k \in \mathcal{K}} \frac{1}{\mu_k} \leq \bar{R}_r - \tau_{N_r} \leq \frac{2r}{1 - \rho} \sum_{k \in \mathcal{K}} \frac{1}{\mu_k}.$$

Taking $d_1 := 2(1 - \rho)^{-1} \sum_{k \in \mathcal{K}} 1/\mu_k$ proves the first part of the lemma.

Proof of (ii). We show that (7.5) holds w.p.1, so that the event E in the statement exists. To this end, we fix $\omega \in \Omega$, and prove the result by constructing a $N_r(\omega)$ -cycle PE-candidate q^ω such that (7.5) holds for the sample path $\bar{Q}(\omega, \cdot)$. To simplify the notation, the values of all the random elements (variables and processes) below are assumed to be realizations corresponding to that fixed ω , although we remove it from the notation (except for the PE-candidate q^ω we construct).

For the limiting process \bar{Q} , let $\bar{B}_{k_j}^{(m)}$, $j = 1, \dots, \dim(\mathbf{k})$, $m = 1, \dots, N_r$, denote the busy time spent serving queue k at stage k_j in the m th server cycle. By definition,

$$\bar{B}_k(\bar{R}_r) = \sum_{m=1}^{N_r} \sum_{j=1}^{\dim(\mathbf{k})} \bar{B}_{k_j}^{(m)}, \quad k \in \mathcal{K},$$

for \bar{B}_k in (7.14). It follows from (7.17) and (7.19) that for $k \in \mathcal{K}$,

$$-\frac{2r}{\mu_k} + \rho_k \left(\tau_{N_r} - \frac{2r}{1 - \rho} \sum_{k \in \mathcal{K}} \frac{1}{\mu_k} \right) \leq \bar{B}_k(\bar{R}_r) \leq \frac{2r}{\mu_k} + \rho_k \left(\tau_{N_r} + \frac{2r}{1 - \rho} \sum_{k \in \mathcal{K}} \frac{1}{\mu_k} \right),$$

so that, for $\delta_k := \bar{B}_k(\bar{R}_r) - \rho_k \tau_{N_r}$, it holds that

$$(7.20) \quad |\delta_k| \leq 2r \left(\frac{1}{\mu_k} + \frac{\rho_k}{1 - \rho} \sum_{k \in \mathcal{K}} \frac{1}{\mu_k} \right), \quad k \in \mathcal{K}.$$

The proof proceeds by explicitly constructing q^ω . To this end, we first characterize a N_r -cycle closed-curve in \mathbb{R}^K , denoted by \bar{Q}' , whose trajectory is sufficiently close to the sample

path of \bar{Q} (corresponding to the sample point ω). However, that closed curve \bar{Q}' is not necessarily a PE-candidate, as its components may achieve negative values. We then show that only a small perturbation of the trajectory \bar{Q}' , such that the perturbed trajectory remains sufficiently close to \bar{Q} , produces a bona-fide N_r -cycle PE-candidate q^ω .

To construct \bar{Q}' , we first take $\bar{Q}'(0) := \bar{Q}(0)$, and then specify the busy times of \bar{Q}' , such that $\bar{Q}'(\tau_{N_r}) = \bar{Q}'(0)$. (We treat \bar{Q}' as a queue process, similarly to our treatment of the fluid models. Thus, by “busy times” of \bar{Q}'_k we mean the times at which the k th component of \bar{Q}' is decreasing.) This can be easily achieved by solving *flow balance equations* which equate the “inflow” to \bar{Q}' over the N_r table cycles, which occurs at a constant rate λ_k throughout, to the “outflow” over the N_r table cycles, which occurs at constant rate $-\mu_k$ only during the busy times. Let $(\bar{B}'_{k_j}(m), j = 1, \dots, \dim(\mathbf{k}), m = 1, \dots, N_r)$ denote those busy times of \bar{Q}'_k .

(1) For queue k with $\delta_k < 0$, we take $\bar{B}'_{k_1}(1) := \bar{B}_{k_1}(1) + |\delta_k|$; and $\bar{B}'_{k_j}(m) := \bar{B}_{k_j}(m)$, for $j = 1, \dots, \dim(\mathbf{k}), m = 1, \dots, N_r$, and $j + m > 2$. Thus, except for its first busy time, all other busy times of \bar{Q}'_k are equal to those of \bar{Q}_k .

(2) For queue k with $\delta_k > 0$, we take $\bar{B}'_{k_j}(\hat{m}) := \bar{B}_{k_j}(\hat{m}) - \delta_k$ for some $\hat{j} \in \{1, \dots, \dim(\mathbf{k})\}$ and $\hat{m} \in \{1, \dots, N_r\}$ with $\bar{B}_{k_j}(\hat{m}) \geq \delta_k$. (Such \hat{j} and \hat{m} exist for sufficiently small r due to (7.18) and (7.19).) and as $\bar{B}'_{k_j}(m) := \bar{B}_{k_j}(m)$ for $j = 1, \dots, \dim(\mathbf{k}), m = 1, \dots, N_r, j \neq \hat{j}$, and $m \neq \hat{m}$. Thus, \bar{Q}'_k and \bar{Q}_k have the same busy times, except for one busy time, which is shorter for \bar{Q}'_k by δ_k .

(3) For queue k with $\delta_k = 0$, we take $\bar{B}'_{k_j}(m) := \bar{B}_{k_j}(m)$ for all $j = 1, \dots, \dim(\mathbf{k})$ and $m = 1, \dots, N_r$. In particular, \bar{Q}'_k and \bar{Q}_k have the same busy times.

Observe that the busy times of \bar{Q}' satisfy the flow balance at all queues, i.e.,

$$\sum_{m=1}^{N_r} \sum_{j=1}^{\dim(\mathbf{k})} \bar{B}'_{k_j}(m) = \rho_k \tau_{N_r} \quad k \in \mathcal{K}.$$

so that $\bar{Q}'(\tau_{N_r}) = \bar{Q}'(0)$. In addition, we will show that

$$(7.21) \quad \|\bar{Q} - \bar{Q}'\|_{\bar{R}_r \vee \tau_{N_r}} \leq \left(\frac{2rK}{1-\rho} \sum_{k \in \mathcal{K}} \frac{1}{\mu_k} \right) \max_{k \in \mathcal{K}} \{\mu_k\}.$$

However, before proving (7.21) we show that we can use this inequality to construct a PE-candidate q^ω as in the statement of the lemma. To this end, let

$$(7.22) \quad q^\omega := \bar{Q}' + \Delta \quad \text{for } \Delta := \left(\frac{2rK}{1-\rho} \sum_{k \in \mathcal{K}} \frac{1}{\mu_k} \right) \max_{k \in \mathcal{K}} \{\mu_k\}.$$

Since $\bar{Q} \geq 0$, it follows from (7.21) that $q^\omega \geq 0$, and since \bar{Q}' is a closed curve in \mathbb{R}^K , so is q^ω . Finally, q^ω clearly satisfies the fluid model equations (3.1), and is therefore a bona-fide PE-candidate. Combining (7.21) and (7.22) gives

$$\|\bar{Q} - q^\omega\|_{\bar{R}_r \vee \tau_{N_r}} \leq \left(\frac{4rK}{1-\rho} \sum_{k \in \mathcal{K}} \frac{1}{\mu_k} \right) \max_{k \in \mathcal{K}} \{\mu_k\},$$

so that (7.5) follows by setting $d_2 := \left(\frac{4K}{1-\rho} \sum_{k \in \mathcal{K}} \frac{1}{\mu_k} \right) \max_{k \in \mathcal{K}} \{\mu_k\}$.

