

When is Eaton's Markov chain irreducible?

James P. Hobert, Aixin Tan and Ruitao Liu

Department of Statistics

University of Florida

Gainesville, FL 32611 USA

November 2006

Abstract

Consider a parametric statistical model, $P(dx|\theta)$, and an improper prior distribution, $\nu(d\theta)$, that together yield a (proper) formal posterior distribution, $Q(d\theta|x)$. The prior is called *strongly admissible* if the generalized Bayes estimator of every bounded function of θ is admissible under squared error loss. ? showed that a sufficient condition for strong admissibility of ν is the *local recurrence* of the Markov chain whose transition function is $R(\theta, d\eta) = \int Q(d\eta|x)P(dx|\theta)$. Applications of this result and its extensions are often greatly simplified when the Markov chain associated with R is irreducible. However, establishing irreducibility can be difficult. In this paper, we provide a characterization of irreducibility for general state space Markov chains, and we use this characterization to develop an easily-checked, necessary and sufficient condition for irreducibility of Eaton's Markov chain. All that is required to check this condition is a simple examination of P and ν . Application of the main result is illustrated using two examples.

1 Introduction

Consider a parametric statistical decision problem with sample space X and parameter space Θ . Both X and Θ are assumed to be Polish spaces equipped with their Borel σ -algebras $\mathcal{B}(X)$ and $\mathcal{B}(\Theta)$. Suppose that $P : \mathcal{B}(X) \times \Theta \rightarrow [0, 1]$ represents a parametric statistical model; i.e., for each θ , $P(\cdot|\theta)$ is a probability measure and, for each A , $P(A|\cdot)$ is a measurable function. As usual, the idea is that we will observe a random element whose distribution is $P(dx|\theta)$ and the goal is to use the observation to make inferences about the unknown parameter θ . This will be done within the Bayesian paradigm using an *improper* prior distribution. In particular, let $\nu(d\theta)$ denote a σ -finite measure with $\nu(\Theta) = \infty$. Define the marginal measure as

$$M(dx) = \int_{\Theta} P(dx|\theta)\nu(d\theta) .$$

? shows that if M is σ -finite, then there exists a formal posterior distribution, Q , defined as follows.

Definition 1. A function $Q : \mathcal{B}(\Theta) \times X \rightarrow [0, 1]$ is called a formal posterior distribution (FPD) if

1. $Q(\cdot|x)$ is a probability measure for each x ,
2. $Q(B|\cdot)$ is a measurable function for each B , and
3. $Q(d\theta|x)M(dx) = P(dx|\theta)\nu(d\theta)$; that is, for all $A \in \mathcal{B}(X)$ and $B \in \mathcal{B}(\Theta)$,

$$\int_A Q(B|x)M(dx) = \int_B P(A|\theta)\nu(d\theta) .$$

The FPD is unique in the sense that if \tilde{Q} is another FPD, then there is an M -null set A_0 such that $x \notin A_0$ implies $Q(d\theta|x) = \tilde{Q}(d\theta|x)$. Throughout this paper, M is assumed to be σ -finite so an FPD is guaranteed to exist. We now briefly describe a method of evaluating improper prior distributions that is due to M.L. Eaton. (For a more in depth review of this area, see ?.)

Consider the problem of estimating a bounded, real-valued function $\gamma(\theta)$ under squared error loss. Of course, the formal Bayes estimator of $\gamma(\theta)$ is $\hat{\gamma}(x) = \int_{\Theta} \gamma(\theta)Q(d\theta|x)$. The risk function of a generic estimator, say δ , is its mean squared error; i.e.,

$$r(\delta, \theta) = \int_X (\delta(x) - \gamma(\theta))^2 P(dx|\theta) .$$

The estimator δ is called *almost- ν -admissible* if for any estimator δ' such that

$$r(\delta', \theta) \leq r(\delta, \theta) \quad \forall \theta \in \Theta ,$$

the set $\{\theta \in \Theta : r(\delta', \theta) < r(\delta, \theta)\}$ has ν -measure zero. If $P(dx|\theta)$ and $\nu(d\theta)$ combine to yield an FPD that generates (almost) admissible estimators for a large class of functions of θ , then we would be willing to endorse ν as a good “all purpose” prior to use in conjunction with the statistical model $P(dx|\theta)$. This idea provides motivation for the following definition.

