On sigma-finite measures related to the Martin boundary of recurrent Markov chains

Joseph Najnudel

Université Paul Sabatier, Toulouse III

August 7, 2015

The Brownian setting

The initial setting we studied with Roynette, Vallois and Yor is the following.

- We consider the Wiener measure $\mathbb{W}$ on the space $\mathcal{C}(\mathbb{R}_+, \mathbb{R})$.
- We let $(\Gamma_t)_{t\geq 0}$ be measurable nonnegative random variables such that

$$0 < \mathbb{E}_{\mathbb{W}}[\Gamma_t] < \infty.$$

- We define probability measures $(\mathbb{Q}_t)_{t\geq 0}$ by

$$\mathbb{Q}_t := \frac{\Gamma_t}{\mathbb{E}_{\mathbb{W}}[\Gamma_t]} \cdot \mathbb{W}.$$

- The setting considered here is similar to what is done in statistical physics, where the weight Γ_t is replaced by $e^{-H/T}$, H being the Hamiltonian (i.e. the energy of the configuration), and T the temperature.
- In our setting, a natural question is the following: does there exists a measure $\mathbb{Q}_\infty$ such that $\mathbb{Q}_t$ tends to $\mathbb{Q}_\infty$ when t goes to infinity?
- Of course, the answer to the question depends on the choice of $(\Gamma_t)_{t\geq 0}$ and the precise notion of convergence which is considered. The definition we take is the following: for all $s \geq 0$ and $\Lambda_s \in \mathcal{F}_s$,

$$\mathbb{Q}_t(\Lambda_s) \underset{t\to\infty}{\longrightarrow} \mathbb{Q}_\infty(\Lambda_s).$$

Here are some easy examples.

- If $\Gamma_t = e^{\lambda X_t}$ for $\lambda \in \mathbb{R}$, the density of $\mathbb{Q}_t$ with respect to $\mathbb{W}$ is $e^{\lambda X_t - t\lambda^2/2}$ and then $\mathbb{Q}_\infty$ exists and is the law of the Brownian motion with drift λ.
- If $\Gamma_t = 1 + X_t^2$, the measure $\mathbb{Q}_\infty$ exists and is equal to $\mathbb{W}$. Indeed, for $\Lambda_s \in \mathcal{F}_s$ and $t \geq s$,

$$\mathbb{Q}_t(\Lambda_s) = \frac{\mathbb{E}[\mathbb{1}_{\Lambda_s}(1 + X_s^2 + t - s)]}{1+t} \underset{t\to\infty}{\longrightarrow} \mathbb{W}(\Lambda_s).$$

- If $\Gamma_t = e^{tX_t}$, $\mathbb{Q}_\infty$ does not converge, since for $t \geq s$, X_s is, under $\mathbb{Q}_t$, a gaussian variable with mean st and variance s.

With Roynette, Vallois and Yor, we have studied many examples, and in all the examples we were interested in, $\mathbb{Q}_\infty$ exists. Here are some of these examples:

- For a function f from $\mathbb{R}$ to $\mathbb{R}$, integrable with respect to the measure $(1+|x|)dx$,
$$\Gamma_t := \exp\left(\int_0^\infty f(X_t)dX_t\right).$$
- For $\lambda \in \mathbb{R}$,
$$\Gamma_t := e^{\lambda L_t^0},$$
where L_t^0 is the local time at time t and level 0 of the canonical trajectory X.
- For a strictly positive, integrable function φ from $\mathbb{R}_+$ to $\mathbb{R}_+$, $\Gamma_t := \varphi(S_t)$, where S_t denotes the supremum of X up to time t.

A more difficult example has the interest that it is related to the Edwards polymer model in statistical physics. It is the following:

$$\Gamma_t = \exp\left(-\frac{1}{T}\int_0^t \delta(X_t - X_s)dsdt\right) := \exp\left(-\frac{1}{T}\int_{-\infty}^{\infty}(L_t^y)^2 dy\right).$$

- For this example, I have proven (the proof is a full article) that the limiting measure $\mathbb{Q}_\infty$ exists
- However, contrarily to the previous examples, I am not able to describe the properties of the trajectory under this measure.
- I conjecture a ballistic behavior, i.e. $|X_t|/t$ converges to a constant (depending only on T), with gaussian fluctuations. Such a behavior has been proven for the value of X_t under $\mathbb{Q}_t$, by van der Hofstad, den Hollander and König.