To finish the proof of the lemma, it remains to justify (7.21). To this end, note that \bar{Q}' and \bar{Q} follow identical trajectories from initialization until some busy time differs, namely, when $\bar{B}'_{\ell_j}(m) \neq \bar{B}_{\ell_j}(m)$, for some $\ell \in \mathcal{K}, j \in \{1, \dots, \dim(\mathbf{k})\}$, and $m \in \{1, \dots, N_r\}$. By construction,

$|\bar{B}'_{\ell_j(m)} - \bar{B}_{\ell_j(m)}| = |\delta_\ell|$. Since queue k in either process decreases at rate $\mu_k - \lambda_k$ during the busy times, and increases at rate λ_k everywhere else, it holds that

$$(7.23) \quad \|\bar{Q}_k - \bar{Q}'_k\|_{\bar{D}_{\ell_j(m)} \vee \bar{D}'_{\ell_j(m)}} = |\delta_\ell| \mu_k, \quad k \in \mathcal{K},$$

where $\bar{D}'_{\ell_j(m)}$ (alternatively, $\bar{D}_{\ell_j(m)}$) is the departure epoch immediate after the busy time $\bar{B}'_{\ell_j(m)}$ (alternatively, $\bar{B}_{\ell_j(m)}$) in \bar{Q}' (alternatively, \bar{Q}). Then (7.23) implies that

$$(7.24) \quad \|\bar{Q} - \bar{Q}'\|_{\bar{D}_{\ell_j(m)} \vee \bar{D}'_{\ell_j(m)}} \leq |\delta_\ell| \max_{k \in \mathcal{K}} \{\mu_k\}.$$

After that departure epoch (i.e., $\bar{D}'_{\ell_j(m)}$ for \bar{Q}' , and $\bar{D}_{\ell_j(m)}$ for \bar{Q}), the trajectories of \bar{Q}' and \bar{Q} increase and decrease at the same rate over the same time intervals, until another busy time differs, i.e., $\bar{B}'_{\hat{\ell}_j(m)} \neq \bar{B}_{\hat{\ell}_j(m)}$, for some $\hat{\ell} \in \mathcal{K}$ ($\hat{\ell} \neq \ell$), $\hat{j} \in \{1, \dots, \dim(\hat{\mathcal{L}})\}$, and $\hat{m} \in \{1, \dots, N_r\}$.

Following similar arguments as above, the second difference in the busy times can *further* enlarge the distance between \bar{Q}' and \bar{Q} (from time zero to the departure epoch after the busy time in consideration) component wise by a maximum of $(|\delta_\ell| + |\delta_{\hat{\ell}}|) \max_{k \in \mathcal{K}} \{\mu_k\}$. In particular, define

$$(7.25) \quad \tilde{Q}' := \bar{Q}' - \left(\bar{Q}'(\bar{D}_{\ell_j(m)} \vee \bar{D}'_{\ell_j(m)}) - \bar{Q}(\bar{D}_{\ell_j(m)} \vee \bar{D}'_{\ell_j(m)}) \right).$$

(Note that \tilde{Q}' is the trajectory "shifted" from \bar{Q}' , so that $\tilde{Q}'(\bar{D}_{\ell_j(m)} \vee \bar{D}'_{\ell_j(m)}) = \bar{Q}(\bar{D}_{\ell_j(m)} \vee \bar{D}'_{\ell_j(m)})$.) It holds that

$$(7.26) \quad \sup_{t \in [\bar{D}_{\ell_j(m)} \vee \bar{D}'_{\ell_j(m)}, \bar{D}_{\hat{\ell}_j(m)} \vee \bar{D}'_{\hat{\ell}_j(m)}]} \max_{k \in \mathcal{K}} |\bar{Q}_k(t) - \tilde{Q}'_k(t)| \leq (|\delta_\ell| + |\delta_{\hat{\ell}}|) \max_{k \in \mathcal{K}} \{\mu_k\}.$$

This fact in (7.26), together with (7.24) and (7.25), gives

$$(7.27) \quad \sup_{t \in [\bar{D}_{\ell_j(m)} \vee \bar{D}'_{\ell_j(m)}, \bar{D}_{\hat{\ell}_j(m)} \vee \bar{D}'_{\hat{\ell}_j(m)}]} \max_{k \in \mathcal{K}} |\bar{Q}_k(t) - \bar{Q}'_k(t)| \leq (2|\delta_\ell| + |\delta_{\hat{\ell}}|) \max_{k \in \mathcal{K}} \{\mu_k\}.$$

Since the right-hand side of (7.27) is larger than that of (7.24), we have that

$$\|\bar{Q} - \bar{Q}'\|_{\bar{D}_{\hat{\ell}_j(m)} \vee \bar{D}'_{\hat{\ell}_j(m)}} \leq (2|\delta_\ell| + |\delta_{\hat{\ell}}|) \max_{k \in \mathcal{K}} \{\mu_k\}.$$

The same arguments continue to the end of the N_r th server cycle. Therefore,

$$\begin{aligned} \|\bar{Q} - \bar{Q}'\|_{\bar{R}_r \vee \tau_{N_r}} &\leq \sum_{k \in \mathcal{K}} K |\delta_k| \max_{k \in \mathcal{K}} \{\mu_k\} \\ &\leq K \sum_{k \in \mathcal{K}} 2r \left(\frac{1}{\mu_k} + \frac{\rho_k}{1 - \rho} \sum_{k \in \mathcal{K}} \frac{1}{\mu_k} \right) \max_{k \in \mathcal{K}} \{\mu_k\} \quad \text{by (7.20)} \\ &= \left(\frac{2rK}{1 - \rho} \sum_{k \in \mathcal{K}} \frac{1}{\mu_k} \right) \max_{k \in \mathcal{K}} \{\mu_k\}. \quad \square \end{aligned}$$

7.2. Proof of Theorem 5. Consider a sequence of systems operating under the sequence of binomial-exhaustive policies π_* , each with parameters (L_*, \mathbf{r}_*) . For each $n \geq 1$, let

$\bar{M}^n(t) := \max\{m \geq 1 : \bar{U}^{(m),n} \leq t\}$. Then

$$\begin{aligned}
\lim_{t \rightarrow \infty} \bar{C}_{\pi_*^n}^n(t) &= \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \psi(\bar{Q}_{\pi_*^n}^n(u)) du \\
&= \lim_{t \rightarrow \infty} \frac{\sum_{m=1}^{\bar{M}^n(t)} \int_{\bar{U}^{(m-1),n}}^{\bar{U}^{(m),n}} \psi(\bar{Q}_{\pi_*^n}^n(u)) du + \int_{\bar{U}^{(\bar{M}^n(t)),n}}^t \psi(\bar{Q}_{\pi_*^n}^n(u)) du}{\sum_{m=1}^{\bar{M}^n(t)} \bar{T}^{(m),n} + (t - \bar{U}^{(\bar{M}^n(t)),n})} \\
&= \lim_{t \rightarrow \infty} \frac{\frac{1}{\bar{M}^n(t)} \sum_{m=1}^{\bar{M}^n(t)} \int_{\bar{U}^{(m-1),n}}^{\bar{U}^{(m),n}} \psi(\bar{Q}_{\pi_*^n}^n(u)) du + \frac{1}{\bar{M}^n(t)} \int_{\bar{U}^{(\bar{M}^n(t)),n}}^t \psi(\bar{Q}_{\pi_*^n}^n(u)) du}{\frac{1}{\bar{M}^n(t)} \sum_{m=1}^{\bar{M}^n(t)} \bar{T}^{(m),n} + \frac{1}{\bar{M}^n(t)} (t - \bar{U}^{(\bar{M}^n(t)),n})} \\
(7.28) \quad &= \frac{\mathbb{E}_{\alpha^n} [\bar{\Psi}_{\pi_*^n}^n]}{\mathbb{E}_{\alpha^n} [\bar{T}_{L_*^n}^n]} \quad w.p.1,
\end{aligned}$$

for $\bar{\Psi}_{\pi_*^n}^n$ in (2.5), where the last equality follows because the embedded DTMC converges to its unique stationary distribution α^n , and because the second terms on both the numerator and denominator converge to 0 w.p.1. Indeed, $\bar{Q}_{\pi_*^n}^n(\infty)$ is bounded and $0 \leq t - \bar{U}^{(\bar{M}^n(t)),n} \leq \bar{T}_{L_*^n}^{(\bar{M}^n(t)+1),n} < \infty$ w.p.1 by virtue of (5.1).