Definition 2. The prior ν is called *strongly admissible* if $\hat{\gamma}$ is almost- ν -admissible for every bounded, real-valued function γ .

? developed a sufficient condition for strong admissibility that involves the Markov transition function $R : \Theta \times \mathcal{B}(\Theta) \rightarrow [0, 1]$ given by

$$R(\theta, d\eta) = \int_{\mathcal{X}} Q(d\eta|x)P(dx|\theta) .$$

Before we can state the result, we need a couple of concepts from general state space Markov chain theory. Let $W = \{W_n\}_{n=0}^{\infty}$ denote the Markov chain on Θ driven by R and let \Pr_{θ} denote the overall probability law governing the chain when $W_0 = \theta$. For $B \in \mathcal{B}(\Theta)$, let σ_B denote the first return to B ; i.e.,

$$\sigma_B = \min \{n \geq 1 : W_n \in B\} ,$$

with the understanding that $\sigma_B = \infty$ if $W_n \in \bar{B}$ for all $n \geq 1$.

Definition 3. The Markov chain W is called *locally- ν -recurrent* if, for each B with $0 < \nu(B) < \infty$, the set

$$\{\theta \in B : \Pr_{\theta}(\sigma_B < \infty) < 1\}$$

has ν -measure 0.

In words, the chain is locally- ν -recurrent if, when started inside the set B , then aside from a set of starting values that has ν -measure 0, the chain returns to B with probability 1. Note that, unlike the standard definition of recurrence (see; e.g., ?, Chapter 8), this definition pertains to both reducible and irreducible chains. Indeed, just before defining local- ν -recurrence on page 1174, ? states: ‘‘The following definition, a modified notion of recurrence, allows us to circumvent a discussion of irreducibility issues while relating our previous admissibility results to the recurrence of W .’’ The following was proven in ?.

Theorem 1. *If W is a locally- ν -recurrent Markov chain, then ν is a strongly admissible prior.*

Establishing local- ν -recurrence directly using the definition (or the characterization based on the Dirichlet form of R) is typically infeasible. However, if W is ν -irreducible; that is, any set B with $\nu(B) > 0$ is accessible from any $\theta \in \Theta$, then recurrence and local- ν -recurrence are equivalent (?). Hence, if ν -irreducibility of W can be demonstrated, then all of the techniques that have been developed for establishing recurrence can be brought to bear on the problem. Indeed, nearly all of the applications of Theorem 1 have involved first demonstrating that W is ν -irreducible and then showing that W is recurrent. Examples can be found in ?, ?, ? and ?. Similarly, ? have recently extended and generalized the theoretical results of ? and ? under the assumption that the chains of interest are irreducible.

There is one very simple sufficient condition for ν -irreducibility of W and this was used in most of the applications mentioned above. If the support of the statistical model does not depend on the

parameter; i.e., if the set $\{A \in \mathcal{B}(X) : P(A|\theta) > 0\}$ is the same for all $\theta \in \Theta$, then W is ν -irreducible (?). Until now, however, there has been no easy way to check for ν -irreducibility of Eaton's Markov chain when this condition fails. In this paper, we provide an easily-checked, necessary and sufficient condition for ν -irreducibility of W . This result cannot be stated precisely at this point, but the sufficiency half, which is the practically important part, can be.

Theorem 2. *The Markov chain W is ν -irreducible if there do not exist two sets $A \in \mathcal{B}(X)$ and $C \in \mathcal{B}(\Theta)$ with the following properties: C is nonempty, $\nu(\overline{C}) > 0$, $P(\overline{A}|\theta) = 0$ for every $\theta \in C$ and $P(A|\theta) = 0$ for ν -almost all $\theta \in \overline{C}$.*

This result allows one to establish ν -irreducibility of Eaton's Markov chain through a simple examination of P and ν . Neither the posterior distribution, Q , nor the Markov transition function, R , is required to check the condition. It is interesting that, if the sets A and C do exist, then $P(dx|\theta)$ and $P(dx|\theta')$ are mutually singular probability measures whenever $\theta \in C$ and $\theta' \in \overline{C}$ (aside from a ν -null set of θ s in \overline{C}). Thus, the statistical model is, in a sense, an artificial concatenation of two different models.

