Chapter

Gibbs measures

This Chapter contains a simplified account of some theory for Gibbs measures, which I learned from the very thorough monograph by Georgii (1988) with a little help from the gentler exposition by Kindermann & Snell (1980).

1. Notation

- Let S be a countably infinite index set of ***sites***. For each i in S, suppose $\mathcal{X}_i$ is a set equippped with a sigma-field $\mathcal{B}_i$. For each $A \subseteq S$ write Ω_A for $\times_{i \in A} \mathcal{X}_i$. Equip Ω_A with its product sigma-field $\mathcal{B}_A = \bigotimes_{i \in A} \mathcal{B}_i$. Abbreviate Ω_S to Ω and $\mathcal{B}_S$ to $\mathcal{B}$.
- Write $\omega_A = (\omega_i : i \in A)$ for both the generic point of Ω_A and for the coordinate projection of a generic ω in Ω onto Ω_A.
- By definition, $\mathcal{B}$ is the smallest sigma-field on Ω for which each coordinate projection $\omega \mapsto \omega_i$ is $\mathcal{B}\backslash\mathcal{B}_i$-measurable. Consequently, for each $A \subseteq S$, the projection map $\omega \mapsto \omega_A$ is $\mathcal{B}\backslash\mathcal{B}_A$-measurable. Write $\mathcal{F}_A$ for the smallest sigma-field on Ω for which the map $\omega \mapsto \omega_A$ is $\mathcal{F}_A\backslash\mathcal{B}_A$-measurable. Each set in $\mathcal{F}_A$ is of the form $B \times \Omega_{S\backslash A}$ with $B \in \mathcal{B}_A$.
- Write $\mathbb{S}$ for the set of all finite, nonempty subsets of S.
- Let $\mathcal{H}$ denote the set of all bounded, real-valued, $\mathcal{B}$-measurable functions on Ω. For $A \subseteq \mathbb{S}$ define

  $$\mathcal{L}_A = \{f \in \mathcal{H} : f \text{ depends on } \omega \text{ only through the coordinates } \omega_A\}$$

  That is, $f \in \mathcal{L}_A$ if and only if $f(\omega) = g(\omega_A)$ for some bounded, $\mathcal{B}_A\backslash\mathcal{B}(\mathbb{R})$-measurable function g on Ω_A. The functions in $\mathcal{L}_A$ generate a sigma-field $\mathcal{F}_A$ on Ω. Write $\mathcal{L}$ for $\cup_{A\in\mathbb{S}}\mathcal{L}_A$, the set of all functions in $\mathcal{H}$ that depend on ω only through some finite subset of coordinates. If $f \in \mathcal{L}$, write $S(f)$ for the smallest A such that $f \in \mathcal{L}_A$.
- Call a sequence $\{A(n) : n \in \mathbb{N}\} \subset \mathbb{S}$ an ***$\mathbb{S}$-cover for S*** if $A(n) \uparrow S$ as $n \to \infty$.

2. A cautionary example

Ignore this Section.

We will be building Gibbs measures from a collection of desired conditional distributions. As you saw in Chapter MRF, it is not a completely trivial task to find Markov kernels that have the consistency properties required for conditional distributions. For a finite S, the construction of a Gibbs measure via a density (defined by a family Ψ of nonnegative functions) guarantees the necessary consistency. We cannot follow exactly the same route when S is infinite, because it is not always possible to define a joint density for ω_S by taking an infinite product of Ψ functions.

better example needed

<1> **Example.** Suppose $\mathcal{X}_i = \{0, 1\}$ and λ^i is the uniform distribution on $\mathcal{X}_i$, for every $i \in S = \mathbb{N}$. Suppose $\mathbb{F} = \{\{i, i+1\} : i \in S\}$ and

$$\Psi_{\{i,i+1\}} = 2\{\omega_i = \omega_{i+1}\} + \{\omega_i \neq \omega_{i+1}\}.$$

The product measure $\lambda^S = \otimes_{i \in S}\lambda^i$ is well defined. We might hope to construct $\mathbb{P}$ by defining

$$\frac{d\mathbb{P}}{d\lambda^S}(\omega_S) = \frac{1}{Z}\prod\nolimits_{i\in S}\Psi_{\{i,i+1\}}(\omega_i, \omega_{i+1})$$
$$\text{where } Z = \lambda^S \prod\nolimits_{i\in S}\Psi_{\{i,i+1\}}(\omega_i, \omega_{i+1}).$$

Unfortunately,

$$Z \geq \prod\nolimits_{j\in\mathbb{N}} \lambda \otimes \lambda\Psi_{\{2j,2j+1\}}(\omega_{2j}, \omega_{2j+1}) = \prod\nolimits_{j\in\mathbb{N}}\left(\tfrac{1}{2} \times 2 + \tfrac{1}{2} \times 1\right) = \infty$$

□ We would end up with ∞/∞.