Lemma 5.2, the continuity of ψ and the continuous mapping theorem imply that

$$\bar{\Psi}_{\pi_*^n}^n \Rightarrow \int_0^{\tau_*} \psi(q_*(u)) du \quad \text{in } \mathbb{R} \text{ as } n \rightarrow \infty.$$

Thus, since $\mathbb{E}_{\alpha^n} [\bar{T}_{L_*^n}^n] = \tau_*$ for all $n \geq 1$, due to (3.6) and (5.1), the assumed UI of $\{\bar{\Psi}_{\pi_*^n}^n : n \geq 1\}$ and (7.28) give

$$\lim_{n \rightarrow \infty} \lim_{t \rightarrow \infty} \bar{C}_{\pi_*^n}^n(t) = \lim_{n \rightarrow \infty} \frac{1}{\tau_*} \mathbb{E}_{\alpha^n} [\bar{\Psi}_{\pi_*^n}^n] = \frac{1}{\tau_*} \int_0^{\tau_*} \psi(q_*(u)) du = c_* \quad w.p.1. \quad \square$$

7.3. Proof of Theorem 2. Throughout this section, we consider a sequence of systems operating under the binomial-exhaustive policy with control parameters $(L, \mathbf{r}) \in \mathbb{N} \times \mathcal{R}$ for each system n , and its fluid limit q (established in Corollary 5.1). We denote the unique PE (the global limit cycle) of that fluid limit by q_e .

Before proving Theorem 2, we state two technical lemmas—Lemma 7.3 and Lemma 7.4 below. These two lemmas are proved in Theorem 4 and Proposition 6.1 in Hu et al. (2020), and are restated here for completeness. Recall that, if \bar{Q}^n is stationary for each $n \geq 1$, then $\bar{Q}^n \Rightarrow q_e$ in D^K as $n \rightarrow \infty$ by Lemma 5.2.

LEMMA 7.3. *Assume that, for all $n \geq 1$, $\bar{Q}^n(0) \stackrel{d}{=} \bar{Q}^n(\infty)$, so that the process Q^n is stationary. Then*

- (i) $\mathbb{E} [\bar{Q}_k^n(\bar{A}_i^n)] = q_{e,k}(a_i)$ for all $n \geq 1$, $k \in \mathcal{K}$, $i \in \mathcal{I}^L$.
- (ii) If (a) $\mathbb{E} [S_k^2] < \infty$ for all $k \in \mathcal{K}$, (b) $\mathbb{E} [(V_i^n)^2] < \infty$ for all $i \in \mathcal{I}^L$, $n \geq 1$, and (c) $\mathbb{E} [(\bar{V}_i^n)^2] \rightarrow s_i^2$ as $n \rightarrow \infty$ for all $i \in \mathcal{I}^L$, then $\mathbb{E} [Q_k^n(A_i^n) Q_j^n(A_i^n)] < \infty$ for all $n \geq 1$ and

$$\lim_{n \rightarrow \infty} \mathbb{E} [\bar{Q}_k^n(\bar{A}_i^n) \bar{Q}_j^n(\bar{A}_i^n)] = q_{e,k}(a_i) q_{e,j}(a_i), \quad \text{for all } k, j \in \mathcal{K}, \text{ and } i \in \mathcal{I}^L.$$

- (iii) If (a) for each $k \in \mathcal{K}$, there exists $\epsilon_k > 0$ such that $\mathbb{E} [e^{tS_k}] < \infty$ for all $t \in (-\epsilon_k, \epsilon_k)$, (b) $\mathbb{E} [e^{tV_i^n}] < \infty$ for all $t \in \mathbb{R}_+$, $i \in \mathcal{I}^L$, $n \geq 1$, (c) $\mathbb{E} [(\bar{V}_i^n)^\ell] \rightarrow s_i^\ell$ as $n \rightarrow \infty$ for all $\ell \geq 3$, $i \in \mathcal{I}^L$, then $\mathbb{E} [Q_k^n(A_i^n)^\ell] < \infty$ for all $n \geq 1$ and

$$\lim_{n \rightarrow \infty} \mathbb{E} [\bar{Q}_k^n(\bar{A}_i^n)^\ell] = (q_{e,k}(a_i))^\ell \quad \text{for all } \ell \geq 3, k \in \mathcal{K}, \text{ and } i \in \mathcal{I}^L.$$

Recall (see Section 5.2) that, for each stage i , $\Theta_{p(i)}^{(\ell)}$ denotes the busy period “generated” by the service of the ℓ th served customer in queue $p(i)$, which is the queue being polled at stage i .

LEMMA 7.4. *Assume that $\bar{Q}^n(0) \stackrel{d}{=} \tilde{Q}^n(\infty)$, so that Q^n is stationary, for all $n \geq 1$, and consider the corresponding sequence of busy times $\{B_i^n : i \in \mathcal{I}^L, n \geq 1\}$ (over a generic stationary server cycle). Then*

- (i) $\mathbb{E}[\bar{B}_i^n] = r_i q_{e,p(i)}(a_i) \mathbb{E}[\Theta_{p(i)}]$ for all $n \geq 1$, $i \in \mathcal{I}^L$.
(ii) If (a) $\mathbb{E}[S_k^2] < \infty$ for all $k \in \mathcal{K}$, (b) $\mathbb{E}[(V_i^n)^2] < \infty$ for all $i \in \mathcal{I}^L$, $n \geq 1$, and (c) $\mathbb{E}[(\bar{V}_i^n)^2] \rightarrow s_i^2$ as $n \rightarrow \infty$ for all $i \in \mathcal{I}^L$, then $\mathbb{E}[(\bar{B}_i^n)^2] < \infty$ for all $n \geq 1$ and

$$\lim_{n \rightarrow \infty} \mathbb{E}[(\bar{B}_i^n)^2] \rightarrow (r_i q_{e,p(i)}(a_i) \mathbb{E}[\Theta_{p(i)}])^2, \quad \text{for all } i \in \mathcal{I}^L.$$

- (iii) If (a) for each $k \in \mathcal{K}$, there exists $\epsilon_k > 0$ such that $\mathbb{E}[e^{tS_k}] < \infty$ for all $t \in (-\epsilon_k, \epsilon_k)$, (b) $\mathbb{E}[e^{tV_i^n}] < \infty$ for all $t \in \mathbb{R}_+$, $i \in \mathcal{I}^L$, $n \geq 1$, (c) $\mathbb{E}[(\bar{V}_i^n)^\ell] \rightarrow s_i^\ell$ as $n \rightarrow \infty$ for all $\ell \geq 3$, $i \in \mathcal{I}^L$, then $\mathbb{E}[(\bar{B}_i^n)^\ell] < \infty$ for all $n \geq 1$ and

$$\lim_{n \rightarrow \infty} \mathbb{E}[(\bar{B}_i^n)^\ell] = (r_i q_{e,p(i)}(a_i) \mathbb{E}[\Theta_{p(i)}])^\ell \quad \text{for all } \ell \geq 3, i \in \mathcal{I}^L.$$

PROOF OF THEOREM 2. We prove that the two assertions in the theorem hold for the binomial-exhaustive policy under any control parameters $(L, \mathbf{r}) \in \mathbb{N} \times \mathcal{R}$, and so, in particular, for (L_*, \mathbf{r}_*) .

Proof of (i). Since $\psi(x) = O(\|x\|^p)$ for some $p > 1$, there exist $x_0 \in \mathbb{R}_+$ and $M \in \mathbb{R}_+$ such that for all x with $\max_{k \in \mathcal{K}} x_k \geq x_0$, we have

$$(7.29) \quad \psi(x) \leq M \|x\|^p \leq M (K (\max_{k \in \mathcal{K}} x_k)^2)^{\frac{p}{2}} \leq MK^{\frac{p}{2}} \sum_{k \in \mathcal{K}} x_k^p.$$

Let $M' := MK^{\frac{p}{2}}$, and $\mathbf{x}_0 := (x_0, \dots, x_0) \in \mathbb{R}_+^{\mathcal{K}}$. Due to (7.29) and the fact that ψ is non-decreasing, $\bar{\Psi}_{(L, \mathbf{r})}^n$ satisfies

$$(7.30) \quad \begin{aligned} \bar{\Psi}_{(L, \mathbf{r})}^n &= \int_0^{\bar{T}_L^n} \psi(\bar{Q}^n(u)) du \\ &\leq \int_0^{\bar{T}_L^n} \left(\psi(\mathbf{x}_0) \mathbf{1}_{\{\max_{k \in \mathcal{K}} \bar{Q}_k^n(u) < x_0\}} + \left(M' \sum_{k \in \mathcal{K}} (\bar{Q}_k^n(u))^p \right) \mathbf{1}_{\{\max_{k \in \mathcal{K}} \bar{Q}_k^n(u) \geq x_0\}} \right) du \\ &\leq \bar{T}_L^n \psi(\mathbf{x}_0) + \int_0^{\bar{T}_L^n} \left(M' \sum_{k \in \mathcal{K}} (\bar{Q}_k^n(u))^p \right) du. \end{aligned}$$