The rest of the paper is organized as follows. Section 2 contains a new characterization of irreducibility for general state space Markov chains. This characterization is used in Section 3 to prove the main result. Section 3 also contains two examples.

2 A characterization of irreducibility for general Markov chains

Let $S(y, dz)$ be a Markov transition function on a general state space $(Y, \mathcal{B}(Y))$ as described, for example, in ?, Section 3.4. Denote the corresponding Markov chain as $Y = \{Y_n\}_{n=0}^\infty$. For $n \in \mathbb{N} := \{1, 2, 3, \dots\}$, let $S^n(y, dz)$ denote the n -step Markov transition function corresponding to S , which is defined inductively by

$$S^{n+1}(y, dz) = \int_Y S^n(w, dz) S(y, dw),$$

where $S^1 \equiv S$. Of course, $S^n(y, A) = \Pr_y(Y_n \in A)$, where $\Pr_y(\cdot)$ denotes the overall law governing Y on Y^∞ assuming that $Y_0 = y$. Let φ denote a non-trivial, σ -finite measure on $(Y, \mathcal{B}(Y))$. Here is a standard definition of irreducibility for general state space Markov chains.

Definition 4. (? , p.87). *The Markov chain Y is called φ -irreducible if, for every measurable A with $\varphi(A) > 0$ and every $y \in Y$, there exists an $n \in \mathbb{N}$ (which may depend on y and A) such that $S^n(y, A) > 0$.*

In words, φ -irreducibility means that every set A with $\nu(A) > 0$ is *accessible* from any $y \in Y$. We will call the Markov chain φ -reducible when it is not φ -irreducible; that is, when there exist y and A with $\varphi(A) > 0$ such that $S^n(y, A) = 0$ for all $n \in \mathbb{N}$. We will sometimes find it convenient to apply the

phrases “ φ -irreducible” and “ φ -reducible” to the Markov transition function, S . Our first result shows that, if the chain is φ -reducible, then we may assume that $y \in \bar{A}$.

Proposition 1. *The Markov chain Y is φ -reducible if and only if there exist $A \in \mathcal{B}(Y)$ with $\varphi(A) > 0$ and $y \in \bar{A}$ such that $S^n(y, A) = 0$ for all $n \in \mathbb{N}$.*

Proof. The sufficiency part is obvious. Now assume that the chain is φ -reducible so there exist $y \in Y$ and $A \in \mathcal{B}(Y)$ with $\varphi(A) > 0$ such that $S^n(y, A) = 0$ for all $n \in \mathbb{N}$. If $y \in \bar{A}$, then there is nothing to prove, so assume that $y \in A$. We will establish the existence of $y' \in \bar{A}$ such that $S^n(y', A) = 0$ for all $n \in \mathbb{N}$. First, for each $m \in \mathbb{N}$, define

$$B_m = \{w \in \bar{A} : S^m(w, A) > 0\},$$

and set $B = \cup_{m=1}^{\infty} B_m$. Now fix $m \in \mathbb{N}$ and note that

$$0 = S^{m+1}(y, A) = \int_Y S^m(z, A)S(y, dz) \geq \int_{B_m} S^m(z, A)S(y, dz).$$

Since $S^m(z, A) > 0$ for $z \in B_m$, we must have $S(y, B_m) = 0$. But this result holds for every $m \in \mathbb{N}$, so it follows that $S(y, B) = 0$. Note that Y can be partitioned into A , B and $\bar{A} - B$, and we know that $S(y, A) = S(y, B) = 0$. Therefore, $S(y, \bar{A} - B) = 1$, which implies that $\bar{A} - B$ is not empty. Clearly, any $y' \in \bar{A} - B$ satisfies $S^n(y', A) = 0$ for all $n \in \mathbb{N}$. \square

In the classical case where Y is countable, the chain is called *irreducible* (with no prefix) if, for each $i, j \in Y$, there is an $n \in \mathbb{N}$ such that $S^n(i, \{j\}) > 0$. This is equivalent to c -irreducibility where c denotes counting measure on Y . In this context, a nonempty set $C \subset Y$ is called *closed* if, once the chain enters C , it cannot leave. Formally, C is closed if $\sum_{j \in C} S(i, \{j\}) = 1$ for all $i \in C$. Obviously, the state space, Y , is closed. In fact, the Markov chain is irreducible if and only if Y has no proper, closed subset (see; e.g., ?, Problem 8.21). We now extend these ideas to handle Markov chains on general state spaces.