---

3. Consistent sets of conditional distributions

For each i in S, suppose λ^i is a sigma-finite measure on $\mathcal{B}_i$. For each A in $\mathbb{S}$, write λ^A for the product measure $\bigotimes_{i\in A}\lambda^i$ on $\mathcal{B}_A$.

For some index set $\mathbb{F} \subseteq \mathbb{S}$, suppose $\Psi := \{\Psi_a : a \in \mathbb{F}\}$ is a collection of nonnegative, $\mathcal{B}$-measurable functions on Ω for which Ψ_a depends on ω only through the coordinates ω_a.

Some important features of Ψ are captured by its factor graph, which has a node for each site i in S and a node for each factor a in $\mathbb{F}$, with i connected to a if and only if i is one of the sites in a. For each $A \subseteq S$ define

$$\partial A := \{a \in \mathbb{F} : A \cap a \neq \emptyset\}$$
$$\mathcal{N}(A) := \{j \in S \backslash A : j \in a \text{ for some } a \text{ in } \partial A\}$$

To avoid problems with infinite products, I will assume that both ∂A and $\mathcal{N}(A)$ are finite if $A \in \mathbb{S}$. For each such A define

$$G_A(\omega) = \prod\nolimits_{a\in\partial A}\Psi_a(\omega_a).$$

Sometimes I will write $G(\omega_A, \omega_{\mathcal{N}(A)})$ to emphasize the fact that G_A depends on ω only through the coordinates ω_i for $i \in A \cup \mathcal{N}(A)$. Similarly, define

$$Z_A(\omega_{S\backslash A}) = Z_A(\omega_{\mathcal{N}(A)}) := \lambda^A G_A(\omega_A, \omega_{\mathcal{N}(A)}),$$

where the λ^A integrates out over the ω_A coordinates to leave a dependence on (a subset of) the $\omega_{S\backslash A}$ coordinates.

> Assume $Z_A(\omega_{S\backslash A}) < \infty$ for each $\omega \in \Omega$. At the risk of some unforeseen complications, I will not assume that Z_A is everywhere strictly positive. However, for the purposes of Stat 606, you could safely restrict yourself to the case where $Z_A(\omega_{S\backslash A}) > 0$ for every ω.

Define probability measures $Q_{S\backslash A}(\cdot \mid \omega_{S\backslash A})$ on $\mathcal{B}$ for each $A \in \mathbb{S}$ and each $\omega_{S\backslash A} \in \Omega_{S\backslash A}$ for which $Z_A(\omega_{S\backslash A}) \neq 0$ by

$$Q_{S\backslash A}(f \mid \omega_{S\backslash A}) = \frac{1}{Z_A(\omega_{S\backslash A})}\lambda^A f(\omega_A, \omega_{S\backslash A}) G_A(\omega_A, \omega_{\mathcal{N}(A)}) \qquad \text{for } f \in \mathcal{H}.$$

Here the $\omega_{S\backslash A}$ fixes both the $\omega_{\mathcal{N}(A)}$ coordinates for G_A and some of the coordinates for f. The λ^A integrates out over the ω_A coordinates. If $Z_A(\omega_{S\backslash A}) = 0$, define $Q_{S\backslash A}(\cdot \mid \omega_{S\backslash A})$ to be the zero measure.

REMARK. If $f \in \mathcal{L}_D$ then $Q_{S\backslash A}(f \mid \omega_{S\backslash A})$ depends on ω only through the coordinates ω_i for $i \in (D\backslash A) \cup \mathcal{N}(A)$.

We could regard $Q_{S\backslash A}$ as a linear map from $\mathcal{H}$ into $\mathcal{L}_{S\backslash A}$, in which case it would be natural to omit the explicit $\omega_{S\backslash A}$ and write just $Q_{S\backslash A} f$.

