TOPICS IN STATISTICAL PHYSICS AND PROBABILITY THEORY

INSTRUCTOR: RON PELED, TEL AVIV UNIVERSITY

1. Infinite-volume Gibbs measures

We discuss some basic facts regarding Gibbs measures on the infinite lattice $\mathbb{Z}^d$. Our discussion is restricted to an explicit case with nearest-neighbor interactions but the reader should note that the general theory allows much more flexibility; see Friedli and Velenik [3, Chapter 6] for additional details.

Preliminaries. Let $(S, \mathcal{S}, \lambda)$ be a measure space (λ is a positive measure, either finite or infinite). We assume that S is a Polish space (metric, separable and complete) with $\mathcal{S}$ its Borel sigma algebra (this is convenient in order to define Gibbs measures through regular conditional probabilities as we do, though it is possible to develop the theory under weaker assumptions). Fix an integer $d \geqslant 1$. We will consider various probability measures on the measurable space

$$\Omega := S^{\mathbb{Z}^d} = \{\varphi \mid \varphi : \mathbb{Z}^d \to S\} \quad \text{equipped with the product topology and Borel sigma algebra } \mathcal{F}.$$

We denote subsets of $\mathbb{Z}^d$ by Λ or Δ. For a subset $\Lambda \subseteq \mathbb{Z}^d$ and $\varphi \in \Omega$ we write φ_Λ for the restriction $\varphi|_\Lambda$. A *cylinder set* in Ω is a set of the form

$$\prod_{v \in \mathbb{Z}^d} E_v \text{ with } E_v \in \mathcal{S} \text{ for all } v \text{ and } E_v = S \text{ for all but finitely many } v.$$

We call a function $f : \Omega \to \mathbb{R}$ *local* if there exists a finite $\Lambda \subseteq \mathbb{Z}^d$ such that $f(\varphi) = f(\varphi')$ whenever $\varphi_\Lambda = \varphi'_\Lambda$. We say that f is *determined by* Λ (for any such Λ). We recall several basic facts about probability measures on Ω:

- A probability measure $\mathbb{P}$ on Ω is determined by the probabilities $\mathbb{P}(E)$ for all cylinder sets E, or by the expectations $\mathbb{P}(f)$ for all bounded, continuous local functions f.
- A sequence of probability measures $\mathbb{P}_n$ on Ω converges (in distribution, i.e., in the weak* topology) to a probability measure $\mathbb{P}$ on Ω if $\mathbb{P}(f) = \lim_{n\to\infty} \mathbb{P}_n(f)$ for all bounded, continuous functions f. This is equivalent to convergence for all bounded, continuous *local* functions f (see, e.g., [2, Chapter 3, Proposition 4.6(b)]).
- If S is compact then Ω and the set of probability measures on Ω is compact. Thus, in this case, any sequence of probability measures $\mathbb{P}_n$ on Ω has a convergent subsequence. Consequently, convergence of $\mathbb{P}_n$ itself (to some limit) is assured once $\lim_{n\to\infty} \mathbb{P}_n(f)$ exists for all bounded, continuous local functions f.
- For each measure $\mathbb{P}$ on Ω and sub-sigma-algebra $\mathcal{G} \subseteq \mathcal{F}$ there exists a regular conditional probability for $\mathbb{P}$ conditioned on $\mathcal{G}$. That is, there exists a function $\mathbb{Q} : \Omega \times \mathcal{F} \to [0, 1]$ satisfying that
  - $\mathbb{Q}(\varphi, \cdot)$ is a probability measure on Ω, except for φ in a $\mathbb{P}$-null set in $\mathcal{G}$.
  - For each $E \in \mathcal{F}$, $\mathbb{Q}(\cdot, E) = \mathbb{P}(E \mid \mathcal{G})$ except on a $\mathbb{P}$-null set in $\mathcal{G}$.

  In addition, if both $\mathbb{Q}_1, \mathbb{Q}_2$ satisfy the above two properties then the probability measures $\mathbb{Q}_1(\varphi, \cdot)$ and $\mathbb{Q}_2(\varphi, \cdot)$ are equal except for φ in a $\mathbb{P}$-null set in $\mathcal{G}$. See [1, Section 10.2] for proofs and additional details.

Finite-volume Gibbs measures. We restrict attention to measures defined via a nearest-neighbor interaction with soft constraints as follows. Let $h : S \times S \to (0, \infty)$ be a measurable function satisfying

$$h(a, b) = h(b, a) \quad \text{for all } a, b \in S. \tag{1}$$

Date: June 11, 2017.