We start by showing that $\{\bar{T}_L^n \psi(\mathbf{x}_0) : n \geq 1\}$ is UI. To do this, note that the steady-state cycle length \bar{T}_L^n can be represented as $\bar{T}_L^n = \sum_{i \in \mathcal{I}^L} (\bar{B}_i^n + \bar{V}_i^n)$. Similar derivation to that of (5.2) gives

$$(7.31) \quad \bar{B}_i^n \Rightarrow r_i q_{e,p(i)}(a_i) \mathbb{E}[\Theta_{p(i)}] \quad \text{as } n \rightarrow \infty.$$

and, by Lemma 7.4 (i),

$$(7.32) \quad \mathbb{E}[\bar{B}_i^n] = r_i q_{e,p(i)}(a_i) \mathbb{E}[\Theta_{p(i)}] \quad \text{for all } n \geq 1.$$

Since $\bar{B}_i^n \geq 0$ for all $n \geq 1$, the two convergence in (7.31) and in (7.32) imply together that the sequence $\{\bar{B}_i^n : n \geq 1\}$ is UI; e.g., (Billingsley, 1968, Theorem 5.4). Together with the fact that $\{\bar{V}_i^n : n \geq 1\}$ is UI by Assumption 1, we get that $\{\bar{T}_L^n : n \geq 1\}$, and thus $\{\bar{T}_L^n \psi(\mathbf{x}_0) : n \geq 1\}$, is UI.

We next prove that the second term in the right-hand side of (7.30) is UI. To this end, let \tilde{B}_i^n denote the busy time if the exhaustive policy is employed at stage i , when the initial queue length at the corresponding polling epoch, the arrival process to the queue, and the service times of all customers served during B_i^n remain unchanged, so that $\bar{B}_i^n \leq \tilde{B}_i^n$ w.p.1 for all $n \geq 1$. Then

$$\begin{aligned}
M' \int_0^{\bar{T}_L^n} \sum_{k \in \mathcal{K}} (\bar{Q}_k^n(u))^p du &= M' \sum_{i \in \mathcal{I}^L} \int_{\bar{A}_i^n}^{\bar{A}_i^n + \bar{B}_i^n + \bar{V}_i^n} \sum_{k \in \mathcal{K}} (\bar{Q}_k^n(u))^p du \\
&\leq M' \sum_{i \in \mathcal{I}^L} (\bar{B}_i^n + \bar{V}_i^n) \sum_{k \in \mathcal{K}} (\bar{Q}_k^n(\bar{A}_i^n) + \bar{\mathcal{P}}_k^n(\bar{B}_i^n + \bar{V}_i^n))^p \\
(7.33) \quad &\leq M' \sum_{i \in \mathcal{I}^L} (\tilde{B}_i^n + \bar{V}_i^n) \sum_{k \in \mathcal{K}} (\bar{Q}_k^n(\bar{A}_i^n) + \bar{\mathcal{P}}_k^n(\tilde{B}_i^n + \bar{V}_i^n))^p,
\end{aligned}$$

where the first inequality is due to the omission of the service process at stage i .

We next show that, for any $\ell \geq 1$,

$$(7.34) \quad \sup_{n \geq 1} \mathbb{E} \left[(\tilde{B}_i^n + \bar{V}_i^n)^\ell \right] < \infty,$$

and

$$(7.35) \quad \sup_{n \geq 1} \mathbb{E} \left[(\bar{Q}_k^n(\bar{A}_i^n) + \bar{\mathcal{P}}_k^n(\tilde{B}_i^n + \bar{V}_i^n))^\ell \right] < \infty,$$

from which it follows that, for any $\epsilon > 0$,

$$\sup_{n \geq 1} \mathbb{E} \left[(\tilde{B}_i^n + \bar{V}_i^n)^{1+\epsilon} (\bar{Q}_k^n(\bar{A}_i^n) + \bar{\mathcal{P}}_k^n(\tilde{B}_i^n + \bar{V}_i^n))^{p(1+\epsilon)} \right] < \infty$$

by virtue of Hölder's inequality, so that the sequence of bounds in (7.33) is UI.

The inequality in (7.34) follows directly from the fact that \tilde{B}_i^n and \bar{V}_i^n are independent, and both are uniformly bounded in n . Indeed, $\sup_n \mathbb{E}[(\tilde{B}_i^n)^\ell] < \infty$ by Lemma 7.4 (i)–(iii), and $\sup_n \mathbb{E}[(\bar{V}_i^n)^\ell] < \infty$ by Assumption 2. To prove (7.35), note that

$$\begin{aligned}
\mathbb{E} \left[(\bar{Q}_k^n(\bar{A}_i^n) + \bar{\mathcal{P}}_k^n(\tilde{B}_i^n + \bar{V}_i^n))^\ell \right] &= \mathbb{E} \left[\sum_{j=0}^{\ell} \binom{\ell}{j} (\bar{Q}_k^n(\bar{A}_i^n))^{\ell-j} (\bar{\mathcal{P}}_k^n(\tilde{B}_i^n + \bar{V}_i^n))^j \right] \\
(7.36) \quad &\leq \sum_{j=0}^{\ell} \binom{\ell}{j} \mathbb{E} \left[(\bar{Q}_k^n(\bar{A}_i^n))^{(\ell-j)\alpha} \right]^{\frac{1}{\alpha}} \mathbb{E} \left[(\bar{\mathcal{P}}_k^n(\tilde{B}_i^n + \bar{V}_i^n))^{j\beta} \right]^{\frac{1}{\beta}},
\end{aligned}$$

where the equality holds the the Binomial Theorem, and the inequality follows from Hölder's inequality, for $\alpha > 1$ and $1/\alpha + 1/\beta = 1$.

Let $\{\cdot\}$ denote the Stirling numbers of the second type, and recall that, for a Poisson random variable Y with mean ν , it holds that $\mathbb{E}[Y^N] = \sum_{j=1}^N \{N \atop j\} \nu^j$, for $N \in \mathbb{N}$. Thus,

$$\begin{aligned}
(7.37) \quad \mathbb{E} \left[\left(\bar{\mathcal{P}}_k^n(\tilde{B}_i^n + \bar{V}_i^n) \right)^{j\beta} \right] &= \mathbb{E} \left[\mathbb{E} \left[\left(\bar{\mathcal{P}}_k^n(\tilde{B}_i^n + \bar{V}_i^n) \right)^{j\beta} \mid \tilde{B}_i^n + \bar{V}_i^n \right] \right] \\
&= \mathbb{E} \left[\mathbb{E} \left[\left(\frac{1}{n} \mathcal{P}_k^n(\tilde{B}_i^n n + \bar{V}_i^n n) \right)^{j\beta} \mid \tilde{B}_i^n + \bar{V}_i^n \right] \right] \\
&= \mathbb{E} \left[\frac{1}{n^{j\beta}} \sum_{m=0}^{j\beta} \left(\lambda_k(\tilde{B}_i^n n + \bar{V}_i^n n) \right)^m \{j\beta \atop m\} \right] \\
&= \mathbb{E} \left[\sum_{m=0}^{j\beta} \frac{1}{n^{j\beta-m}} \lambda_k^m(\tilde{B}_i^n + \bar{V}_i^n)^m \{j\beta \atop m\} \right] \\
&= \sum_{m=0}^{j\beta} \frac{1}{n^{j\beta-m}} \lambda_k^m \{j\beta \atop m\} \mathbb{E} \left[(\tilde{B}_i^n + \bar{V}_i^n)^m \right] \\
&= \lambda_k^{j\beta} \mathbb{E} \left[(\tilde{B}_i^n + \bar{V}_i^n)^{j\beta} \right] + o(1).
\end{aligned}$$

Plugging (7.37) into (7.36), we get

$$\begin{aligned}
&\mathbb{E} \left[\left(\bar{Q}_k^n(\bar{A}_i^n) + \bar{\mathcal{P}}_k^n(\tilde{B}_i^n + \bar{V}_i^n) \right)^\ell \right] \\
&\leq \sum_{j=0}^{\ell} \binom{\ell}{j} \mathbb{E} \left[\left(\bar{Q}_k^n(\bar{A}_i^n) \right)^{(\ell-j)\alpha} \right]^{\frac{1}{\alpha}} \left(\lambda_k^{j\beta} \mathbb{E} \left[(\tilde{B}_i^n + \bar{V}_i^n)^{j\beta} \right] + o(1) \right)^{\frac{1}{\beta}},
\end{aligned}$$

which is uniformly bounded in n due to Lemma 7.3 (i)–(iii), Lemma 7.4 (i)–(iii), Assumption 2, and the independence of \tilde{B}_i^n and \bar{V}_i^n . Thus, (7.35) holds.