Definition 5. *A set $C \in \mathcal{B}(Y)$ is called closed if it is nonempty and $S(y, \bar{C}) = 0$ for all $y \in C$.*

Here is the general state space version of the result in ?’s (?) Problem 8.21.

Theorem 3. *The Markov chain Y is φ -reducible if and only if there exists a closed set C with $\varphi(\bar{C}) > 0$.*

Proof. To prove sufficiency, suppose C is a closed set with $\varphi(\bar{C}) > 0$. Assume that for some $n \in \mathbb{N}$, $S^n(y, \bar{C}) = 0$ for all $y \in C$. Then, for any $y \in C$, we have

$$S^{n+1}(y, \bar{C}) = \int_Y S^n(z, \bar{C})S(y, dz) = \int_C S^n(z, \bar{C})S(y, dz) = 0.$$

Hence, by induction, $S^n(y, \overline{C}) = 0$ for all $y \in C$ and all $n \in \mathbb{N}$. Therefore, since C is nonempty and $\varphi(\overline{C}) > 0$, the chain is φ -reducible.

Now to prove necessity, assume that the chain is φ -reducible. By Proposition 1, there exist a measurable A with $\varphi(A) > 0$ and a $y \in \overline{A}$ such that $S^n(y, A) = 0$ for all $n \in \mathbb{N}$. For each $m \in \mathbb{N}$, define

$$B_m = \{w \in Y : S^m(w, A) > 0\},$$

and set $B = \cup_{m=1}^{\infty} B_m$. We will show that the measurable set $C := \overline{A} \cap \overline{B}$ is the closed set that we seek. First, $y \in C$ so C is nonempty. Now suppose $y' \in C$ is such that $S(y', \overline{C}) > 0$. Then

$$S(y', A) + S(y', B) \geq S(y', A \cup B) = S(y', \overline{C}) > 0.$$

Since $y' \in C \Rightarrow y' \in \overline{B} \Rightarrow y' \in \overline{B}_1$, we know $S(y', A) = 0$, so it must be the case that $S(y', B) > 0$. Hence, there must exist a $k \in \mathbb{N}$ such that $S(y', B_k) > 0$, which implies that

$$S^{k+1}(y', A) = \int_Y S^k(z, A) S(y', dz) = \int_{B_k} S^k(z, A) S(y', dz) > 0.$$

But this implies that $y' \in B_{k+1}$, which contradicts the fact that $y' \in C$. Therefore, it must be the case that $S(y, \overline{C}) = 0$ for all $y \in C$, and this implies that C is closed. Finally, $\varphi(\overline{C}) \geq \varphi(A) > 0$. \square

In the next section, we use Theorem 3 to develop an easily-checked, necessary and sufficient condition for ν -irreducibility of Eaton's Markov chain.

3 Conditions for irreducibility of Eaton's Markov chain

? studied the Markov transition function

$$R(\theta, d\eta) = \int_{\mathcal{X}} Q(d\eta|x) P(dx|\theta),$$

where P is the statistical model and Q is a FPD. Since any FPD can be used to construct R , R is not unique. For example, if $\tilde{Q} \neq Q$ is another FPD, then

$$\tilde{R}(\theta, d\eta) = \int_{\mathcal{X}} \tilde{Q}(d\eta|x) P(dx|\theta)$$

is an equally valid version of the Markov transition function. However, the following result shows that R enjoys a uniqueness property similar to the uniqueness property of the FPD that was discussed in Section 1.

Proposition 2. *If R and \tilde{R} are two different versions of Eaton's Markov transition function, then there exists a ν -null set $B_0 \in \mathcal{B}(\Theta)$ such that $\theta \notin B_0$ implies $R(\theta, \cdot) = \tilde{R}(\theta, \cdot)$.*

Proof. Since Q and \tilde{Q} are both FPDs, there exists an M -null set A_0 such that $x \notin A_0 \Rightarrow Q(d\theta|x) = \tilde{Q}(d\theta|x)$. Fix $B \in \mathcal{B}(\Theta)$ and note that

$$R(\theta, B) - \tilde{R}(\theta, B) = \int_{A_0} [Q(B|x) - \tilde{Q}(B|x)] P(dx|\theta) .$$