For a finite set $\Lambda \subseteq \mathbb{Z}^d$ and $\eta : \mathbb{Z}^d \to S$ we let $\mathbb{P}^\eta_\Lambda$ be the probability measure on Ω defined by

$$d\mathbb{P}^\eta_\Lambda(\varphi) := \frac{1}{Z^\eta_\Lambda} \prod_{\substack{u \sim v\\ \{u,v\}\cap\Lambda\neq\emptyset}} h(\varphi_u, \varphi_v) \prod_{v\in\Lambda} d\lambda(\varphi_v) \prod_{v\in\mathbb{Z}^d\setminus\Lambda} d\delta_{\eta_v}(\varphi_v) \tag{2}$$

where δ_s is the Dirac delta measure at s, so that the measure $\mathbb{P}^\eta_\Lambda$ is supported on configurations φ satisfying $\varphi_{\Lambda^c} = \eta_{\Lambda^c}$, and where

$$Z^\eta_\Lambda := \int \prod_{\substack{u \sim v\\ \{u,v\}\cap\Lambda\neq\emptyset}} h(\varphi_u, \varphi_v) \prod_{v\in\Lambda} d\lambda(\varphi_v) \prod_{v\in\mathbb{Z}^d\setminus\Lambda} d\delta_{\eta_v}(\varphi_v)$$

is a normalizing constant. We call $\mathbb{P}^\eta_\Lambda$ a *finite-volume Gibbs measure in* Λ *with boundary condition* η. We will also use the same notation with $\eta : \Lambda^c \to S$ (as $\mathbb{P}^\eta_\Lambda$ clearly depends only on η_{Λ^c}). We assume throughout that the interaction function h and single-site measure λ satisfy suitable integrability conditions to ensure that for any finite $\Lambda \subseteq \mathbb{Z}^d$ and $\eta : \mathbb{Z}^d \to S$,

$$Z^\eta_\Lambda < \infty$$

so that $\mathbb{P}^\eta_\Lambda$ is well defined, and

$$\mathbb{P}^\eta_\Lambda(f) \text{ is a continuous function of } \eta, \text{ for any bounded, continuous function } f. \tag{3}$$

The following lemma describes two fundamental properties of Gibbs measures.

Lemma 1.1. *Let $\Delta, \Lambda \subseteq \mathbb{Z}^d$ be finite sets satisfying $\Delta \subseteq \Lambda$ and let $\eta : \mathbb{Z}^d \to S$.*

(i) *(Domain Markov property). The measure $\mathbb{P}^\eta_\Lambda$ depends on η only through $\eta_{\partial_{ext}\Lambda}$, where $\partial_{ext}\Lambda := \{v \in \mathbb{Z}^d : v \notin \Lambda$ and there exists $u \in \Lambda$ adjacent to $v\}$.*

(ii) *(Gibbs property). Suppose φ is sampled from $\mathbb{P}^\eta_\Lambda$. Then the distribution of φ conditioned on φ_{Δ^c} (in the sense of regular conditional probabilities) equals $\mathbb{P}^{\varphi_{\Delta^c}}_\Delta$, almost surely.*

Both claims follow directly from the definition (2) of $\mathbb{P}^\eta_\Lambda$.

Examples. We list a few examples of the above setup.

(i) **Ising and Potts models**. In the Ising model $S = \{-1, 1\}$ with the discrete topology and uniform probability measure and, for a given inverse temperature $\beta \in \mathbb{R}$ (negative β yields the anti-ferromagnetic Ising model),

$$h(a, b) = \exp(\beta ab).$$

More generally, given an integer $q \geqslant 2$, the Potts model is defined by taking $S = \{1, 2, \dots, q\}$ with the discrete topology and uniform probability measure and, for a given inverse temperature $\beta \in \mathbb{R}$,

$$h(a, b) = \exp(-\beta\delta_{ab}),$$

where δ_{ab} is the Kronecker delta function.

(ii) **Spin O(n) model**. For a given integer $n \geqslant 1$, the spin $O(n)$ model has $S = \mathcal{S}^{n-1}$, the sphere of dimension $n - 1$, with the topology inherited from the embedding $\mathcal{S}^{n-1} \subseteq \mathbb{R}^n$ and the uniform (i.e., rotationally invariant) probability measure. For a given inverse temperature $\beta \in \mathbb{R}$,

$$h(a, b) = \exp(\beta\langle a, b\rangle),$$

where $\langle\cdot,\cdot\rangle$ denotes the standard inner product in $\mathbb{R}^n$. The case $n = 1$ is exactly the Ising model. The case $n = 2$ is called the XY model, or plane rotator model, and the case $n = 3$ is called the Heisenberg model.

