the service of the i customers arriving at the times $\tau_0^{M_n}, \dots, \tau_{i-1}^{M_n}$ starts without delay, and immediately before time $\tau_i^{M_n}$ all customers that arrived before time $\tau_0^{M_n}$ have left. Thus immediately before time $\tau_i^{M_n}$ there are [due to $\mathbf{P}(\beta_1 < \alpha_1 + \cdots + \alpha_i) = 0$] i customers in the system, they all arrived at or after time $\tau_0^{M_n}$, and their service started without delay. Thus the future of the system after time $\tau_i^{M_n}$ is not affected by its past before time $\tau_0^{M_n}$ and behaves distributionally in the same way for all $n \geqslant 0$. Now note that due to $\mathbf{P}(\alpha_1 + \cdots + \alpha_i \leqslant l) = 1$, we have $\tau_i^{M_n} \leqslant \tau_0^{M_n} + l$. Thus the future of the system after time $S_n = \tau_0^{M_n} + l$ is not affected by its past before time $S_n - l = \tau_0^{M_n}$ and behaves distributionally in the same way for all $n \geqslant 0$. Thus $(Q_s, R_s)_{s \in [0,\infty)}$ is lag-l regenerative with regeneration times S_n, $n \geqslant 0$.

5 Time-Inhomogeneous Regeneration

In this section we extend the classical regeneration concept by allowing the regeneration to depend on the time when it occurs. This is the kind of regeneration found in *time-inhomogeneous* Markov chains (Markov chains with transition probabilities that depend on the time of transition). When such a chain visits a recurrent reference state, it only starts anew *conditionally on the time of visit*, that is, *time-inhomogeneous* regeneration takes place. We shall also extend wide-sense regeneration in the same way by allowing the future, given the time of regeneration, to be conditionally independent not necessarily of the full past but only of the past regeneration times.

Time-inhomogeneity allows the environment in which the process develops to change deterministically with time, as is often the case in real world situations. For instance, the traffic intensity in actual queueing systems can vary drastically with time of day. Adapting a periodic model with period the day is not very helpful, since the relevant time scale is often minutes rather than days. A time-inhomogeneous model thus seems more appropriate. In spite of this, the mathematical theory of time-inhomogeneous models is poorly developed. Hopefully, this section and the next three (though abstract) will be a small contribution to such a theory.

Note that if two independent versions of a time-inhomogeneous Markov chain enter a fixed reference state at the same time, then we can let the two chains run together from that time on without affecting their distributions, that is, we would have created an exact coupling. Note also that it is not natural to look for shift-coupling or epsilon-couplings of time-inhomogeneous Markov chains because the behaviour after regeneration can differ drastically depending on the time of regeneration.

The same observations apply to time-inhomogeneous regenerative processes. Therefore, our main task here is to find conditions under which two versions of such a process can be forced to regenerate simultaneously (or simultaneously in the distributional sense). We shall carry out the main part of the construction of a successful (distributional) exact coupling in the next section, but the resulting analogue of Theorems 3.3 and 4.3 is stated in this section as Theorem 5.3.

5.1 Definitions

Call a one-sided shift-measurable stochastic process $Z = (Z_s)_{s\in[0,\infty)}$ *time-inhomogeneous regenerative* with *regeneration times* S if the future after regeneration $\theta_{S_n}(Z,S)$ depends on the past $((Z_s)_{s\in[0,S_n)}, S_0, \dots, S_n)$ only through the time of regeneration S_n, and the conditional distribution of $\theta_{S_n}(Z,S)$ given the value of S_n is regular and does not depend on $n \geqslant 0$. In other words, Z is time-inhomogeneous regenerative with regeneration times S if there is a $(([0,\infty), \mathcal{B}[0,\infty)), (H \times L, \mathcal{H} \otimes \mathcal{L}))$ probability kernel $p(\cdot|\cdot)$ such that for $n \geqslant 0$ and $A \in \mathcal{H} \otimes \mathcal{L}$,

$$\mathbf{P}(\theta_{S_n}(Z,S) \in A | (Z_s)_{s\in[0,S_n)}, S_0, \dots, S_n) = p(A|S_n) \quad \text{a.s.} \tag{5.1}$$

Call the pair (Z,S) *time-inhomogeneous regenerative* of *type* $p(\cdot|\cdot)$ if this holds. Let the negative random variable S_{-1} be such that (Z°, S°) depends on (D, S_{-1}) only through S_0.