Proof of (ii). For $x \in \mathbb{R}^K$, let $f(x) = \sum_{k=1}^K c_k x_k$. Since $\psi(x) = O(f(x))$, there exist $x_0 \in \mathbb{R}_+$ and $M \in \mathbb{R}_+$, such that for all x with $\max_{k \in \mathcal{K}} x_k \geq x_0$,

$$\psi(x) \leq M f(x) = M \sum_{k \in \mathcal{K}} c_k x_k.$$

As in (7.30), this implies that

$$(7.38) \quad \bar{\Psi}_{(L,r)}^n = \int_0^{\bar{T}_L^n} \psi(\bar{Q}^n(u)) du \leq \bar{T}_L^n \psi(\mathbf{x}_0) + \int_0^{\bar{T}_L^n} \left(M \sum_{k \in \mathcal{K}} c_k \bar{Q}_k^n(u) \right) du.$$

Since $\{\bar{T}_L^n \psi(\mathbf{x}_0) : n \geq 1\}$ was shown to be UI in the proof of part (i) above, we need only show that the sequence corresponding to the second term on the right-hand side of (7.38) is UI. Similarly to the derivation of (7.33), we can bound this term from above via

$$\begin{aligned}
(7.39) \quad &\int_0^{\bar{T}_L^n} \left(M \sum_{k \in \mathcal{K}} c_k \bar{Q}_k^n(u) \right) du \\
&\leq M \sum_{k \in \mathcal{K}} c_k \sum_{i \in \mathcal{I}^L} \left(\bar{Q}_k^n(\bar{A}_i^n) + \bar{\mathcal{P}}_k^n(\tilde{B}_i^n + \bar{V}_i^n) \right) \left(\tilde{B}_i^n + \bar{V}_i^n \right) \\
&= M \sum_{k \in \mathcal{K}} c_k \sum_{i \in \mathcal{I}^L} \left(\bar{Q}_k^n(\bar{A}_i^n) \tilde{B}_i^n + \bar{Q}_k^n(\bar{A}_i^n) \bar{V}_i^n + \bar{\mathcal{P}}_k^n(\tilde{B}_i^n + \bar{V}_i^n) (\tilde{B}_i^n + \bar{V}_i^n) \right).
\end{aligned}$$

We next show that the sequence corresponding to each term in the right-hand side of (7.39) is UI. First, by Lemma 5.2, $\bar{Q}_k^n(\bar{A}_i^n) \Rightarrow q_{e,k}(a_i)$ in \mathbb{R}_+ as $n \rightarrow \infty$. In addition, it follows from (7.31) (setting $r_i = 1$) that

$$(7.40) \quad \tilde{B}_i^n \Rightarrow q_{e,p(i)}(a_i) \mathbb{E} [\Theta_{p(i)}] \quad \text{as } n \rightarrow \infty,$$

so that, by Slutsky's theorem,

$$(7.41) \quad \bar{Q}_k^n(\bar{A}_i^n) \tilde{B}_i^n \Rightarrow q_{e,k}(a_i) q_{e,p(i)}(a_i) \mathbb{E} [\Theta_{p(i)}] \quad \text{as } n \rightarrow \infty.$$

Now,

$$(7.42) \quad \begin{aligned} \mathbb{E} \left[\bar{Q}_k^n(\bar{A}_i^n) \tilde{B}_i^n \right] &= \mathbb{E} \left[\mathbb{E} \left[\bar{Q}_k^n(\bar{A}_i^n) \frac{1}{n} \sum_{j=1}^{Q_{p(i)}^n(A_i^n)} \Theta_{p(i)}^{(j)} \middle| Q^n(A_i^n) \right] \right] \\ &= \mathbb{E} \left[\bar{Q}_k^n(\bar{A}_i^n) \bar{Q}_{p(i)}^n(\bar{A}_i^n) \right] \mathbb{E} [\Theta_{p(i)}] \\ &\rightarrow q_{e,k}(a_i) q_{e,p(i)}(a_i) \mathbb{E} [\Theta_{p(i)}] \quad \text{as } n \rightarrow \infty \quad \text{by Lemma 7.3 (ii)}. \end{aligned}$$

It follows from (7.41), (7.42), and the fact that both $\bar{Q}_k^n(\bar{A}_i^n) \tilde{B}_i^n$ and $q_{e,k}(a_i) q_{e,p(i)}(a_i) \mathbb{E} [\Theta_{p(i)}]$ are non-negative, that the sequence $\{\bar{Q}_k^n(\bar{A}_i^n) \tilde{B}_i^n : n \geq 1\}$ is UI. Second, $\bar{Q}_k^n(\bar{A}_i^n)$ and \bar{V}_i^n being independent implies that

$$\mathbb{E} \left[(\bar{Q}_k^n(\bar{A}_i^n) \bar{V}_i^n)^2 \right] = \mathbb{E} \left[(\bar{Q}_k^n(\bar{A}_i^n))^2 \right] \mathbb{E} \left[(\bar{V}_i^n)^2 \right].$$

Because $\mathbb{E} \left[(\bar{V}_i^n)^2 \right] < \infty$ under Assumption 2, and $\mathbb{E} \left[(\bar{Q}_k^n(\bar{A}_i^n))^2 \right] < \infty$ given $\mathbb{E} [S_k^2] < \infty$ by Lemma 7.3 (ii), the second moment of $\bar{Q}_k^n(\bar{A}_i^n) \bar{V}_i^n$ is finite, implying that $\{\bar{Q}_k^n(\bar{A}_i^n) \bar{V}_i^n : n \geq 1\}$ is UI.

Lastly, for $\bar{\mathcal{P}}_k^n(\tilde{B}_i^n + \bar{V}_i^n)(\tilde{B}_i^n + \bar{V}_i^n)$, note that $\tilde{B}_i^n \Rightarrow q_{e,p(i)}(a_i) \mathbb{E} [\Theta_{p(i)}]$ by (7.40), and $\bar{V}_i^n \Rightarrow s_i$ as $n \rightarrow \infty$. By the FWLLN for Poisson processes, we have

$$(7.43) \quad \bar{\mathcal{P}}_k^n(\tilde{B}_i^n + \bar{V}_i^n) \Rightarrow \lambda_k (q_{e,p(i)}(a_i) \mathbb{E} [\Theta_{p(i)}] + s_i) \quad \text{as } n \rightarrow \infty.$$

By Slutsky's theorem, $\bar{\mathcal{P}}_k^n(\tilde{B}_i^n + s_i)(\tilde{B}_i^n + \bar{V}_i^n) \Rightarrow \lambda_k (q_{e,p(i)}(a_i) \mathbb{E} [\Theta_{p(i)}] + s_i)^2$ as $n \rightarrow \infty$. Next,

$$(7.44) \quad \begin{aligned} \mathbb{E} \left[\bar{\mathcal{P}}_k^n(\tilde{B}_i^n + \bar{V}_i^n)(\tilde{B}_i^n + \bar{V}_i^n) \right] &= \mathbb{E} \left[\mathbb{E} \left[\bar{\mathcal{P}}_k^n(\tilde{B}_i^n + \bar{V}_i^n)(\tilde{B}_i^n + \bar{V}_i^n) \middle| \tilde{B}_i^n + \bar{V}_i^n \right] \right] \\ &= \mathbb{E} \left[\lambda_k (\tilde{B}_i^n + \bar{V}_i^n)^2 \right] \\ &\rightarrow \lambda_k (q_{e,p(i)}(a_i) \mathbb{E} [\Theta_{p(i)}] + s_i)^2 \quad \text{as } n \rightarrow \infty, \end{aligned}$$

where the limit follows from Lemma 7.4 (i)–(ii) and Assumption 2. Since both the pre-limit and limit in (7.43) are non-negative, Theorem 5.4 in Billingsley (1968), (7.43) and (7.44) imply that $\{\bar{\mathcal{P}}_k^n(\tilde{B}_i^n + \bar{V}_i^n)(\tilde{B}_i^n + \bar{V}_i^n) : n \geq 1\}$ is UI, and in turn, so is $\{\bar{\Psi}_{(L,r)}^n : n \geq 1\}$. \square

8. Summary and Future Research. We considered the optimal-control problem of polling systems with large switchover times. Under the large-switchover-time scaling, we established that the binomial-exhaustive policy, with properly chosen control parameters, is asymptotically optimal. Those optimal control parameters are computed by solving an FCP for a related deterministic relaxation, which is described via an HDS, and arises as the fluid limit for a sequence of stochastic systems operating under the corresponding binomial-exhaustive policy. For the important special case in which the basic table is cyclic and the

cost function is separable convex and has at most a polynomial growth, we showed that the exhaustive policy is asymptotically optimal.