- A question we asked is why convergence occurs for many different examples. We have partially answered to the question in our monograph with Roynette and Yor.
- In part of the examples (not all of them) we considered, the limiting measure $\mathbb{Q}_\infty$ is absolutely continuous with respect to a common σ-finite measure $\mathcal{W}$ on the space $\mathcal{C}(\mathbb{R}_+, \mathbb{R})$. This σ-finite measure can be characterized in several ways.
- The measure $\mathcal{W}$ can be decomposed into a sum of two measures $\mathcal{W}_+$ and $\mathcal{W}_-$, the latter being obtained from the former simply by replacing the canonical trajectory by its opposite. Let us then focus on $\mathcal{W}_+$.

The measure $\mathcal{W}_+$ is the unique σ-finite measure on $\mathcal{C}(\mathbb{R}_+, \mathbb{R})$ satisfying the following properties:

- Almost every trajectory tends to $+\infty$ at infinity.
- For all $a \in \mathbb{R}$, $s \geq 0$, $\Lambda_s \in \mathcal{F}_s$, we have the following relation between $\mathbb{W}$ and $\mathcal{W}_+$:

$$\mathcal{W}(\Lambda_s, g_a \leq s) = \mathbb{E}_{\mathbb{W}}[\mathbb{1}_{\Lambda_s}(X_s - a)_+],$$

where g_a denotes the last hitting time of a by the canonical trajectory.

There are several ways to describe the canonical trajectory under $\mathcal{W}_+$.

- For $\ell \geq 0$, let $\mathbb{P}_\ell$ be the law of the concatenation of a Brownian motion stopped at its inverse local time at local time ℓ and level zero, and an independent Bessel process of dimension 3. Then for any measurable event Λ,
$$\mathcal{W}_+(\Lambda) = \int_0^\infty \mathbb{P}_\ell(\Lambda) d\ell.$$
- For $a \geq 0$, let $\mathbb{P}'_a$ be the law of the concatenation of a Brownian motion stopped at its first hitting time of $-a$, and an independent BES(3) process shifted by $-a$:
$$\mathcal{W}_+(\Lambda) = \int_0^\infty \mathbb{P}'_a(\Lambda) da.$$

- For $t \geq 0$, let $\mathbb{P}''_t$ be the law of the concatenation of a Brownian bridge of length t and a BES(3) process:

$$\mathcal{W}_+(\Lambda) = \int_0^\infty \mathbb{P}''_t(\Lambda) \frac{dt}{\sqrt{2\pi t}}.$$

One has also the following invariance property related to the Wiener measure: for all $t \geq 0$, if $\mathbb{W}_t$ denotes the law of a Brownian motion on $[0, t]$, then $\mathcal{W}_+$ is the image of $\mathbb{W}_t \otimes \mathcal{W}_+$ by the operation of concatenation of the trajectories. The only σ-finite measures we know they satisfy this property are $\mathbb{W}$, $\mathcal{W}_+$, $\mathcal{W}_-$ and their linear combinations (question: are there others such measures?)

The construction of $\mathcal{W}_+$ has been generalized in different settings in our monograph: the two-dimensional Brownian motion, more general linear diffusions on $\mathbb{R}$, and recurrent Markov chains satisfying some general assumptions stated below.
In the sequel of the talk, we will consider the setting of discrete Markov chains.

The setting of Markov chains

From now, we consider the following setting:

- We take a countable set E, and the canonical process $(X_n)_{n\geq 0}$ on $E^{\mathbb{N}}$.
- We take the family $(\mathbb{P}_x)_{x\in E}$ of probability measures corresponding to a Markov chain on E.
- The Markov chain is assumed to be irreducible and recurrent.
- We suppose that for all $x \in E$, the transition probability $p_{x,y}$ vanishes for all but finitely many elements $y \in E$.