<2> **Lemma.** *Suppose $A, B \in \mathbb{S}$ with $A \subseteq B$. For each $\omega_{S\backslash B} \in \Omega_{S\backslash B}$ the following two properties hold.*

(i) $Q_{S\backslash B}\{\omega : Z_A(\omega_{S\backslash A}) = 0 \mid \omega_{S\backslash B}\} = 0$

(ii) for each $f \in \mathcal{H}$,

$$Q_{S\backslash B}(f \mid \omega_{S\backslash B}) = Q_{S\backslash B}\left[Q_{S\backslash A}\left(f(\omega_A, \omega_{B\backslash A}, \omega_{S\backslash B}) \mid \omega_{S\backslash A}\right) \mid \omega_{S\backslash B}\right]$$

REMARK. On the right-hand side in (ii), the $Q_{S\backslash A}$ integrates over ω_A with the coordinates $\omega_{B\backslash A}, \omega_{S\backslash B}$ being fixed by $\omega_{S\backslash A}$. The $Q_{S\backslash B}$ then integrates over $\omega_{B\backslash A}$ for fixed $\omega_{S\backslash B}$.

Proof. We have only to consider the case of a fixed $\omega_{S\backslash B}$ for which $Z_B(\omega_{S\backslash B})$ is nonzero, for otherwise both assertions are trivially true.

Temporarily write N for $\mathcal{N}(A)$. Note that G_B factorizes as

<3> $$G_B(\omega) = G_A(\omega_A, \omega_N) H(\omega_{S\backslash A}) \qquad \text{where } H(\omega_{S\backslash A}) = \prod_{a \in \partial B\backslash\partial A} \Psi_a(\omega_a).$$

In fact, H depends only on coordinates ω_i for $i \in (B \cup \mathcal{N}(B))\backslash A$.

Use the fact that $\lambda^B = \lambda^{B\backslash A} \otimes \lambda^A$ to write Z_B times the left-hand side of (i) as

$$\begin{aligned} &\lambda^{B\backslash A}\lambda^A G_A(\omega_A, \omega_N) H(\omega_{S\backslash A})\{Z_A(\omega_{S\backslash A}) = 0\} \\ &\qquad = \lambda^{B\backslash A} Z_A(\omega_{S\backslash A}) H(\omega_{S\backslash A})\{Z_A(\omega_{S\backslash A}) = 0\} = 0. \end{aligned}$$

By virtue of (i), neither side of (ii) is changed if we replace f by $\{Z_A(\omega_{S\backslash A}) \neq 0\} f$. Define

$$\begin{aligned} F(\omega_{S\backslash A}) &= \{Z_A(\omega_{S\backslash A}) \neq 0\} Q_{S\backslash A}\left(f \mid \omega_{S\backslash A}\right) \\ &= \frac{\{Z_A(\omega_{S\backslash A}) \neq 0\}}{Z_A(\omega_{S\backslash A})}\lambda^A G_A(\omega_A, \omega_N) f(\omega). \end{aligned}$$

Then

$$\begin{aligned} Z_B(\omega_{S\backslash B}) \times \text{RHS of (ii)} &= \lambda^{B\backslash A}\lambda^A G_A(\omega_A, \omega_N) H(\omega_{S\backslash A}) F(\omega_{S\backslash A}) \\ &= \lambda^{B\backslash A} H(\omega_{S\backslash A}) F(\omega_{S\backslash A}) Z_A(\omega_{S\backslash A}) \\ &= \lambda^{B\backslash A} H(\omega_{S\backslash A})\{Z_A \neq 0\}\lambda^A \left(G_A(\omega_A, \omega_N) f(\omega)\right) \\ &= \lambda^B G_B(\omega)\{Z_A \neq 0\} f(\omega) \qquad \text{by <3>} \\ &= Z_B(\omega_{S\backslash B}) \times \text{LHS of (ii)}. \end{aligned}$$

□

REMARK. Most of the theory depends on Ψ only through the $Q_{S\backslash A}$ measures, which are often referred to as a ***specification***. See Georgii (1988, page 16), for example.

4. Existence of Gibbs measures

A probability measure $\mathbb{P}$ on $\mathcal{B}$ is said to be a ***Gibbs measure*** for the family Ψ if it has the Q's from the previous Section as its conditional distributions, that is, for every A in $\mathbb{S}$ and at least for $f \in \mathcal{H}$,

$$\mathbb{P}(f \mid \mathcal{F}_{S\setminus A}) = Q_{S\setminus A}\left(f \mid \omega_{S\setminus A}\right) \qquad \text{a.e. } [\mathbb{P}],$$

Equivalently,

<4> $$\mathbb{P}f = \mathbb{P}Q_{S\setminus A}\left(f \mid \omega_{S\setminus A}\right) \qquad \text{for all } f \in \mathcal{H}.$$

The set of all Gibbs probability measures for a given Ψ is denoted by $\mathbb{G}(\Psi)$.