The analytical tools in this paper can be useful in characterizing asymptotically optimal controls in other settings. For example, the Stochastic Economic Lot Scheduling Problem (SELSP), can be modeled as a polling system in which backlogged demand implies that the buffer content can be negative; see, e.g., [Federguen and Katalan \(1996\)](#). Further, the stability region of the fluid model for polling systems is easier to characterize than that of the underlying stochastic system, and can therefore be used to study the stability of stochastic polling systems under controls that do not adhere to the conditions in [Fricker and Jaibi \(1994\)](#).

Acknowledgement. We thank Professor Hanoch Levy for valuable conversations regarding the binomial-exhaustive policy, which he proposed in [Levy \(1988\)](#).

APPENDIX A: PROOF OF PROPOSITION 4.1

The proof of Proposition 4.1 involves approximating ψ with piecewise linear functions. Since ψ is assumed to be separable convex, each of its components ψ_k is an increasing convex function mapping \mathbb{R}_+ into itself, and can therefore be approximated over any compact interval by piecewise linear functions of the form

$$(A.1) \quad p_k(x) = \begin{cases} p_k^{(1)}x & \text{if } \alpha_k^{(1)} \leq x < \alpha_k^{(2)} \\ p_k^{(2)}x & \text{if } \alpha_k^{(2)} \leq x < \alpha_k^{(3)} \\ \vdots & \\ p_k^{(N_k-1)}x & \text{if } \alpha_k^{(N_k-1)} \leq x < \alpha_k^{(N_k)} \\ p_k^{(N_k)}x & \text{if } \alpha_k^{(N_k)} \leq x. \end{cases}$$

where $0 \leq p_k^{(1)} < p_k^{(2)} \cdots < p_k^{(N_k)}$ in \mathbb{R}_+ and $0 = \alpha_k^{(1)} < \alpha_k^{(2)} \cdots < \alpha_k^{(N_k)}$ in \mathbb{R}_+ . We refer to $\{\alpha_k^{(\ell)} : 1 \leq \ell \leq N_k\}$ as the *irregular points* (of p_k), since p_k is differentiable everywhere except at those points, and to $\{p_k^{(\ell)} : 1 \leq \ell \leq N_k\}$ as the *coefficients* of p_k .

The next lemma is the key to proving Proposition 4.1; its proof appears in Section A.1 below. We say that $f : \mathbb{R}_+^K \rightarrow \mathbb{R}_+$ is separable and piecewise linear if for $x \in \mathbb{R}_+^K$, $f(x) = \sum_{k=1}^K f_k(x_k)$ and each f_k is a piecewise linear function, mapping \mathbb{R}_+ into itself.

LEMMA A.1. *If the basic table is cyclic and ψ is separable and piecewise linear (in addition to being nondecreasing and continuous), then q_{exh} solves (4.4).*

PROOF OF PROPOSITION 4.1. Fix $L \in \mathbb{N}$. For any L -cycle PE-candidate q_e^L with cycle length τ_L , the trajectory of queue k is bounded from below by 0 and from above by $M_k := \frac{1}{2}(1 - \rho_k)\lambda_k\tau_L^2$, so that

$$\|\psi_k(q_{e,k}^L)\|_{\tau_L} \leq \psi_k(M_k), \quad k \in \mathcal{K}.$$

Hence, for any $\epsilon > 0$, there exist piecewise linear functions p_k and h_k , both mapping \mathbb{R}_+ to itself, such that for all $y \in [0, \psi_k(M_k)]$,

$$0 < p_k(y) - \psi_k(y) < \epsilon/K \quad \text{and} \quad 0 < \psi_k(y) - h_k(y) < \epsilon/K, \quad k \in \mathcal{K}.$$

Let $p, h : \mathbb{R}_+^K \rightarrow \mathbb{R}_+$ be defined for $x \in \mathbb{R}_+^K$ via $p(x) := \sum_{k \in \mathcal{K}} p_k(x_k)$ and $h(x) := \sum_{k \in \mathcal{K}} h_k(x_k)$. It follows that

$$0 < \frac{1}{\tau_L} \int_0^{\tau_L} p(q_e^L(u)) du - \frac{1}{\tau_L} \int_0^{\tau_L} \psi(q_e^L(u)) du < \epsilon,$$

and

$$0 < \frac{1}{\tau_L} \int_0^{\tau_L} \psi(q_e^L(u)) du - \frac{1}{\tau_L} \int_0^{\tau_L} h(q_e^L(u)) du < \epsilon.$$

Consider two L -cycle optimization problems, denoted by P_p and P_h , which replace the objective function ψ in problem (4.5) with p and h , respectively. Since q_{exh} is optimal for both P_p and P_h by Lemma A.1 and ϵ is arbitrary, q_{exh} is a solution to L -cycle optimization problem (4.5). As the latter statement holds for all $L \in \mathbb{N}$, q_{exh} is a solution to the global optimization problem (4.4). \square

A.1. Proof of Lemma A.1. For each $k \in \mathcal{K}$, let ψ_k be in the form of (A.1), with irregular points $\{\alpha_k^{(\ell)} : 1 \leq \ell \leq N_k\}$ and coefficients $\{\psi_k^{(\ell)} : 1 \leq \ell \leq N_k\}$, for some $N_k \in \mathbb{N}$. Fix $L \in \mathbb{N}$. To show that q_{exh} is a solution of the L -cycle optimization problem (4.5), we consider a relaxed problem, in which each queue is optimized without consideration of all other queues. To this end, for each $k \in \mathcal{K}$, we consider the following relaxation to (4.5):

$$(A.2) \quad \min_{q^L \in \mathcal{Q}^L} \frac{1}{\tau_L} \int_0^{\tau_L} \psi_k(q_k^L(u)) du.$$

Note that we have K different optimization problems of the form (A.2)—one for each $k \in \mathcal{K}$. For each of these K optimization problems, let $q_{e,k}^L$ denote the k th component of a solution to the problem (A.2). The closed curve $q_e^L := (q_{e,k}^L, k \in \mathcal{K})$ necessarily gives a lower bound for the optimal objective value in the L -cycle optimization problem (4.5). (However, q_e^L needs not be an element of \mathcal{Q}^L , as it may not be a bona-fide L -cycle PE-candidate.) The statement of the lemma will therefore follow by proving that q_{exh} , which is a feasible solution to (4.5), is a solution to (A.2) for each $k \in \mathcal{K}$.