(iii) **Discrete Gaussian free field**. Here, $S = \mathbb{R}$ with its standard topology and Lebesgue measure and, for a given inverse temperature $\beta > 0$,

$$h(a, b) = \exp(-\beta(a - b)^2).$$

Note that here, unlike the previous examples, the space S is not compact.

(iv) **Random-cluster model**. Given a finite graph $G = (V(G), E(G))$ and reals $0 < p < 1$, $q > 0$, the random-cluster (or FK) model is the probability measure $\mathbb{P}_{G,p,q}$ over subsets of $E(G)$ given by

$$\mathbb{P}_{G,p,q}(\omega) := \frac{1}{Z_{G,p,q}} p^{|\omega|}(1-p)^{|E(G)\setminus\omega|} q^{\mathcal{C}(\omega)}, \quad \omega \subseteq E(G),$$

where $\mathcal{C}(\omega)$ is the number of connected components in the graph $(V(G), \omega)$ and $Z_{G,p,q}$ normalizes $\mathbb{P}_{G,p,q}$ to be a probability measure. The random-cluster model does *not* fall under the framework of the above discussion as it is not a nearest-neighbor model (for $q \neq 1$, due to the factor $q^{\mathcal{C}(\omega)}$), but is included in a more comprehensive discussion of infinite-volume Gibbs measures, see, e.g., Friedli and Velenik [3, Chapter 6].

Infinite-volume Gibbs measures. We proceed to define infinite-volume Gibbs measures following the Dobrushin, Lanford, Ruelle (DLR) formalism.

Definition. A probability measure $\mathbb{P}$ on Ω is called a *Gibbs measure* (corresponding to the finite-volume Gibbs measures given in (2)) if it satisfies the following condition:
Let φ be sampled from $\mathbb{P}$. For each finite $\Lambda \subseteq \mathbb{Z}^d$, the conditional distribution of φ given φ_{Λ^c} equals $\mathbb{P}_\Lambda^{\varphi_{\Lambda^c}}$.
Equivalently but in more formal terms: a regular conditional probability for $\mathbb{P}$ conditioned on the sigma-algebra generated by φ_{Λ^c} is given by the mapping $(\varphi, E) \mapsto \mathbb{P}_\Lambda^{\varphi_{\Lambda^c}}(E)$.

Does there exist a Gibbs measure? Is it unique? It is clear that the set of Gibbs measures is convex (a mixture of Gibbs measures is again a Gibbs measure). Hence, there can either be zero, one, or infinitely many Gibbs measures. We first consider the existence question.

Lemma 1.2. *Let (Λ_n) be a sequence of finite domains which increases to $\mathbb{Z}^d$ (i.e., $\Lambda_n \subseteq \Lambda_{n+1}$ and $\cup\Lambda_n = \mathbb{Z}^d$) and let (η_n) be a sequence of functions $\eta_n : \Lambda_n^c \to S$.*

(i) *If the sequence of measures $\mathbb{P}_{\Lambda_n}^{\eta_n}$ converges then its limit is a Gibbs measure. (such a limiting procedure is called a* thermodynamic limit*, or an* infinite-volume limit*).*
(ii) *If S is compact then there exists a subsequence $n_k \uparrow \infty$ such that $\mathbb{P}_{\Lambda_{n_k}}^{\eta_{n_k}}$ converges.*

Proof. The second part follows immediately from compactness of the set of probability measures on Ω. We proceed to prove the first part. Denote by $\mathbb{P}$ the limit of $\mathbb{P}_{\Lambda_n}^{\eta_n}$. It suffices to show that for any finite $\Lambda \subseteq \mathbb{Z}^d$ and any bounded, continuous function $f : \Omega \to \mathbb{R}$,

$$\mathbb{P}(f \,|\, \varphi_{\Lambda^c}) = \mathbb{P}_\Lambda^{\varphi_{\Lambda^c}}(f). \tag{4}$$