> REMARK. Some authors would call $Q_{S\setminus A}(\cdot \mid \omega_{S\setminus A})$ a ***regular conditional distribution*** for $\mathbb{P}$ given $\mathcal{F}_{S\setminus A}$. In my opinion, it is a backward step to express the conditioning properties of $\mathbb{P}$ in terms of Kolmogorov conditional expectations when we know that regular conditional distributions exist.

At an ω for which $Z_D(\omega_{S\setminus D}) = 0$, the conditional distribution $Q_{S\setminus D}$ is the zero measure and not a probability. This is of no major importance because, from part (i) of Lemma <2>,

$$\mathbb{P}\{\omega : Z_D(\omega_{S\setminus D}) = 0\} = \mathbb{P}Q_{S\setminus A}\{Z_A(\omega_{S\setminus D}) = 0\} = 0$$

if $\mathbb{P}$ satisfies <4> for $f = \{Z_D = 0\}$, for some $A \supseteq D$.

It is not completely obvious that Gibbs measures exist—that $\mathbb{G}(\Psi)$ is nonempty—for any particular Ψ. At least when each $\mathcal{X}_i$ is a finite set, it is very easy to prove existence because (as a special case of the Kolmogorov extension theorem) there is a one-to-one correspondence between probability measures on $\mathcal{B}$ and increasing linear functionals $\mathbb{P} : \mathcal{L} \to \mathbb{R}$ for which $\mathbb{P}1 = 1$.

see Kolmogorov.pdf

<5> **Theorem.** *Suppose each coordinate space $\mathcal{X}_i$ is finite. Suppose also that there exists an $\mathbb{S}$-cover $\{A(n) : n \in \mathbb{N}\}$ for which each of the sets $F_n := \{\omega_{S\setminus A(n)} \in \Omega_{S\setminus A(n)} : Z_{A(n)}(\omega_{S\setminus A(n)}) > 0\}$ is nonempty. Then $\mathbb{G}(\Psi) \neq \emptyset$.*

Proof. Define $S(n) := S\setminus A(n)$. Let ν_n be any probability measure on $\mathcal{B}_{S(n)}$ for which $\nu_n(F_n) = 1$. Define increasing linear functionals μ_n on $\mathcal{L}$ by

$$\mu_n f := \nu_n Q_{S(n)}(f \mid \omega_{S(n)}).$$

Note that $\mu_n 1 = 1$ for every n because $Q_{S(n)}(\Omega \mid \omega_{S(n)}) = 1$ for all $\omega_{S(n)} \in F_n$.

Identify μ_n with a point in the product space

$$\mathbb{K} = \mathsf{X}_{f\in\mathcal{L}}[-m_f, m_f] \qquad \text{where } m_f := \sup_\omega |f(\omega)|.$$

When equipped with its product topology (the weakest topology that makes each coordinate map $\kappa \mapsto \kappa(f)$ continuous), the space $\mathbb{K}$ is compact. The sequence $\{\mu_n : n \in \mathbb{N}\}$ has a cluster point, $\mathbb{P}$, in $\mathbb{K}$.

It is easy to show that $\mathbb{P}$ inherits from the μ_n's the linearity and increasing properties and that $\mathbb{P}1 = 1$; by the Kolmogorov extension theorem, it corresponds to a probability measure on $\mathcal{B}$.

To establish the defining property <4> for a Gibbs measure, consider first an f in some $\mathcal{L}_A$, with $A \in \mathbb{S}$, and an n so large that $A(n) \supseteq A$. From Lemmma <2>,

$$Q_{S(n)}(f \mid \omega_{S(n)}) = Q_{S(n)}\left[Q_{S\setminus A}\left(f \mid \omega_{S\setminus A}\right) \mid \omega_{S(n)}\right].$$

Integrate both sides with respect to ν_n to get

$$\mu_n f = \mu_n g \qquad \text{where } g(\omega_{S\setminus A}) = Q_{S\setminus A}\left(f \mid \omega_{S\setminus A}\right).$$