Now since $0 = M(A_0) = \int_{\Theta} P(A_0|\theta)\nu(d\theta)$, there exists a ν -null set B_0 such that $P(A_0|\theta) = 0$ for all $\theta \in \bar{B}_0$. Hence, $\theta \notin B_0 \Rightarrow R(\theta, B) = \tilde{R}(\theta, B)$. Finally, note that A_0 is determined by Q and \tilde{Q} , and B_0 is determined by A_0 . Therefore, B_0 doesn't depend on the set B , so $\theta \notin B_0 \Rightarrow R(\theta, \cdot) = \tilde{R}(\theta, \cdot)$. \square

Remark 1. An important consequence of Proposition 2 is that either all versions of R are locally- ν -recurrent, or none of them is.

Every FPD satisfies $P(dx|\theta)\nu(d\theta) = Q(d\theta|x)M(dx)$, and it follows that every version of R is reversible (or symmetric) with respect to ν ; that is, if f and g are bounded, real-valued functions on Θ , then

$$\int_{\Theta} \int_{\Theta} f(\theta)g(\eta)R(\theta, d\eta)\nu(d\theta) = \int_{\Theta} \int_{\Theta} f(\theta)g(\eta)R(\eta, d\theta)\nu(d\eta) .$$

This property is key in the proof of our main result, which we now state and prove.

Theorem 4. *There exists a ν -reducible version of R if and only if there exist a nonempty set $C \in \mathcal{B}(\Theta)$ with $\nu(\bar{C}) > 0$ and another set $A \in \mathcal{B}(X)$ such that $P(\bar{A}|\theta) = 0$ for every $\theta \in C$ and $P(A|\theta) = 0$ for ν -a.e. $\theta \in \bar{C}$.*

Proof. To prove sufficiency, suppose that A and C exist. Using property 3 of Definition 1, we have

$$\int_A Q(\bar{C}|x)M(dx) = \int_{\bar{C}} P(A|\theta)\nu(d\theta) = 0 .$$

Hence, if

$$D = \{x \in A : Q(\bar{C}|x) > 0\} ,$$

then $M(D) = 0$ and, obviously, $Q(\bar{C}|x) = 0$ for all $x \in A - D$. Fix $\theta_0 \in C$ and let $\delta_{\theta_0}(d\theta)$ denote a probability measure concentrated on the point θ_0 . Now define

$$\tilde{Q}(d\theta|x) = \begin{cases} Q(d\theta|x) & \text{if } x \notin D \\ \delta_{\theta_0}(d\theta) & \text{if } x \in D \end{cases}$$

Clearly, \tilde{Q} satisfies the first and third properties of Definition 1. Moreover, we show in the Appendix that, for any $B \in \mathcal{B}(\Theta)$, $\tilde{Q}(B|\cdot)$ is measurable. Therefore, \tilde{Q} is a FPD and we now show that

$$\tilde{R}(\theta, d\eta) = \int_X \tilde{Q}(d\eta|x)P(dx|\theta)$$

is the ν -reducible version of R that we seek. By construction, $\tilde{Q}(\overline{C}|x) = 0$ for all $x \in A$. It follows that, for every $\theta \in C$,

$$\tilde{R}(\theta, \overline{C}) = \int_{\mathsf{X}} \tilde{Q}(\overline{C}|x)P(dx|\theta) = \int_A \tilde{Q}(\overline{C}|x)P(dx|\theta) = 0 .$$

Consequently, C is a closed set with $\nu(\overline{C}) > 0$, and it follows from Theorem 3 that \tilde{R} is ν -reducible.

To prove necessity, assume that R is ν -reducible. By Theorem 3 there exists a closed set C with $\nu(\overline{C}) > 0$. Using the reversibility of R , we have

$$\int_{\overline{C}} R(\theta, C)\nu(d\theta) = \int_C R(\theta, \overline{C})\nu(d\theta) = 0 .$$

This, of course, implies that $R(\theta, C) = 0$ for ν -a.e. $\theta \in \overline{C}$. Now define $F_1 = \{x \in \mathsf{X} : Q(C|x) > 0\}$ and $F_2 = \{x \in \mathsf{X} : Q(\overline{C}|x) > 0\}$. Since $Q(C|x) + Q(\overline{C}|x) = 1$ for every $x \in \mathsf{X}$, we have $F_1 \cup F_2 = \mathsf{X}$. Consequently, $\overline{F_2} \subset F_1$. Since C is closed, we know that for any $\theta \in C$,

$$0 = R(\theta, \overline{C}) = \int_{\mathsf{X}} Q(\overline{C}|x)P(dx|\theta) = \int_{F_2} Q(\overline{C}|x)P(dx|\theta) .$$