Fix such a Λ and f. The equality (4) is equivalent to

$$\mathbb{P}(g \cdot \mathbb{P}(f \,|\, \varphi_{\Lambda^c})) = \mathbb{P}(g \cdot \mathbb{P}_\Lambda^{\varphi_{\Lambda^c}}(f)) \tag{5}$$

for any bounded, continuous function $g : \Omega \to \mathbb{R}$ determined by Λ^c (i.e., satisfying $g(\varphi) = g(\varphi')$ whenever $\varphi_{\Lambda^c} = \varphi'_{\Lambda^c}$). We shall obtain the equality (5) by developing both its sides. First, the left-hand side of (5) satisfies

$$\mathbb{P}(g \cdot \mathbb{P}(f \,|\, \varphi_{\Lambda^c})) = \mathbb{P}(g \cdot f) \tag{6}$$

by the properties of conditional expectation (as g is determined by Λ^c). Second, the right-hand side of (5) satisfies

$$\mathbb{P}(g \cdot \mathbb{P}_\Lambda^{\varphi_{\Lambda^c}}(f)) = \lim_{n\to\infty} \mathbb{P}_{\Lambda_n}^{\eta_n}(g \cdot \mathbb{P}_\Lambda^{\varphi_{\Lambda^c}}(f)) = \lim_{n\to\infty} \mathbb{P}_{\Lambda_n}^{\eta_n}(g \cdot \mathbb{P}_{\Lambda_n}^{\eta_n}(f \,|\, \varphi_{\Lambda^c})) = \lim_{n\to\infty} \mathbb{P}_{\Lambda_n}^{\eta_n}(g \cdot f) = \mathbb{P}(g \cdot f), \tag{7}$$

where the first equality follows since $\mathbb{P}_\Lambda^{\varphi_{\Lambda^c}}(f)$ is a continuous function of φ_{Λ^c} by (3), the second equality follows from the Gibbs property of $\mathbb{P}_{\Lambda_n}^{\eta_n}$ given in Lemma 1.1, the third equality follows from the properties of conditional expectation and the last equality follows since f and g are continuous. Putting together (6) and (7) finishes the proof of (5). □

Thus in the case that S is compact there is always at least one Gibbs measure. In the non-compact case it may be that no Gibbs measures exist and this is in fact the case for the discrete Gaussian free field in dimensions $d \in \{1, 2\}$. The low-temperature Ising model admits two different Gibbs measures (and hence infinitely many Gibbs measures, by taking mixtures of the two), one obtained as a thermodynamic limit when η_n is the constant $+1$ configuration and the other obtained when η_n is the constant -1 configuration.

Extremal Gibbs measures. The convexity of the set of Gibbs measures gives it additional structure: We say that a Gibbs measure $\mathbb{P}$ is *extremal* if it holds that when $\mathbb{P} = \alpha\mathbb{P}_1 + (1-\alpha)\mathbb{P}_2$ for Gibbs measures $\mathbb{P}_1, \mathbb{P}_2$ and $0 < \alpha < 1$ then necessarily $\mathbb{P} = \mathbb{P}_1 = \mathbb{P}_2$.

As will be discussed, extremal Gibbs measures are the building blocks for all Gibbs measures and it is thus desirable to obtain additional properties of them. To this end let us define the *tail σ-algebra* $\mathcal{T}$ as follows,

$$\mathcal{T} := \cap_{\Lambda\subseteq\mathbb{Z}^d \text{ finite}}\sigma(\{\varphi_v : v \notin \Lambda\}).$$

The following characterization is fundamental.

Lemma 1.3. *Let $\mathbb{P}$ be a Gibbs measure. Then $\mathbb{P}$ is extremal if and only if $\mathbb{P}(A) \in \{0,1\}$ for all $A \in \mathcal{T}$ (that is, $\mathbb{P}$ is tail-trivial).*

We require two preliminary claims.

Claim 1.4. *If $\mathbb{P}$ is a Gibbs measure and $A \in \mathcal{T}$ has $\mathbb{P}(A) > 0$ then $\frac{1_A}{\mathbb{P}(A)}\mathbb{P}$ is also a Gibbs measure.*

Proof. Denote $\mathbb{Q} := \frac{1_A}{\mathbb{P}(A)}\mathbb{P}$. Let $\Lambda \subseteq \mathbb{Z}^d$ be finite and let $E \subseteq \Omega$ be measurable. We need to show that

$$\mathbb{Q}(E \,|\, \varphi_{\Lambda^c}) = \mathbb{P}_\Lambda^{\varphi_{\Lambda^c}}(E), \quad \mathbb{Q}\text{-almost surely,}$$

or equivalently, that for each event $E_{\Lambda^c} \subseteq \Omega$ which is measurable with respect to $\sigma(\varphi_{\Lambda^c})$,

$$\mathbb{Q}(1_{E_{\Lambda^c}} \cdot \mathbb{P}_\Lambda^{\varphi_{\Lambda^c}}(E)) = \mathbb{Q}(E_{\Lambda^c} \cap E).$$