Call a one-sided shift-measurable process Z *time-inhomogeneous wide-sense regenerative* with *regeneration times* S if the future after regeneration $\theta_{S_n}(Z,S)$ depends on the past regeneration times $(S_0, \dots, S_n)$ only through the time of regeneration S_n, and the conditional distribution of $\theta_{S_n}(Z,S)$ given the value of S_n is regular and does not depend on $n \geqslant 0$. In other words, Z is time-inhomogeneous wide-sense regenerative with regeneration times S if there is a $(([0,\infty), \mathcal{B}[0,\infty)), (H \times L, \mathcal{H} \otimes \mathcal{L}))$ probability kernel $p(\cdot|\cdot)$ such that for $n \geqslant 0$ and $A \in \mathcal{H} \otimes \mathcal{L}$,

$$\mathbf{P}(\theta_{S_n}(Z,S) \in A | S_0, \dots, S_n) = p(A|S_n) \quad \text{a.s.} \tag{5.2}$$

Call the pair (Z,S) *time-inhomogeneous wide-sense regenerative* of *type* $p(\cdot|\cdot)$ if this holds. Let the negative random variable S_{-1} be such that (Z°, S°) depends on S_{-1} only through S_0.

With $l \geqslant 0$, call a time-inhomogeneous wide-sense regenerative (Z,S) *time-inhomogeneous lag-l regenerative* if (5.2) can be strengthened to: for $n \geqslant 0$ and $A \in \mathcal{H} \otimes \mathcal{L}$,

$$\mathbf{P}(\theta_{S_n}(Z,S) \in A | (Z_s)_{s\in[0,(S_n-r)^+]}, S_0, \dots, S_n) = p(A|S_n) \ \text{ a.s.;} \tag{5.3}$$

and *time-inhomogeneous lag-$l+$ regenerative* if (5.2) can only be strengthened to: for $n \geqslant 0$ and $A \in \mathcal{H} \otimes \mathcal{L}$,

$$\mathbf{P}(\theta_{S_n}(Z,S) \in A | (Z_s)_{s\in[0,(S_n-l)^+)}, S_0, \dots, S_n) = p(A|S_n) \quad \text{a.s.}$$

Therefore, time-inhomogeneous lag-0+ regeneration is the same as time-inhomogeneous regeneration, while lag-0 regeneration implies further that Z_{S_n} is a measurable mapping of S_n.

A pair (Z', S') is a *version* of a time-inhomogeneous regenerative (Z, S) if (Z', S') is time-inhomogeneous regenerative of the same type as (Z, S). A pair (Z', S') is a version of a time-inhomogeneous wide-sense (or lag-l, or lag-$l+$) regenerative (Z, S) if (Z', S') is time-inhomogeneous wide-sense (or lag-l, or lag-$l+$) regenerative of the same type as (Z, S). Note that in these cases the zero-delayed (Z°, S°) is in general *not* a version of (Z, S).

Call a wide-sense time-inhomogeneous regenerative (Z, S) of type $p(\cdot|\cdot)$ *time-homogeneous* if $p(\cdot|s)$ does not depend on s, that is, if

$$p(\cdot|s) = p(\cdot) := \mathbf{P}((Z^\circ, S^\circ) \in \cdot), \quad s \in [0, \infty).$$

Thus if a time-inhomogeneous regenerative (Z, S) is time-homogeneous, then it is classical regenerative. And if a time-inhomogeneous wide-sense (lag-l, lag-$l+$) regenerative (Z, S) is time-homogeneous, then it is wide-sense (lag-l, lag-$l+$) regenerative.

5.2 The Regeneration Times S Are Time-*Homogeneous* Markov

If (Z, S) is time-inhomogeneous regenerative (wide-sense or not), then in general it is neither true that the cycles are i.i.d. nor that S forms a renewal process. However, the following holds.