Thus, $P(F_2|\theta) = 0$ for all $\theta \in C$. Similarly, for ν -a.e. $\theta \in \overline{C}$, we have

$$0 = R(\theta, C) = \int_{\mathsf{X}} Q(C|x)P(dx|\theta) = \int_{F_1} Q(C|x)P(dx|\theta) .$$

Therefore, $P(F_1|\theta) = 0$ for ν -a.e. $\theta \in \overline{C}$, and since $\overline{F_2} \subset F_1$, it follows that $P(\overline{F_2}|\theta) = 0$ for ν -a.e. $\theta \in \overline{C}$. Letting $A = \overline{F_2}$, we have $P(\overline{A}|\theta) = 0$ for all $\theta \in C$ and $P(A|\theta) = 0$ for ν -a.e. $\theta \in \overline{C}$. \square

Obviously, a sufficient condition for ν -irreducibility of R is the nonexistence of the sets A and C in Theorem 4. This is precisely what Theorem 2 says. However, as will be clear from the examples that follow, even when the sets A and C do exist, it is often the case that some versions of R are ν -irreducible. One way to establish local- ν -recurrence (of all versions) of R is to identify a single version of R that is ν -irreducible and then show that this version of R is recurrent. The result then follows from the equivalence of recurrence and local- ν -recurrence (under irreducibility) mentioned in Section 1. We now illustrate the use of Theorem 4 with two examples.

Example 1. Let $\mathsf{X} = \Theta = \mathbb{R}$ and $P(dx|\theta) = p(x|\theta)dx$ where

$$p(x|\theta) = I_{(\theta, \theta+1)}(x) ,$$

and dx denotes Lebesgue measure on X . Take the prior distribution to be $\nu(d\theta) = d\theta$. A simple calculation shows that the marginal measure is $M(dx) = dx$, which is clearly σ -finite. This is all the information we require to apply Theorem 4. Let $C \in \mathcal{B}(\Theta)$ be a nonempty set such that $\nu(\overline{C}) > 0$. We claim that there exist $\theta' \in C$ and $D \subset \overline{C}$ such that $\nu(D) > 0$ and $|\theta - \theta'| < 1/4$ for every $\theta \in D$. To

see this, let $D_i = [\frac{i}{8}, \frac{i+1}{8})$ for all $i \in \mathbb{Z} := \{\dots, -1, 0, 1, \dots\}$. Then $\Theta = \cup_{i \in \mathbb{Z}} D_i$. If there exists an i_0 such that $C \cap D_{i_0}$ is not empty and $\nu(\overline{C} \cap D_{i_0}) > 0$, then we can simply take θ' to be any point in $C \cap D_{i_0}$ and D to be the set $\overline{C} \cap D_{i_0}$. Otherwise, for every $i \in \mathbb{Z}$, either $D_i \subset \overline{C}$ or $\nu(\overline{C} \cap D_i) = 0$. Since $\nu(\overline{C}) > 0$, there must exist an i_0 such that $\nu(\overline{C} \cap D_{i_0}) > 0$, and it follows that $D_{i_0} \subset \overline{C}$. Assume, without loss of generality, that $C \cap [\frac{i_0}{8}, \infty)$ is nonempty and let $i_1 = \min\{i > i_0 : C \cap D_i \text{ is nonempty}\}$, which is clearly finite. Now, any point in $C \cap D_{i_1}$ and the set D_{i_1-1} play the roles of θ' and D . These arguments show that there exist $\theta \in \overline{C}$ and $\theta' \in C$ such that $|\theta - \theta'| < 1/4$ where θ can be chosen outside of any subset of \overline{C} having Lebesgue measure zero. It follows that if $I_1 = (\theta, \theta + 1)$, $I_2 = (\theta', \theta' + 1)$ and $I = I_1 \cap I_2$, then $\nu(I) > 0$. Now, if there exists an $A \in \mathcal{B}(X)$ such that $P(A|\theta) = 0$ and $P(\overline{A}|\theta') = 0$, then

$$\nu(I) = \nu(A \cap I) + \nu(\overline{A} \cap I) \leq \nu(A \cap I_1) + \nu(\overline{A} \cap I_2) = 0 ,$$

which is a contradiction. Hence, such an A cannot exist and it follows from Theorem 4 that every version of Eaton's Markov chain is ν -irreducible.