The last equality can be verified by noting that

$$\begin{aligned}\mathbb{Q}(1_{E_{\Lambda^c}} \cdot \mathbb{P}_\Lambda^{\varphi_{\Lambda^c}}(E)) &= \mathbb{P}\left(\frac{1_A}{\mathbb{P}(A)} \cdot 1_{E_{\Lambda^c}} \cdot \mathbb{P}_\Lambda^{\varphi_{\Lambda^c}}(E)\right) \\ &= \frac{1}{\mathbb{P}(A)}\mathbb{P}(1_{A\cap E_{\Lambda^c}} \cdot \mathbb{P}(E \,|\, \varphi_{\Lambda^c})) = \frac{1}{\mathbb{P}(A)}\mathbb{P}(\mathbb{P}(A \cap E_{\Lambda^c} \cap E \,|\, \varphi_{\Lambda^c})) \\ &= \frac{1}{\mathbb{P}(A)}\mathbb{P}(A \cap E_{\Lambda^c} \cap E) = \mathbb{Q}(E_{\Lambda^c} \cap E),\end{aligned}$$

where the first equality follows from the definition of $\mathbb{Q}$, the second equality follows as $\mathbb{P}$ is a Gibbs measure, the third equality follows as $A \in \mathcal{T} \subseteq \sigma(\varphi_{\Lambda^c})$ and $E_{\Lambda^c} \in \sigma(\varphi_{\Lambda^c})$, the fourth equality follows from the properties of conditional expectation and the last equality follows again from the definition of $\mathbb{Q}$. □

Claim 1.5. *If $\mathbb{P}_1, \mathbb{P}_2$ are Gibbs measures satisfying that $\mathbb{P}_1(A) = \mathbb{P}_2(A)$ for every $A \in \mathcal{T}$ then $\mathbb{P}_1 = \mathbb{P}_2$.*

Proof. Let $\mathbb{P}_1, \mathbb{P}_2$ be distinct Gibbs measures and let $E \subseteq \Omega$ be a measurable set for which $\mathbb{P}_1(E) \neq \mathbb{P}_2(E)$. Let (Λ_n) be an increasing sequence of finite subsets of $\mathbb{Z}^d$ with $\cup\Lambda_n = \mathbb{Z}^d$. Define a random variable $X : \Omega \to \mathbb{R}$ by

$$X(\varphi) := \limsup_{n\to\infty} \mathbb{P}_{\Lambda_n}^{\varphi_{\Lambda_n^c}}(E).$$

As X is measurable with respect to $\mathcal{T}$ it suffices to show that $\mathbb{P}_1(X) \neq \mathbb{P}_2(X)$.

For any Gibbs measure $\mathbb{P}$ and any finite $\Lambda \subseteq \mathbb{Z}^d$,

$$\mathbb{P}_\Lambda^{\varphi_{\Lambda^c}}(E) = \mathbb{P}(E \,|\, \varphi_{\Lambda^c}), \quad \mathbb{P}\text{-almost surely.}$$

Thus, Lévy's downward theorem shows that

$$\begin{aligned}X &= \lim_{n\to\infty} \mathbb{P}_1(E \,|\, \varphi_{\Lambda_n^c}) = \mathbb{P}_1(E \,|\, \mathcal{T}), \quad \mathbb{P}_1\text{-almost surely,} \\ X &= \lim_{n\to\infty} \mathbb{P}_2(E \,|\, \varphi_{\Lambda_n^c}) = \mathbb{P}_2(E \,|\, \mathcal{T}), \quad \mathbb{P}_2\text{-almost surely.}\end{aligned}$$

We conclude that $\mathbb{P}_1(X) = \mathbb{P}_1(E) \neq \mathbb{P}_2(E) = \mathbb{P}_2(X)$ finishing the proof of the claim. □

Proof of Lemma 1.3. First, suppose that there is some $A \in \mathcal{T}$ for which $0 < \mathbb{P}(A) < 1$. Clearly,

$$\mathbb{P} = \mathbb{P}(A) \cdot \frac{1_A}{\mathbb{P}(A)}\mathbb{P} + (1 - \mathbb{P}(A)) \cdot \frac{1_{A^c}}{\mathbb{P}(A^c)}\mathbb{P}$$

so that $\mathbb{P}$ is non-extremal, as both $\frac{1_A}{\mathbb{P}(A)}\mathbb{P}$ and $\frac{1_{A^c}}{\mathbb{P}(A^c)}\mathbb{P}$ are Gibbs measures by Claim 1.4.