Theorem 5.1. *If (Z, S) is time-inhomogeneous wide-sense regenerative, then the sequence S is a time-homogeneous Markov process.*

PROOF. According to (5.2), for each $n \geqslant 0$, $(S_0, \dots, S_n)$ depends only on $(S_{n+1} - S_n, S_{n+2} - S_n, \dots)$ through S_n and thus only on $(S_{n+1}, S_{n+2}, \dots)$ through S_n, that is, S is a Markov process. Also, according to (5.2), the conditional distribution of $(S_{n+1} - S_n, S_{n+2} - S_n, \dots)$ given the value of S_n does not depend on n, and therefore the conditional distribution of $(S_{n+1}, S_{n+2}, \dots)$ given the value of S_n does not depend on n, that is, S is time-homogeneous. □

Let F_s be the conditional distribution of X_{k+1} given $S_k = s$, that is, for $s \in [0, \infty)$ and $A \in \mathcal{B}[0, \infty)$,

$$F_s(A) := p(H \times (A \times [0, \infty)^\infty) \cap L|s) = \mathbf{P}(X_{k+1} \in A | S_k = s).$$

We shall view S as a 'renewal process' that is time-inhomogeneous in the sense that if a 'renewal' occurs at time $S_k = s$, then the next recurrence time X_{k+1} is governed by a distribution that may depend on s, namely F_s. Call F_s the *recurrence distribution at s* and define, for $1 \leqslant n < \infty$, the *n-step recurrence distribution* at s by

$$\begin{aligned} F_s^n(A) &:= p(H \times ([0, \infty)^{n-1} \times A \times [0, \infty)^\infty) \cap L|s) \\ &= \mathbf{P}(S_{k+n} - S_k \in A | S_k = s), \quad s \in [0, \infty), A \in \mathcal{B}[0, \infty). \end{aligned}$$

Let F_s and F_s^n also denote the conditional distribution functions of X_{k+1} and $X_{k+1} + \cdots + X_{k+n}$, respectively, given $S_n = s$, that is, for $s \in [0, \infty)$ and $x \in [0, \infty)$,

$$F_s(x) := F_s([0, x]) = \mathbf{P}(X_{k+1} \leqslant x | S_k = s),$$
$$F_s^n(x) := F_s^n([0, x]) = \mathbf{P}(X_{k+1} + \cdots + X_{k+n} \leqslant x | S_k = s).$$

The n-step transition probabilities of S are

$$\mathbf{P}(S_{k+n} \in A | S_k = s) = F_s^n(s + A), \quad s \in [0, \infty), \ \ A \in \mathcal{B}[0, \infty).$$

If (Z, S) is time-homogeneous, then S is a renewal process, and we have $F_s = F$ independently of s, where F is the common distribution of the i.i.d. recurrence times, and $F_s^n = F^n$, where F^n is the distribution of the sum $X_1 + \cdots + X_n$ (F^n is the nth convolution power of F).

5.3 Examples

Let $S = (S_k)_0^\infty$ be a discrete-time Markov process with state space $[0, \infty)$, increasing strictly to infinity. Put $N_{s-} = \lim_{t \uparrow s} N_t$. Then $(S_{N_{s-}})_{s \in [0,\infty)}$ is time-inhomogeneous lag-0 regenerative with regeneration times S, but certainly $(S_{N_{s-}})_{s \in [0,\infty)}$ is not time-homogeneous, not even when S is a renewal process. Also, $(S_{N_{s-}}, A_s, B_s, D_s, U_s)_{s \in [0,\infty)}$ is time-inhomogeneous lag-0+ (but not lag-0) regenerative with regeneration times S. Moreover, for any $l > 0$, the process $(S_{N_{s-}}, A_s, B_s, D_s, U_s, N_{s+l} - N_s)_{s \in [0,\infty)}$ is time-inhomogeneous lag-l regenerative. If Z is wide-sense regenerative (or lag-l regenerative) with regeneration times S, then so is the stochastic process $(Z_s, S_{N_{s-}}, A_s, B_s, D_s, U_s, N_{s+l} - N_s)_{s \in [0,\infty)}$.

Consider a discrete-time Markov process $Y = (Y_k)_0^\infty$ with state space $\mathbb{R}$ and with $\lim_{k \to \infty} Y_k = \infty$. Put $S_n = Y_{K_n}$, where $K_0 = \inf\{k \geqslant 0 : Y_k \geqslant 0\}$ and, recursively for $n \geqslant 1$, $K_n = \inf\{k \geqslant 0 : Y_k > S_{n-1}\}$. Take $l > 0$ and let $Z = (Z_s)_{s \in [0,\infty)}$ be the process with Z_s the number of times the Markov process Y visits the interval $(s, s + l]$. Then (Z, S) is time-inhomogeneous wide-sense regenerative, but in general not lag-l.