The function g depends on $\omega_{S\setminus A}$ only through the coordinates in $\mathcal{N}(A)$. Thus $g \in \mathcal{L}$. Let n tend to infinity (along a subsequence) to deduce that $\mathbb{P}f = \mathbb{P}g$.

see lambda-space.pdf

That is, the equality in <4> holds at least for $f \in \mathcal{L}$. An appeal to the π-λ-theorem for functions extends the equality to all f in $\mathcal{H}$. The cluster point $\mathbb{P}$ is a Gibbs probability measure. □

5. Representation of Gibbs measures as mixtures

Suppose $\mathbb{G}(\Psi)$ is nonempty. It follows directly from the defining property <4> that $\mathbb{G}(\Psi)$ is convex. A Gibbs measure $\mathbb{P}$ is said to be an ***extremal element*** of $\mathbb{G}(\Psi)$ if it cannot be written as a proper convex combination of two other Gibbs measures: if $\mathbb{P} = \theta\mathbb{P}_1 + (1-\theta)\mathbb{P}_2$ with $0 < \theta < 1$ and $\mathbb{P}_1, \mathbb{P}_2 \in \mathbb{G}(\Psi)$ then we must have $\mathbb{P}_1 = \mathbb{P}_2 = \mathbb{P}$. Write $\mathbb{G}_{\text{ex}}(\Psi)$ for the set of extreme elements of $\mathbb{G}(\Psi)$. This Section will show that there is a very simple way to characterize the extreme Gibbs measures and that, in an appropriate sense, $\mathbb{G}(\Psi)$ is the closed convex hull of $\mathbb{G}_{\text{ex}}(\Psi)$.

Georgii (1988, Chapter 7)

<6> **Definition.** *The tail sigma-field on Ω is defined as $\cap_{A\in\mathbb{S}}\mathcal{F}_{S\backslash A}$. Write $\mathcal{H}_{\text{tail}}$ for the set of all $\mathcal{T}$-measurable functions in $\mathcal{H}$.*

Remark. From the fact that $\mathcal{F}_{S\backslash A} \supseteq \mathcal{F}_{S\backslash B}$ when $A \subseteq B$, it is easy to see that if $\{A(n) : n \in \mathbb{N}\}$ is an $\mathbb{S}$-cover for S then $\mathcal{T} = \cap_{n\in\mathbb{N}}\mathcal{F}_{S\backslash A(n)}$.

<7> **Lemma.** *Suppose $\mathbb{P} \in \mathbb{G}(\Psi)$ and μ is another probability measure on $\mathcal{B}$ that is absolutely continuous with respect to $\mathbb{P}$. Then $\mu \in \mathbb{G}(\Psi)$ if and only if there exists a $\mathcal{T}$-measurable version of the density $\phi(\omega) = d\mu/d\mathbb{P}$.*

Proof. Suppose $\phi(\omega)$ is a $\mathcal{T}$-measurable version of $d\mu/d\mathbb{P}$. Then for each $f \in \mathcal{H}$ and each $A \in \mathbb{S}$,

$$\begin{aligned}
\mu f &= \mathbb{P}\phi f \\
&= \mathbb{P}\left(Q_{S\backslash A}(\phi(\omega)f(\omega) \mid \omega_{S\backslash A})\right) \\
&= \mathbb{P}\left(\phi(\omega)Q_{S\backslash A}(f(\omega) \mid \omega_{S\backslash A})\right) && \text{because } \phi \text{ is also } \mathcal{F}_{S\backslash A}\text{-measurable} \\
&= \mu Q_{S\backslash A}(f \mid \omega_{S\backslash A}).
\end{aligned}$$

It follows that $\mu \in \mathbb{G}(\Psi)$.

Conversely, suppose μ is a Gibbs measure. Let $S(n) := S\backslash A(n)$ for some $\mathbb{S}$-cover $\{A(n) : n \in \mathbb{N}\}$. The restriction of μ to $\mathcal{F}_{S(n)}$ is dominated by the restriction of $\mathbb{P}$ to $\mathcal{F}_{S(n)}$. Let $\phi_n(\omega_{S(n)})$ be an $\mathcal{F}_{S(n)}$-measurable choice for the corresponding density. For every n and $f \in \mathcal{H}$ we have

$$\begin{aligned}
\mathbb{P}f\phi = \mu f = \mu g \qquad & \text{where } g(\omega_{S(n)}) = Q_{S(n)}(f \mid \omega_{S(n)}) \\
&= \mathbb{P}\phi_n(\omega_{S(n)})g(\omega_{S(n)}) && \text{because } g \text{ is } \mathcal{F}_{S(n)}\text{-measurable} \\
&= \mathbb{P}Q_{S(n)}\left(f(\omega)\phi_n(\omega_{S(n)}) \mid \omega_{S(n)}\right) && \text{because } \phi_n \text{ is } \mathcal{F}_{S(n)}\text{-measurable} \\
&= \mathbb{P}f\phi_n.
\end{aligned}$$