For the sake of comparison, we now explain what is required to establish irreducibility in this situation if we make no appeal to Theorem 4. A version of the posterior is given by $Q(d\theta|x) = q(\theta|x)d\theta$ where

$$q(\theta|x) = I_{(x-1,x)}(\theta) .$$

It follows that $R(\theta, d\eta) = r(\eta|\theta)d\eta$ where

$$r(\eta|\theta) = (1 + (\eta - \theta))I_{(-1,0)}(\eta - \theta) + (1 - (\eta - \theta))I_{(0,1)}(\eta - \theta) .$$

Since θ is a location parameter in the density $r(\eta|\theta)$, the Markov chain W can be expressed as a random walk

$$W_{n+1} = W_n + Z_{n+1} ,$$

where Z_1, Z_2, \dots is an independent and identically distributed (iid) sequence of random variables with density given by

$$f(z) = (1 + z)I_{(-1,0)}(z) + (1 - z)I_{(0,1)}(z) .$$

For any $\varepsilon \in (0, 1)$, we have

$$P(Z_1 \in (0, \varepsilon)) = P(Z_1 \in (-\varepsilon, 0)) > 0 .$$

This implies that the chain W can make arbitrarily small jumps in either direction. While this argument makes it intuitively clear that the chain is ν -irreducible, a formal proof requires a technical argument similar to that used in Section 4.3.3 of ?.

It turns out that this random walk is recurrent, which implies that W is locally- ν -recurrent. Hence, by Theorem 1, ν is strongly admissible. In fact, ? shows that, under very mild conditions, Lebesgue measure is a strongly admissible prior for one and two-dimensional location problems. This ends Example 1.

Example 2. Let $\Theta = \mathbb{R}_+ := [0, \infty)$ and $\mathsf{X} = \mathbb{R}_+^n$. Suppose that when $\theta > 0$ the statistical model $P(dx|\theta)$ has a density (with respect to Lebesgue measure on X) given by

$$p(x|\theta) = \prod_{i=1}^n \frac{1}{\theta} I_{[0,\theta)}(x_i)$$

where $x = (x_1, \dots, x_n)$. In words, our statistical model stipulates that X_1, \dots, X_n are iid random variables from the uniform distribution on $[0, \theta)$. Take the prior distribution to be $\nu(d\theta) = \frac{d\theta}{\theta}$ where $d\theta$ denotes Lebesgue measure on Θ . While we have yet to define $P(dx|\theta)$ in the case where $\theta = 0$, from a practical (statistical) standpoint this definition is irrelevant since $\nu(\{0\}) = 0$. However, technically speaking, this distribution must be specified to complete the model. (Note that we cannot simply remove the point $\{0\}$ from Θ because $(0, \infty)$ is not a Polish space.) We consider two alternatives for $P(dx|0)$:

1. a unit point mass at $(0, \dots, 0) \in \mathbb{R}^n$, and
2. n iid exponential random variables with unit scale.

We now apply Theorem 4. In case 1, there does exist a ν -reducible version of R since we can take A to be the point $(0, \dots, 0) \in \mathbb{R}^n$ and $C = \{0\}$. However, in case 2, the sets A and C do not exist. Indeed, let μ denote Lebesgue measure on X and note that it is impossible to satisfy the conditions of Theorem 4 if $\mu(\bar{A}) = 0$ since this yields $P(A|\theta) = 1$ for all $\theta \in \Theta$. Now if $\mu(\bar{A}) > 0$ and $P(\bar{A}|\theta) = 0$, then it must be true that $\theta > 0$ and $\mu(\bar{A} \cap [0, \theta)^n) = 0$. Similarly, if $\theta^* > 0$ and $P(A|\theta^*) = 0$, then $\mu(A \cap [0, \theta^*)^n) = 0$. Now let $\theta_0 = \min\{\theta, \theta^*\}$ and note that

$$\mu([0, \theta_0)^n) = \mu(A \cap [0, \theta_0)^n) + \mu(\bar{A} \cap [0, \theta_0)^n) \leq \mu(A \cap [0, \theta^*)^n) + \mu(\bar{A} \cap [0, \theta)^n) = 0.$$

This is a contradiction, which implies that the sets A and C do not exist. It follows that every version of R is ν -irreducible. Note that, by defining the statistical model carefully on an irrelevant (ν -null) set of θ s, we were able to employ Theorem 4 to show that all versions of R are ν -irreducible.