A time-inhomogeneous Markov chain $Z = (Z_s)_{s \in [0,\infty)}$ with a recurrent state is time-inhomogeneous regenerative with regeneration times S formed by the successive entrances to this state. More generally, let Z be a time-inhomogeneous continuous-time general state space shift-measurable Markov process, let A be a recurrent set and let the time-homogeneous *space-time process* $(Z_s, s)_{s \in [0,\infty)}$ be strong Markov. If the transition probabilities of Z are the same from all states in A, then Z is time-inhomogeneous regenerative with regeneration times S formed by the successive entrances to A.

Still more generally, Theorem 4.6 has a time-inhomogeneous counterpart. Let $Z = (Z_s)_{s \in [0,\infty)}$ be a time-inhomogeneous general state space shift-measurable Markov process such that the space-time process $(Z_s, s)_{s \in [0,\infty)}$

is strong Markov. Suppose Z has a set of states A such that τ_A (the hitting time of A) is measurable and finite with probability one for all initial distributions and $Z_{\tau_A} \in A$, and such that for some $l > 0$, $p \in (0,1]$ and a probability kernel $\mu(\cdot,\cdot)$,

$$\mathbf{P}(Z_{t+l} \in \cdot | Z_t = x) \geqslant p\mu(t,\cdot), \quad x \in A, \ t \in [0,\infty).$$

Then, with $T_0 = \tau_A$ and $T_{k+1} = \inf\{t \geqslant T_k + l : Z_t \in A\}$ for $k \geqslant 0$, we have

$$\mathbf{P}(Z_{T_k+l} \in \cdot | Z_{T_k} = x, T_k = t) \geqslant p\mu(t,\cdot), \quad x \in A, \ t \in [0,\infty).$$

This allows us to extend the underlying probability space by conditional splitting: apply Theorem 5.1 in Chapter 3 recursively to obtain i.i.d. 0-1 variables $I_0, I_1, \ldots$ such that for $k \geqslant 0$,

$$(Z, I_0, \ldots, I_{k-1}) \text{ depends on } I_k \text{ only through } (Z_{T_k}, T_k, Z_{T_k+l}),$$
$$\mathbf{P}(I_k = 1 | Z_{T_k}, T_k) = p \quad \text{and} \quad \mathbf{P}(Z_{T_k+l} \in \cdot | Z_{T_k}, T_k, I_k) = \mu(T_k, \cdot).$$

Let K_n be the $(n+1)$th index k such that $I_k = 1$. Conditionally on the randomized stopping time $S_n = T_{K_n} + l$, the time-inhomogeneous Markov process Z has conditional distribution $\mu(S_n, \cdot)$ at time S_n and is conditionally independent of its state at time $S_n - l$. Thus, conditionally on S_n, the future of Z after time S_n is independent of the past before time $S_n - l$, and the conditional distribution of the future given the value of S_n does not depend on n. This argument can be sharpened along the lines of the proof of Theorem 4.6. Thus Z is time-inhomogeneous lag-l regenerative.

Finally, consider the following time-inhomogeneous version of the $GI/GI/k$ queueing model, $1 \leqslant k \leqslant \infty$: customers arrive to a k-server station at times forming a Markov sequence with state space $[0,\infty)$ and increasing strictly to infinity, and line up to be served under the first-come-first-served discipline with service times that depend on the time of arrival and/or the time when the service starts. A simple special case is the time-inhomogeneous version of the $M/GI/k$ queue obtained by allowing the Poisson arrivals (M stands for *memoryless*, the Poisson process property) to be nonstationary. If the system empties infinitely often with probability one, then the queue length and ordered remaining service times process $(Q_s, R_s)_{s\in[0,\infty)}$ is time-inhomogeneous regenerative with the times of arrivals to an idle system as regeneration times. If the system cannot empty but $(Q_s)_{s\in[0,\infty)}$ has a recurrent state $i < k$ and the event at (4.30) occurs for infinitely many m, then the argument in Section 4.6 shows that $(Q_s, R_s)_{s\in[0,\infty)}$ is time-inhomogeneous lag-l regenerative.

5.4 The Key Coupling Result

The key coupling result from the time-homogeneous case (Theorems 3.2 and 4.2) extends as follows to the time-inhomogeneous case. We leave out