Fix $k \in \mathcal{K}$. For $\psi_k^{(0)} := 0$, let

$$h_k^{(\ell)} := \psi_k^{(\ell)} - \psi_k^{(\ell-1)}, \quad 1 \leq \ell \leq N_k.$$

The objective function in (A.2) satisfies

$$\begin{aligned} \frac{1}{\tau_L} \int_0^{\tau_L} \psi_k(q_k^L(s)) ds &= \frac{1}{\tau_L} \int_0^{\tau_L} \left[\psi_k^{(1)} q_k^L(s) \mathbf{1}_{\{q_k^L(s) \geq \alpha_k^{(1)}\}} + (\psi_k^{(2)} - \psi_k^{(1)}) q_k^L(s) \mathbf{1}_{\{q_k^L(s) \geq \alpha_k^{(2)}\}} \right. \\ &\quad \left. + \dots + (\psi_k^{(N_k)} - \psi_k^{(N_k-1)}) q_k^L(s) \mathbf{1}_{\{q_k^L(s) \geq \alpha_k^{(N_k)}\}} \right] ds \\ &= \frac{1}{\tau_L} \int_0^{\tau_L} \left[h_k^{(1)} q_k^L(s) \mathbf{1}_{\{q_k^L(s) \geq \alpha_k^{(1)}\}} + h_k^{(2)} q_k^L(s) \mathbf{1}_{\{q_k^L(s) \geq \alpha_k^{(2)}\}} \right. \\ &\quad \left. + \dots + h_k^{(N_k)} q_k^L(s) \mathbf{1}_{\{q_k^L(s) \geq \alpha_k^{(N_k)}\}} \right] ds \\ &= \frac{1}{\tau_L} \sum_{\ell=1}^{N_k} h_k^{(\ell)} \int_0^{\tau_L} q_k^L(s) \mathbf{1}_{\{q_k^L(s) \geq \alpha_k^{(\ell)}\}} ds. \end{aligned}$$

Then for

$$A^{(\ell)}(q_k^L) := \int_0^{\tau_L} q_k^L(s) \mathbf{1}_{\{q_k^L(s) \geq \alpha_k^{(\ell)}\}} ds, \quad 1 \leq \ell \leq N_k,$$

(A.2) is equivalent to

$$(A.3) \quad \begin{aligned} \min_{q^L \in \mathcal{Q}^L} \quad & \frac{1}{\tau_L} \sum_{\ell=1}^{N_k} h_k^{(\ell)} A^{(\ell)}(q_k^L) \\ \text{s.t.} \quad & A^{(\ell)}(q_k^L) = \int_0^{\tau_L} q_k^L(s) \mathbf{1}_{\{q_k^L(s) \geq \alpha_k^{(\ell)}\}} ds, \quad 1 \leq \ell \leq N_k. \end{aligned}$$

Now, $A^{(\ell)}(q_k^L)$ can be further partitioned into L sub-areas over each table cycle. In particular, let

$$a^{(\ell,m)}(q_k^L) := \int_{u^{(m-1)}}^{u^{(m)}} q_k^L(s) \mathbf{1}_{\{q_k^L(s) \geq \alpha_k^{(\ell)}\}} ds, \quad 1 \leq m \leq L, \quad 1 \leq \ell \leq N_k,$$

where $u^{(m-1)}$ denotes the beginning epoch of the m th table cycle. (To facilitate the notation henceforth, the superscript (m) in $u^{(m)}$ is an index for table cycles of q^L over the cycle length $[0, \tau_L]$. This is different from the convention elsewhere in the paper, e.g., in (3.1), where (m) was indexing server cycles.) We can then write

$$A^{(\ell)}(q_k^L) = \sum_{m=1}^L a^{(\ell,m)}(q_k^L), \quad 1 \leq \ell \leq N_k,$$

so that (A.3) can be equivalently written as

$$(A.4) \quad \min_{q^L \in \mathcal{Q}^L} \frac{1}{\tau_L} \sum_{\ell=1}^{N_k} h_k^{(\ell)} \sum_{m=1}^L a^{(\ell,m)}(q_k^L)$$

$$s.t. \quad a^{(\ell,m)}(q_k^L) = \int_{u^{(m-1)}}^{u^{(m)}} q_k^L(s) \mathbf{1}_{\{q_k^L(s) \geq \alpha_k^{(\ell)}\}} ds, \quad 1 \leq m \leq L, \quad 1 \leq \ell \leq N_k.$$

For

$$\tau^{(\ell,m)}(q_k^L) := \int_{u^{(m-1)}}^{u^{(m)}} \mathbf{1}_{\{q_k^L(s) \geq \alpha_k^{(\ell)}\}} ds, \quad 1 \leq m \leq L, \quad 1 \leq \ell \leq N_k,$$

we have that

$$(A.5) \quad a^{(\ell,m)}(q_k^L) \geq \frac{1}{2}(1 - \rho_k) \lambda_k \left(\tau^{(\ell,m)}(q_k^L) \right)^2,$$

and

$$(A.6) \quad \sum_{m=1}^L \tau^{(\ell,m)}(q_k^L) \geq M^{(\ell)},$$

where

$$M^{(\ell)} := \max \left\{ \tau_L - L \left(\frac{\alpha_k^{(\ell)}}{\lambda_k} + \frac{\alpha_k^{(\ell)}}{\mu_k - \lambda_k} \right), 0 \right\}.$$

Adding the inequalities in (A.5) and (A.6) to the constraints of problem (A.4) does not change its feasible region, and yields the following equivalent formulation

$$(A.7) \quad \min_{q^L \in \mathcal{Q}^L} \frac{1}{\tau_L} \sum_{\ell=1}^{N_k} h_k^{(\ell)} \sum_{m=1}^L a^{(\ell,m)}(q_k^L)$$

$$s.t. \quad a^{(\ell,m)}(q_k^L) = \int_{u^{(m-1)}}^{u^{(m)}} q_k^L(s) \mathbf{1}_{\{q_k^L(s) \geq \alpha_k^{(\ell)}\}} ds,$$

$$\tau^{(\ell,m)}(q_k^L) = \int_{u^{(m-1)}}^{u^{(m)}} \mathbf{1}_{\{q_k^L(s) \geq \alpha_k^{(\ell)}\}} ds,$$

$$a^{(\ell,m)}(q_k^L) \geq \frac{1}{2}(1 - \rho_k) \lambda_k \left(\tau^{(\ell,m)}(q_k^L) \right)^2, \quad 1 \leq m \leq L, \quad 1 \leq \ell \leq N_k$$

$$\sum_{m=1}^L \tau^{(\ell,m)}(q_k^L) \geq M^{(\ell)}, \quad 1 \leq \ell \leq N_k.$$

Next, let

$$\boldsymbol{\tau} := (\tau^{(\ell,m)}, 1 \leq m \leq L, 1 \leq \ell \leq N_k) \quad \text{and} \quad \mathbf{a} := (a^{(\ell,m)}, 1 \leq m \leq L, 1 \leq \ell \leq N_k).$$

We consider the following relaxed problem by dropping the first three constraints in problem (A.7)

$$(A.8) \quad \begin{aligned} \min_{\boldsymbol{\tau}, \mathbf{a}} \quad & \frac{1}{\tau_L} \sum_{\ell=1}^{N_k} h_k^{(\ell)} \sum_{m=1}^L a^{(\ell,m)} \\ \text{s.t.} \quad & a^{(\ell,m)} \geq \frac{1}{2} (1 - \rho_k) \lambda_k \left(\tau^{(\ell,m)} \right)^2, \quad 1 \leq m \leq L, \quad 1 \leq \ell \leq N_k \\ & \sum_{m=1}^L \tau^{(\ell,m)} \geq M^{(\ell)}, \quad 1 \leq \ell \leq N_k. \end{aligned}$$

It follows from observation (and can be verified by solving the Karush-Kuhn Tucker equations) that the solution to problem (A.8), denoted by $(\boldsymbol{\tau}_*, \mathbf{a}_*)$, has elements

$$\tau_*^{(\ell,m)} = M^{(\ell)} / L \quad \text{and} \quad a_*^{(\ell,m)} = \frac{1}{2} (1 - \rho_k) \lambda_k \left(\tau_*^{(\ell,m)} \right)^2, \quad 1 \leq m \leq L, \quad 1 \leq \ell \leq N_k.$$

Note that problem (A.8) is a relaxation of problem (A.7) because for any feasible solution $(\boldsymbol{\tau}, \mathbf{a})$ to (A.8), there does not necessarily exist a corresponding $q^L \in \mathcal{Q}_k^L$ such that, for all $1 \leq m \leq L$ and $1 \leq \ell \leq N_k$,

$$(A.9) \quad a^{(\ell,m)} = \int_{u^{(m-1)}}^{u^{(m)}} q_k^L(s) \mathbf{1}_{\{q_k^L(s) \geq \alpha_k^{(\ell)}\}} ds \quad \text{and} \quad \tau^{(\ell,m)} = \int_{u^{(m-1)}}^{u^{(m)}} \mathbf{1}_{\{q_k^L(s) \geq \alpha_k^{(\ell)}\}} ds.$$

Hence, if there exists a $q_e^L \in \mathcal{Q}_k^L$ such that (A.9) holds for $(\boldsymbol{\tau}_*, \mathbf{a}_*)$, then q_e^L is a solution to (A.7).