It follows that $\phi_n = \phi$ a.e. $[\mathbb{P}]$ for each n and hence $\liminf_n \phi_n$ is a $\mathcal{T}$-measurable version of the density. □

<8> **Definition.** *A probability measure $\mathbb{P}$ on $\mathcal{B}$ is said to be* ***trivial on*** *$\mathcal{T}$ if $\mathbb{P}F$ is either 0 or 1 for each F in $\mathcal{T}$. Equivalently, $\mathbb{P}(f \mid \mathcal{H}_{\text{tail}}) = \mathbb{P}f$ a.e. $[\mathbb{P}]$ for each f in $\mathcal{H}$.*

<9> **Theorem.** *Suppose $\mathbb{P} \in \mathbb{G}(\Psi)$.*

(i) $\mathbb{P} \in \mathbb{G}_{\text{ex}}(\Psi)$ if and only if $\mathbb{P}$ is trivial on $\mathcal{T}$.

(ii) If $\mu \in \mathbb{G}(\Psi)$ and $\mu F = \mathbb{P}F$ for all F in $\mathcal{T}$ then $\mu = \mathbb{P}$, as measures on $\mathcal{B}$.

(iii) If $\mathbb{P}, \mu \in \mathbb{G}_{\text{ex}}(\Psi)$ and $\mathbb{P} \neq \mu$ then the two measures are mutually singular.

Proof. Suppose $F_0 \in \mathcal{T}$ with $0 < \mathbb{P}F_0 < 1$. Define $\mu_i(\cdot) = \mathbb{P}(\cdot \mid F_i)$, where $F_1 = \Omega\backslash F_0$. That is, $d\mu_i/d\mathbb{P} = \{\omega \in F_i\}/\mathbb{P}F_i$, a $\mathcal{T}$-measurable function of ω for $i = 0, 1$. By Lemma <7>, the μ_i are distinct (because $\mu_0 F_0 = 1 = \mu_1 F_1$) Gibbs measures for which $\mathbb{P} = (\mathbb{P}F_0)\mu_0 + (\mathbb{P}F_1)\mu_1$. Thus $\mathbb{P}$ is not extremal.

Conversely, suppose $\mathbb{P}$ is trivial on $\mathcal{T}$ but $\mathbb{P}$ can be written as a convex combination of two Gibbs measures, $\theta\mathbb{P}_0 + (1-\theta)\mathbb{P}_1$. Again by Lemma <7>, there must exist $\mathcal{T}$-measurable versions of the densities $\phi_i = d\mathbb{P}_i/d\mathbb{P}$. Triviality implies $\phi_i = \mathbb{P}\phi_i = 1$ a.e. $[\mathbb{P}]$, ensuring that $\mathbb{P}_0 = \mathbb{P}_1 = \mathbb{P}$. The Gibbs measure $\mathbb{P}$ must be extremal.

For (ii), note that both μ and $\mathbb{P}$ are absolutely continuous with respect to the Gibbs measure $\mathbb{P}_0 = (\mu + \mathbb{P})/2$. Let $\phi(\omega)$ be a $\mathcal{T}$-measurable version of the density $d\mu/d\mathbb{P}_0$. For an f in $\mathcal{H}$ let $F = \mathbb{P}_0(f \mid \mathcal{T})$. Then

$$\begin{aligned}\mu f = \mathbb{P}_0\phi(\omega)f(\omega) &= \mathbb{P}_0\phi(\omega)F(\omega)\\ &= \tfrac{1}{2}(\mu + \mathbb{P})(\phi F)\\ &= \mu F \qquad \text{because } \mu = \mathbb{P} \text{ on } \mathcal{T}.\end{aligned}$$

Similarly, $\mathbb{P}f = \mathbb{P}F$. The equality $\mu F = \mathbb{P}F$, for the integrals of the $\mathcal{T}$-measurable function F, then implies $\mu f = \mathbb{P}f$.