We then fix a point $x_0 \in E$, and a function φ from E to $\mathbb{R}_+$, not identically zero, and satisfying the following properties:

- φ vanishes at x_0,
- φ is harmonic at every point $x \neq x_0$, i.e.

$$\mathbb{E}_x[\varphi(X_1)] = \varphi(x).$$

Then, for $r \in (0,1)$, we define ψ_r by

$$\psi_r(x) = \varphi(x) + \frac{r}{r-1}\mathbb{E}_{x_0}[\varphi(X_1)].$$

The function ψ_r is harmonic everywhere except at x_0, and one has

$$\psi_r(x_0) = r\mathbb{E}_x[\psi_r(X_1)].$$

From this relation, we deduce that $\left(\psi_r(X_n) r^{L_{n-1}^{x_0}}\right)_{n \geq 0}$ is a martingale under $\mathbb{P}_x$, where $L_{n-1}^{x_0}$ denotes the number of hitting times of x_0 by X between times 0 and $n-1$.
We then construct a measure $\mu_x^{(r)}$ whose density with respect to $\mathbb{P}_x$ is equal to this martingale at time n after restriction to $\mathcal{F}_n$, for all $n \geq 0$.
Under $\mu_x^{(r)}$, the total local time at x_0 is finite almost surely, and then we define:

$$\mathbb{Q}_x := r^{-L_\infty^{x_0}} \mu_x^{(r)}.$$

One can prove that $\mathbb{Q}_x$ does not depend on r.

The family of σ-finite measures $(\mathbb{Q}_x)_{x \in E}$ can be considered as an analog of $\mathcal{W}$ in the present setting. Note that the point x_0 is less important than one may believe, since the same measures can be recovered from any other point $y_0 \in E$ if the function φ is changed to the function $\varphi^{[y_0]}$ described below. Here are some properties similar to those of $\mathcal{W}$:

- Under $\mathbb{Q}_x$, almost every trajectory $(X_n)_{n \geq 0}$ is transient, i.e. it visits each element of E finitely many times.
- For all $y_0 \in E$, and $\Lambda_n \in \mathcal{F}_n$,

$$\mathbb{Q}_x[\Lambda_n, g_{y_0} < n] = \mathbb{E}_x[\mathbb{1}_{\Lambda_n} \varphi^{[y_0]}(X_n)],$$

  where g_{y_0} is the last hitting time of y_0 and $\varphi^{[y_0]}(y)$ is the total measure, under $\mathbb{Q}_y$, of all the trajectories which do not visit y_0. Note that $\varphi^{[x_0]} = \varphi_x$.

The measure $\mathbb{Q}_x$ can be decomposed in the following way:

$$\mathbb{Q}_x = \mathbb{Q}_x^{[y_0]} + \sum_{k \geq 1} \mathbb{P}_x^{\tau_k^{y_0}} \circ \widetilde{\mathbb{Q}}_{y_0},$$

where $\mathbb{Q}_x^{[y_0]}$ is the restriction of $\mathbb{Q}_x$ to the trajectories which do not hit y_0, $\widetilde{\mathbb{Q}}_{y_0}$ is the restriction of $\mathbb{Q}_{y_0}$ to the trajectories which do not return at y_0, $\mathbb{P}_x^{\tau_k^{y_0}}$ is the law of the initial Markov chain stopped at the k-th hitting time of y_0, and $\circ$ denotes the operation on the measures which corresponds to the concatenation of the trajectories.

One has also a similar invariance property between $(\mathbb{P}_x)_{x \in E}$ and $(\mathbb{Q}_x)_{x \in E}$ as in the Brownian setting. More precisely, the following holds

$$\mathbb{Q}_x = \sum_{y \in E} \mathbb{1}_{X_n = y} \cdot (\mathbb{P}_x^{(n)} \circ \mathbb{Q}_y)$$

where $\mathbb{P}_x^{(n)}$ is the law of the initial Markov chain starting at x and stopped at time n.
Note that the same result holds if we replace $\mathbb{Q}_x$ and $\mathbb{Q}_y$ by $\mathbb{P}_x$ and $\mathbb{P}_y$. We don't know if there are families of measures satisfying the same properties, which are not linear combinations of $(\mathbb{P}_x)_{x \in E}$ and $(\mathbb{Q}_x)_{x \in E}$ for a suitable choice of x_0 and φ.