Regardless of how $P(dx|0)$ is defined, the marginal measure is given by

$$M(dx) = \frac{dx}{nx_{(n)}^n},$$

where $x_{(n)} := \max\{x_1, \dots, x_n\}$ and dx denotes Lebesgue measure on X . Consider case 1 again. For x such that $x_{(n)} > 0$, define

$$q(\theta|x) = \frac{nx_{(n)}^n}{\theta^{n+1}} I_{(x_{(n)}, \infty)}(\theta).$$

Since the point $(0, \dots, 0) \in \mathbb{R}^n$ has M -measure 0, $Q(d\theta|x)$ can be chosen essentially arbitrarily when $x_{(n)} = 0$. We consider two different choices. Let $\delta_0(d\theta)$ denote the probability measure concentrated at $\{0\}$ and let $\delta_1(d\theta)$ denote a probability measure with support \mathbb{R}_+ . (It seems more appropriate to take

$Q(d\theta|x)$ equal to $\delta_0(d\theta)$ when $x_{(n)} = 0$.) Two different versions of the posterior distribution are given by

$$Q_i(d\theta|x) = \begin{cases} q(\theta|x)d\theta & x_{(n)} > 0 \\ \delta_i(d\theta) & x_{(n)} = 0 \end{cases}$$

for $i \in \{0, 1\}$. The version of Eaton's chain associated with Q_0 is ν -reducible since, if the chain is started at $\theta = 0$, it stays there forever. On the other hand, ? show that the version associated with Q_1 is ν -irreducible, and they go on to show that this chain is recurrent. Consequently, this version of Eaton's chain is locally- ν -recurrent and it follows from Theorem 1 that ν is strongly admissible. This ends Example 2.

A On the measurability of \tilde{Q}

Here we establish the measurability of \tilde{Q} . Fix $B \in \mathcal{B}(\Theta)$. It suffices to show that, for any $t \in \mathbb{R}$, the set

$$\tilde{G}_t = \{x \in \mathcal{X} : \tilde{Q}(B|x) < t\}$$

is in $\mathcal{B}(\mathcal{X})$. We will accomplish this using the partition

$$\tilde{G}_t = (\tilde{G}_t \cap D) \cup (\tilde{G}_t \cap \overline{D}).$$

The set $G_t = \{x \in \mathcal{X} : Q(B|x) < t\} \in \mathcal{B}(\mathcal{X})$ since $Q(B|\cdot)$ is a measurable function. The measures Q and \tilde{Q} agree on \overline{D} , and hence

$$\tilde{G}_t \cap \overline{D} = \{x \in \overline{D} : \tilde{Q}(B|x) < t\} = \{x \in \overline{D} : Q(B|x) < t\} = G_t \cap \overline{D} \in \mathcal{B}(\mathcal{X}).$$

It remains to show that $\tilde{G}_t \cap D \in \mathcal{B}(\mathcal{X})$. There are four possible cases:

1. $\theta_0 \in B$ and $t \leq 1$
2. $\theta_0 \in B$ and $t > 1$
3. $\theta_0 \notin B$ and $t \leq 0$
4. $\theta_0 \notin B$ and $t > 0$

In the first two cases, $\tilde{Q}(B|x) = 1$ for all $x \in D$. Hence, in the first case, $\tilde{G}_t \cap D = \emptyset \in \mathcal{B}(\mathcal{X})$, and in the second, $\tilde{G}_t \cap D = D \in \mathcal{B}(\mathcal{X})$. In the last two cases, $\tilde{Q}(B|x) = 0$ for all $x \in D$. Hence, in the third case, $\tilde{G}_t \cap D = \emptyset \in \mathcal{B}(\mathcal{X})$, and in the last $\tilde{G}_t \cap D = D \in \mathcal{B}(\mathcal{X})$.

Acknowledgments

Hubert's research partially supported by NSF Grant DMS-05-03648.