For (iii), the measures μ and $\mathbb{P}$ must have different restrictions to $\mathcal{T}$. That is, there exists some $F \in \mathcal{T}$ such that $\mu F < \mathbb{P}F$. As both measures are extremal, we must have $\mu F = 0$ and $\mathbb{P}F = 1$. That is, $\mathbb{P}$ concentrates on F and μ concentrates on F^c. □

<10> **Theorem.** *Suppose each coordinate space $\mathcal{X}_i$ is finite.*

(i) There exists a set $\Omega_0 \in \mathcal{T}$ for which $\mathbb{P}\Omega_0 = 1$ for every $\mathbb{P}$ in $\mathbb{G}(\Psi)$.

(ii) There exists a collection $\{\lambda_\omega : \omega \in \Omega_0\} \subseteq \mathbb{G}_{\text{ex}}(\Psi)$ for which $\omega \mapsto \lambda_\omega f$ is $\mathcal{T}$-measurable and $\mathbb{P}f = \mathbb{P}\lambda_\omega f$ for each $f \in \mathcal{H}$ and each $\mathbb{P} \in \mathbb{G}(\Psi)$.

Proof. Let $S(n)) := S\backslash A(n)$ for a fixed $\mathbb{S}$-cover $\{A(n) : N \in \mathbb{N}\}$. Finiteness of each $\Omega_{A(n)}$ implies that the vector space $\mathcal{L}_{A(n)}$ is spanned by a finite set of functions. Taking all rational combinations of the basis functions we get a countable subset $\mathcal{L}_n^\circ$ that is uniformly dense in $\mathcal{L}_{A(n)}$. The countable set $\mathcal{L}^\circ := \cup_{n\in\mathbb{N}}\mathcal{L}_n^\circ$ is uniformly dense in $\mathcal{L}$.

Suppose $\mu \in \mathbb{G}$ and $f \in \mathcal{H}$. By the convergence theorem for reverse martingales,

$$Q_{S(n)}(f \mid \omega_{S(n)}) = \mu(f \mid \mathcal{F}_{S(n)}) \to \mu(f \mid \mathcal{T}) \qquad \text{a.e. } [\mu]$$

Put another way, μ gives measure one to the set

$$\Omega_f := \{\omega \in \Omega : \limsup_n Q_{S(n)}(f \mid \omega_{S(n)}) = \liminf_n Q_{S(n)}(f \mid \omega_{S(n)})\}$$

and on Ω_f the limit $\lambda_\omega := \lim_n Q_{S(n)}(f \mid \omega_{S(n)})$ exists and equals $\mu(f \mid \mathcal{T})$ a.e. $[\mu]$. The sets

$$\Omega_2 := \bigcap_{f\in\mathcal{L}^\circ}\Omega_f \qquad \text{and} \qquad \Omega_1 := \{\omega \in \Omega_2 : \limsup_n Q_{S(n)}(\Omega \mid \omega_{S(n)}) = 1\}$$

are both $\mathcal{T}$-measurable and have μ measure one, for each $\mu \in \mathbb{G}$.

Convergence of $Q_{S(n)}(f \mid \omega_{S(n)})$ for each f in the uniformly dense set $\mathcal{L}^\circ$ implies convergence for all f in $\mathcal{L}$. For each $\omega \in \Omega_1$, the functional λ_ω, as a map from $\mathcal{L}$ to $\mathbb{R}$, inherits linearity and the increasing property and $\lambda_\omega 1 = 1$. We may extend λ_ω to a probability measure on $\mathcal{B}$.

REMARK. Note: I am not asserting that $\lim_n Q_{S(n)}(f \mid \omega_{S(n)})$ exists for every f in $\mathcal{H}$. The operations needed to extend λ_ω from $\mathcal{L}$ to $\mathcal{H}$ need not commute with the limit operation that defines $\lambda_\omega f$ for $f \in \mathcal{L}$.

For each $f \in \mathcal{L}$, each $\omega \in \Omega_1$, and each $A \in \mathbb{S}$, Lemma <2> shows that

$$\begin{aligned} \lambda_\omega f &= \lim_n Q_{S(n)}\left(Q_{S\setminus A}(f \mid \omega_{S\setminus A}) \mid \omega_{S(n)}\right) \\ &= \lambda_\omega\left(Q_{S\setminus A}(f \mid \omega_{S\setminus A})\right) \qquad \text{because } Q_{S\setminus A} f \in \mathcal{L}. \end{aligned}$$

Thus $\lambda_\omega \in \mathbb{G}$ for each $\omega \in \Omega_0$.