This property of invariance means, in some sense, that if we take finite parts of the trajectories, the measure $\mathbb{Q}_x$ starts like $\mathbb{P}_x$. Since the canonical trajectory is transient under $\mathbb{Q}_x$, one can interpret (very informally!) the measure $\mathbb{Q}_x$ as follows: it looks like the law of "a recurrent Markov chain, conditionned to be transient".

With this interpretation, it remains to see what is the role played by the function φ used to construct $\mathbb{Q}_x$. As we will see in the last section, φ is related to the way the trajectories go to infinity under $\mathbb{Q}_x$; more precisely, to the Martin boundary of the Markov chain considered at the beginning.

Link with the Martin boundary

For a recurrent Markov chain on E, one can define the so-called *Martin boundary* of E in the following way (this construction is essentially due to Kemeny and Snell, after a similar, more classical construction for transient chains by Doob and Hunt).

- For some fixed $x_0 \in E$, we define G_{x_0} as the Green function of the Markov chain stopped just before the first strictly positive hitting time T'_{x_0} of x_0:

$$G_{x_0}(x,y) = \mathbb{E}_x[L^y_{T'_{x_0}-1}].$$

- We then define the following function:

$$L_{x_0}(x,y) := \frac{G_{x_0}(x,y)}{G_{x_0}(x_0,y)}.$$

- The function L_{x_0} defines a distance on E, given by

$$\delta_{x_0,w}(x,y) = \sum_{z \in E} w_z \frac{|L_{x_0}(z,x) - L_{x_0}(z,y)| + |\mathbb{1}_{z=x} - \mathbb{1}_{z=y}|}{1 + \sup_{y \in E} L_{x_0}(z,y)},$$

  where $w = (w_z)_{z \in E}$ is a summable family of elements of $\mathbb{R}_+^*$.
- The completion $\overline{E}$ of $(E, \delta_{x_0,w})$ is called the *Martin compactification of E*, and its topological structure does not depend on the choice of x_0 and w.
- The boundary ∂E of $\overline{E}$ is called the *Martin boundary* of E.
- By continuity, one can define $L_{x_0}(x, \alpha)$ for all $\alpha \in \partial E$.

- For all $\alpha \in \partial E$, the function $\varphi_\alpha : x \mapsto L_{x_0}(x,\alpha)\mathbb{1}_{x \neq x_0}$ is nonnegative, and harmonic everywhere except at x_0.
- If φ_α is minimal among the functions satisfying these properties, i.e. any smaller such function is equal to $c\varphi_\alpha$ for some $c \in [0,1]$, then one says that α is in the *minimal Martin boundary* $\partial_m E$.
- One then has the following result: for any function φ from E to $\mathbb{R}_+$ which vanishes at x_0 and is harmonic everywhere else, there exists a finite measure μ_{φ,x_0} on $\partial_m E$, such that for all $x \neq x_0$,

$$\varphi(x) = \int_{\partial_m E} L_{x_0}(x,\alpha) d\mu_{\varphi,x_0}(\alpha).$$

This property implies that one can completely characterize the families $(\mathbb{Q}_x)_{x \in E}$ in terms of the minimial Martin boundary $\partial_m E$.

- We consider a stationnary nonnegative measure $(\beta(y))_{y \in E}$ for the Markov chain. It is uniquely determined up to a multiplicative constant.
- For $x_0 \in E$, $\alpha \in \partial_m E$, one can define a family of measures $(\mathbb{Q}_x^{(x_0,\alpha)})_{x \in E}$ related to the point x_0 and the function

$$\varphi(x) := \frac{L_{x_0}(x,\alpha)}{\beta(x_0)} \mathbb{1}_{x \neq x_0}.$$

- One can show that $(\mathbb{Q}_x^{(\alpha)} = \mathbb{Q}_x^{(x_0,\alpha)})_{x \in E}$ does not depend on x_0.

Now, we have the following result: if $(\mathbb{Q}_x)_{x\in E}$ is a family of measure constructed as before, then there exists a unique finite measure μ on $\partial_m E$ such that for all measurable events Λ, and all $x \in E$,

$$\mathbb{Q}_x(\Lambda) = \int_{\partial_m E} \mathbb{Q}_x^{(\alpha)}(\Lambda) d\mu(\alpha).$$

We then have the following result: for all $\alpha \in \partial_m E$, $x \in E$, almost every trajectory tends to α (for the topology of $\overline{E}$) at infinity, under the measure $\mathbb{Q}_x^{\alpha}$. Informally, the family $(\mathbb{Q}_x^{\alpha})_{x\in E}$ corresponds to the "law of the initial Markov chain conditionned to tend to α at infinity".