Second, suppose that $\mathbb{P}$ is non-extremal, so that there exist distinct Gibbs measures $\mathbb{P}_1, \mathbb{P}_2$ and $0 < \alpha < 1$ with $\mathbb{P} = \alpha\mathbb{P}_1 + (1-\alpha)\mathbb{P}_2$. As $\mathbb{P}_1 \neq \mathbb{P}_2$ it follows from Claim 1.5 that there exists an event $A \in \mathcal{T}$ for which $\mathbb{P}_1(A) \neq \mathbb{P}_2(A)$. Thus, $\mathbb{P}(A) = \alpha\mathbb{P}_1(A) + (1-\alpha)\mathbb{P}_2(A) \notin \{0,1\}$. □

Extremal decomposition. For any Gibbs measure there *exists* a *unique* way to express the measure as an 'average' of extremal Gibbs measures. This is a consequence of abstract principles such as the Krein-Milman theorem and Choquet's theorem. It may also be seen directly, by starting with a Gibbs measure $\mathbb{P}$ and considering the regular conditional probability obtained by conditioning $\mathbb{P}$ on the tail sigma-algebra $\mathcal{T}$. One checks that $\mathbb{P}$-almost surely, the resulting conditional probability distribution is an extremal Gibbs measure and this then gives the required decomposition. We do not elaborate further on this decomposition here and refer to Friedli and Velenik [3, Section 6.8.4].

Translation-invariant and ergodic Gibbs measures. A probability measure $\mathbb{P}$ is called translation-invariant if

$$\mathbb{P}(A) = \mathbb{P}(\theta_v A) \quad \text{for each measurable set } A \subseteq \Omega \text{ and } v \in \mathbb{Z}^d,$$

where, for a configuration $\varphi \in \Omega$, $\theta_v\varphi$ is the shifted configuration defined by $(\theta_v\varphi)_w = \varphi(w-v)$, and $\theta_v A = \{\theta_v\varphi : \varphi \in A\}$. The set of translation-invariant Gibbs measures is often simpler to study and has special relevance to the physics of the model. The set of translation-invariant Gibbs measures is itself convex and hence it is natural to try and characterize its extremal elements. Here, we mean that a translation-invariant Gibbs measure $\mathbb{P}$ is extremal within translation-invariant Gibbs measures if when $\mathbb{P} = \alpha\mathbb{P}_1 + (1-\alpha)\mathbb{P}_2$ for translation-invariant Gibbs measures $\mathbb{P}_1, \mathbb{P}_2$ and $0 < \alpha < 1$ then necessarily $\mathbb{P} = \mathbb{P}_1 = \mathbb{P}_2$. Of course, if $\mathbb{P}$ is extremal (for all Gibbs measures) then it is also extremal within translation-invariant Gibbs measures. However, the converse need not hold.

The following lemma gives a useful characterization of extremality within translation-invariant Gibbs measures. A measurable set $A \subseteq \Omega$ is called *translation-invariant* if $\theta_v A = A$ for all $v \in \mathbb{Z}^d$. A translation-invariant probability measure $\mathbb{P}$ on Ω is called *ergodic* if $\mathbb{P}(A) \in \{0,1\}$ for all translation-invariant $A \subseteq \Omega$.

Lemma 1.6. *A translation-invariant Gibbs measure $\mathbb{P}$ is extremal within translation-invariant Gibbs measures if and only if it is ergodic.*

The proof of the lemma is left as an exercise to the reader.

References

[1] Dudley, Richard M. Real analysis and probability. Vol. 74. Cambridge University Press, 2002.
[2] Ethier, Stewart N., and Thomas G. Kurtz. Markov processes: characterization and convergence. Vol. 282. John Wiley & Sons, 2005.
[3] Friedli, Sacha, and Yvan Velenik. "Equilibrium statistical mechanics of classical lattice systems: A concrete introduction." preparation, available at http://www. unige. ch/math/folks/velenik/smbook (2017).