Let $q_{exh,k}$ denote the trajectory of queue k in q_{exh} . (Note that $q_{exh} \in \mathcal{Q}_k^L$.) It can be verified that (A.9) indeed holds for $(q_{exh,k}, \boldsymbol{\tau}_*, \mathbf{a}_*)$, namely, for all $1 \leq m \leq L$ and $1 \leq \ell \leq N_k$,

$$a_*^{(\ell,m)} = \int_{u^{(m-1)}}^{u^{(m)}} q_{exh,k}(s) \mathbf{1}_{\{q_{exh,k}(s) \geq \alpha_k^{(\ell)}\}} ds \quad \text{and} \quad \tau_*^{(\ell,m)} = \int_{u^{(m-1)}}^{u^{(m)}} \mathbf{1}_{\{q_{exh,k}(s) \geq \alpha_k^{(\ell)}\}} ds.$$

Thus, q_{exh} solves (A.7), and therefore also the equivalent problem (A.2). In turn, q_{exh} is a solution to the L -cycle optimization problem (4.5). Since the arguments hold for each $L \in \mathbb{N}$, q_{exh} is optimal to (4.4). \square

REFERENCES

- Altman, E., A. Khamisy, U. Yechiali. 1992. On elevator polling with globally gated regime. *Queueing Systems* **11**(1-2) 85–90.
- Atar, R., C. Giat, N. Shimkin. 2011. On the asymptotic optimality of the $c\mu/\theta$ rule under ergodic cost. *Queueing Systems* **67**(2) 127–144.
- Billingsley, P. 1968. *Convergence of probability measures*. John Wiley & Sons.
- Boon, M., O. J. Boxma, E. M. Winands. 2011a. On open problems in polling systems. *Queueing Systems* **68**(3-4) 365.
- Boon, M. A., R. Van der Mei, E. M. Winands. 2011b. Applications of polling systems. *Surveys in Operations Research and Management Science* **16**(2) 67–82.
- Borst, S., O. J. Boxma. 2018. Polling: past, present, and perspective. *Top* **26**(3) 335–369.
- Boxma, O. J. 1986. Models of two queues: a few new views. In *Teletraffic Analysis and Computer Performance Evaluation*, O.J. Boxma, J.W. Cohen and H.C. Tijms, Eds. Elsevier Science Publishers, Amsterdam, 75–98.
- Bux, W. 1981. Local-area subnetworks: A performance comparison. *IEEE Transactions on Communications* **29**(10) 1465–1473.

- Chase, C., J. Serrano, P. J. Ramadge. 1993. Periodicity and chaos from switched flow systems: contrasting examples of discretely controlled continuous systems. *IEEE Transactions on Automatic Control* **38**(1) 70–83.
- Cicin-Sain, M., C. E. Pearce, J. Sunde. 2001. On the application of a polling model with non-zero walk times and priority processing to a medical emergency-room environment. *Proceedings of the 23rd International Conference on Information Technology Interfaces, 2001. ITI 2001*. IEEE, 49–56.
- Coffman Jr, E., A. Puhalskii, M. Reiman. 1995. Polling systems with zero switchover times: a heavy-traffic averaging principle. *The Annals of Applied Probability* 681–719.
- Coffman Jr, E., A. Puhalskii, M. Reiman. 1998. Polling systems in heavy traffic: a Bessel process limit. *Mathematics of Operations Research* **23**(2) 257–304.
- Duenyas, I., M. P. Van Oyen. 1996. Heuristic scheduling of parallel heterogeneous queues with set-ups. *Management Science* **42**(6) 814–829.
- Federgruen, A., Z. Katalan. 1994. Approximating queue size and waiting time distributions in general polling systems. *Queueing Systems* **18**(3) 353–386.
- Federgruen, A., Z. Katalan. 1996. The stochastic economic lot scheduling problem: cyclical base-stock policies with idle times. *Management Science* **42**(6) 783–796.
- Fricker, C., M. R. Jaibi. 1994. Monotonicity and stability of periodic polling models. *Queueing systems* **15**(1-4) 211–238.
- Harrison, J. M. 1988. Brownian models of queueing networks with heterogeneous customer populations. *Stochastic differential systems, stochastic control theory and applications*. Springer, 147–186.
- Harrison, J. M., J. A. Van Mieghem, et al. 1997. Dynamic control of Brownian networks: state space collapse and equivalent workload formulations. *The Annals of Applied Probability* **7**(3) 747–771.
- Hofri, M., K. W. Ross. 1987. On the optimal control of two queues with server setup times and its analysis. *SIAM Journal on computing* **16**(2) 399–420.
- Hu, Y., J. Dong, O. Perry. 2020. On the existence and asymptotic approximation of moments for polling systems under the binomial-exhaustive policy. Available in *arXiv:XXX*.
- Levy, H. 1988. Optimization of polling systems: the fractional exhaustive service method. *Report, Tel-Aviv University*.
- Levy, H. 1989. Analysis of cyclic polling systems with binomial gated service. *Performance of Distributed and Parallel Systems* 127–139.
- Levy, H., M. Sidi. 1990. Polling systems: applications, modeling, and optimization. *IEEE Transactions on communications* **38**(10) 1750–1760.
- Levy, H., M. Sidi, O. J. Boxma. 1990. Dominance relations in polling systems. *Queueing systems* **6**(1) 155–171.
- Liu, Z., P. Nain, D. Towsley. 1992. On optimal polling policies. *Queueing Systems* **11**(1-2) 59–83.
- Mack, C., T. Murphy, N. Webb. 1957. The efficiency of n machines uni-directionally patrolled by one operative when walking time and repair times are constants. *Journal of the Royal Statistical Society: Series B (Methodological)* **19**(1) 166–172.
- Matveev, A., V. Feoktistova, K. Bolshakova. 2016. On global near optimality of special periodic protocols for fluid polling systems with setups. *Journal of Optimization Theory and Applications* **171**(3) 1055–1070.
- Meyer, C. D. 2000. *Matrix analysis and applied linear algebra*, vol. 71. Siam.
- Nahmias, S., M. H. Rothkopf. 1984. Stochastic models of internal mail delivery systems. *Management science* **30**(9) 1113–1120.
- Olsen, T. L. 1999. A practical scheduling method for multiclass production systems with setups. *Management Science* **45**(1) 116–130.
- Olsen, T. L. 2001. Limit theorems for polling models with increasing setups. *Probability in the Engineering and Informational Sciences* **15**(01) 35–55.
- Perry, O., W. Whitt. 2016. Chattering and congestion collapse in an overload switching control. *Stochastic Systems* **6**(1) 132–210.
- Resing, J. 1993. Polling systems and multitype branching processes. *Queueing Systems* **13**(4) 409–426.
- Szpankowski, W., V. J. Rego. 1987. Ultimate stability conditions for some multidimensional distributed systems.
- Takagi, H. 1988. Queueing analysis of polling models. *ACM Computing Surveys (CSUR)* **20**(1) 5–28.
- Takagi, H. 1997. Queueing analysis of polling models: progress in 1990-1994. *Frontiers in Queueing: Models, Methods and Problems* 119–146.
- Tijms, H. C. 2003. *A first course in stochastic models*. John Wiley and sons.
- Van den Broek, M. S., J. Van Leeuwen, I. J. Adan, O. J. Boxma. 2006. Bounds and approximations for the fixed-cycle traffic-light queue. *Transportation Science* **40**(4) 484–496.
- Van der Mei, R. 1999. Delay in polling systems with large switch-over times. *Journal of applied probability* **36**(01) 232–243.
- Van der Mei, R., H. Levy. 1997. Polling systems in heavy traffic: Exhaustiveness of service policies. *Queueing Systems* **27**(3) 227–250.

- Vishnevskii, V. M., O. V. Semenova. 2006. Mathematical methods to study the polling systems. *Automation and Remote Control* **67**(2) 173–220.
- Vlasiou, M., I. J. Adan, O. J. Boxma. 2009. A two-station queue with dependent preparation and service times. *European Journal of Operational Research* **195**(1) 104–116.
- Whitt, W. 2002. *Stochastic-process limits: an introduction to stochastic-process limits and their application to queues*. Springer Science & Business Media.
- Winands, E. M. 2007. On polling systems with large setups. *Operations research letters* **35**(5) 584–590.
- Winands, E. M. 2011. Branching-type polling systems with large setups. *OR Spectrum* **33**(1) 77–97.
- Winands, E. M., I. J. Adan, G.-J. van Houtum. 2011. The stochastic economic lot scheduling problem: A survey. *European Journal of Operational Research* **210**(1) 1–9.