Some examples

Let us consider the simple random walk on $\mathbb{Z}$. In this case, one has the following:

$$G_0(0,y) = L_0(0,y) = 1,$$

and for $x \neq 0$,

$$G_0(x,y) = L_0(x,y) = 2(|x| \wedge |y|)\mathbb{1}_{xy>0}.$$

One can deduce that the Martin boundary, and the minimal Martin boundary, have exactly two points, denoted $-\infty$ and $+\infty$, such that

$$L_0(x,+\infty) = 2x_+ + \mathbb{1}_{x=0},\ L_0(x,-\infty) = 2x_- + \mathbb{1}_{x=0}.$$

One then gets two families of σ-finite measures, $(\mathbb{Q}_x^{+\infty})_{x\in\mathbb{Z}}$ and $(\mathbb{Q}_x^{-\infty})_{x\in\mathbb{Z}}$.

- The measure $\mathbb{Q}_0^{+\infty}$ is the sum, for $k \geq 1$, of the law of a simple random walk stopped at the k-th hitting time of zero, followed by an independent random walk on $\mathbb{Z}_+$ with transitions $p_{x,x+1} = \frac{x+1}{2x}$, $p_{x,x-1} = \frac{x-1}{2x}$ (this last random walk is the discrete analog of the BES(3) process).
- The measure $\mathbb{Q}_x^{+\infty}$ is obtained from $\mathbb{Q}_0^{+\infty}$ by translating the trajectories by x.
- The measure $\mathbb{Q}_{-x}^{-\infty}$ is obtained by changing the trajectories under $\mathbb{Q}_x^{+\infty}$ to their opposite.
- The canonical trajectory tends to $+\infty$ under $\mathbb{Q}_x^{+\infty}$ and to $-\infty$ under $\mathbb{Q}_x^{-\infty}$.

Let us now consider the simple random walk on $\mathbb{Z}^2$. In this case, one can prove that there exists a unique nonnegative function φ which is harmonic at every non-zero points, and such that $\varphi(0,0) = 0$, $\varphi(0,1) = 1$. It has the same symmetry as the lattice $\mathbb{Z}^2$, and for all $n \geq 1$,

$$\varphi(n,n) = \frac{4}{\pi} \sum_{j=1}^{n} \frac{1}{2j-1}.$$

One can directly compute each value of φ by using the previous properties: in particular it is always in $\mathbb{Q} + \mathbb{Q}/\pi$. One has the asymptotic expansion at infinity:

$$\varphi(x) = \frac{2}{\pi} \log(||x||) + \frac{2\gamma_{\text{Euler}} + \log 8}{\pi} + O(1/||x||^2).$$

Hence, we can construct only a unique family of measures $(\mathbb{Q}_x)_{x\in\mathbb{Z}^2}$, up to a multiplicative constant, and the Martin boundary of $\mathbb{Z}^2$ has a unique point ∞ for the simple random walk.
Under $\mathbb{Q}_x$, the trajectories tend to ∞ for the topology of the Martin compactification of $\mathbb{Z}^2$, which means that their norm tends to infinity in the usual sense.

Let us now consider a random walk on an infinite binary tree, with transitions probabilities $1/2$, $1/2$ from the root to each of its children, $1/4$, $1/4$ from each other vertex to each of their children, $1/2$ from each vertex (except the root) to its father. Note that with these transitions, the distance to the root is the absolute value of a simple random walk on $\mathbb{Z}$, and then the Markov chain is recurrent.

In this case, one can prove that the Martin boundary is equal to the minimal Martin boundary, and is uncountable: it is indexed by the leafs of the tree, i.e. by the infinite simple paths starting from the root. A function φ corresponding the a given leaf λ takes the value $2^p - 1$ at a vertex x, where p denotes the number of common edges in the simple path from the root to x and the simple path from the root to λ.

Thank you for your attention!