Similarly, the equalities $\mu(f \mid \mathcal{T}) = \lambda_\omega f$ a.e. $[\mu]$ for each $f \in \mathcal{L}^\circ$ extend to $\mathcal{L}$ and then, by the usual sort of π-λ argument, to all f in $\mathcal{H}$. Consequently,

$$\mu f = \mu^\omega \lambda_\omega f \qquad \text{for all } f \in \mathcal{H} \text{ and all } \mu \in \mathbb{G}$$

It remains only to show that $\lambda_\omega \in \mathbb{G}_{\text{ex}}$ for μ almost all ω.

First we need to use $\mathcal{L}^\circ$ to characterize the measures in $\mathbb{G}_{\text{ex}}$. From Theorem <9>, for a Gibbs measure μ,

$$\mu \in \mathbb{G}_{\text{ex}} \qquad \text{iff} \qquad \lambda_\omega f = \mu(f \mid \mathcal{T}) = \mu f \quad \text{a.e. } [\mu] \text{ for each } f \in \mathcal{H}.$$

Uniform approximation followed by a π-λ argument gives

$$\{\omega \in \Omega_1 : \lambda_\omega f = \mu f \text{ for all } f \in \mathcal{H}\} = \cap_{f \in \mathcal{L}^\circ}\{\omega \in \Omega_1 : \lambda_\omega f = \mu f\}.$$

Define $F_f(\omega) := \{\omega \in \Omega_1\}\lambda_\omega f$. For a fixed f, we have $F_f(\omega) = \mu f$ a.e. $[\mu]$ if and only if $\mu\left(F_f(\omega) - \mu f\right)^2 = 0$. Thus

<11> $$\mu \in \mathbb{G}_{\text{ex}} \qquad \text{iff} \qquad \mu(F_f - \mu f)^2 = 0 \qquad \text{for all } f \in \mathcal{L}^\circ.$$

Specializing to the case $\mu = \lambda_{\omega'}$, we get

$$\begin{aligned} \Omega_0 := \{\omega' \in \Omega_1 : \lambda_{\omega'} \in \mathbb{G}_{\text{ex}}\} &= \cap_{f \in \mathcal{L}^\circ}\{\omega' \in \Omega_1 : \lambda_{\omega'}^\omega\left(F_f(\omega) - F_f(\omega')\right)^2 = 0\} \\ &= \cap_{f \in \mathcal{L}^\circ}\{\omega' \in \Omega_1 : \lambda_{\omega'} F_f^2 = F_f(\omega')^2\} \end{aligned}$$

The last representation shows that $\Omega_0 \in \mathcal{T}$.

Finally, note that $\mathbb{P}\Omega_0 = 1$ if $\mathbb{P} \in \mathbb{G}$ because, for each $f \in \mathcal{L}^\circ$,

$$\begin{aligned} \mathbb{P}^{\omega'}\lambda_{\omega'}^\omega\left(F_f(\omega) - F_f(\omega')\right)^2 &= \mathbb{P}^{\omega'}\lambda_{\omega'}^\omega F_f(\omega)^2 - \mathbb{P}^{\omega'} F_f(\omega')^2 \\ &= \mathbb{P}F_f^2 - \mathbb{P}F_f^2 = 0 \end{aligned}$$

Thus $\mathbb{P}f = \mathbb{P}^\omega\left(\{\omega \in \Omega_0\}\lambda_\omega f\right)$ is a representation of $\mathbb{P}$ as a mixture of extremal Gibbs measures. □

REMARK. Equip $\mathbb{G}$ with the smallest sigma-field $\mathcal{G}$ for which each of the maps $\mu \mapsto \mu f$, for $f \in \mathcal{H}$, is $\mathcal{G}\setminus\mathcal{B}(\mathbb{R})$-measurable. Then $\mathbb{G}_{\text{ex}} \in \mathcal{G}$ and $\ell : \omega \to \lambda_\omega$ is a $\mathcal{T}\setminus\mathcal{G}$-measurable map from Ω_0 into $\mathbb{G}_{\text{ex}}$. The image of a Gibbs measure $\mathbb{P}$ under the map ℓ is a probability measure π on $\mathbb{G}_{\text{ex}}$. We could think of $\mathbb{P}$ as a π average over $\mathbb{G}_{\text{ex}}$.

Uniqueness? cf Gheorghii p. 132.