STATIONARY MEASURES ON THE TORUS AFTER BENOIST-QUINT

ALEX WRIGHT

Abstract. We give an exposition of Benoist-Quint's "exponential drift" argument in the case of the two dimensional torus $\mathbb{T}^2$. These are the author's notes for three informal talks he gave at the University of Chicago in March 2013.

1. Statement of results, stationary measures, and backwards random walk

1.1. **Introduction.** The work of Benoist-Quint lies in the field of ergodic theory on homogenous spaces, which seeks among other things to understand orbit closures and invariant measures for actions of subgroups of Lie groups on homogeneous spaces. The most prominent success in direction is

Theorem 1.1 (Ratner's Theorem)**.** *Let G be a Lie group, Γ a lattice in G, and u_t a one parameter unipotent subgroup of G. Then*

1. *the u_t–orbit closure of every point in G/Γ is* homogeneous,
2. *every u_t–orbit in G/Γ equidistributes in its orbit closure,*
3. *every u_t–invariant ergodic probability measure on G/Γ is* homogeneous.

A closed subset $S \subset G/\Gamma$ is called *homogeneous* if there is some Lie subgroup $L \subset G$ whose image in G/Γ is S. It this case it is automatic that $\Gamma \cap L$ is a lattice in L. A probability measure on G/Γ is called *homogeneous* if it is the pushforward of Haar measure on a subgroup L of G for which $\Gamma \cap L$ is a lattice in L.

Idea of Ratner's Theorem Part 3. All such results are based on the idea of *additional invariance.* That is, one begins with a u_t–invariant ergodic probability measure ν on G/Γ, and show that ν is invariant under more and more one parameter subgroups of G. At the end of the argument, ν is invariant under a subgroup $L \subset G$ (generated by all these one parameter subgroups of G) and also supported on a single L orbit, and it follows that ν is homogeneous.

Let us give the idea of how this argument begins. Pick two very nearby points x, y which are in the support of ν but which are not in the same ν orbit. It is possible to do this unless ν is supported on a closed u_t orbit, in which case we are done.)

After some large time T we see that the orbits of x and y have finally drifted apart, and for the first time the distance between $u_T(x)$ and $u_T(y)$ is now 1. A key point is that the distance between $u_t(x)$ and $u_t(y)$ grows polynomially. That is, the orbits have undergone *polynomial drift*. This is important because that it guarantees that in fact the orbits spent a large amount of time almost at distance 1 before finally reaching distance 1 from one another. If the orbits had been separating at exponential speed, the time during which they are about distance one from each other would be at most $O(1)$. But in this polynomial drift situation, the time during which the orbits are about distance one from each other is in fact some fraction of T. That is, the orbits are within distance $1-\delta$ during times $t \in [0.99T, T]$.

Since T is very large, this is in fact a very large window of times. So large in fact that the segment $u_t(x), t \in [0.99T, T]$ equidistributes with respect to ν. Every point $u_t(x)$ in this interval has a "friend" $u_t(y)$ which is distance about 1 away. It is possible to assume that in fact $u_t(y)$ is distance 1 from $u_t(x)$ in the direction of some one parameter subgroup of G. In this way we prove that ν, which is very well approximated by the segment $u_t(x), t \in [0.99T, T]$, is invariant by translation by 1 in the direction of this one parameter subgroup. By replacing 1 with smaller numbers, we see that ν is invariant by this one parameter subgroup.

This argument has a number of common features.

- A *starting condition* must be satisfied to run the argument. Above we had to find appropriate x and y.
- The *direction* in which the drift is occurring must be understood. Above it was in the direction of the one parameter subgroup.
- The *speed* of drift must be understood, and slow drift is much easier than fast. Exponential drift cripples the argument above.

Although we have used equidistribution (of large chunks of the horocycle) in the above sketch, it is not required. Indeed, all that is required is for most points to have a "friend" in a certain direction. It seems to be common for additional invariance arguments to use some equidistribution statement as intuition, but actually use something weaker in the proof. The intuitive equidistribution statements are usually true, but sometimes only provable after the fact.

Idea of Ratner's Theorem Parts 1 and 2. These parts follow from part 3. If we wish to consider the u_t–orbit of x, we consider any weak-$*$ limit of Lebesgue probability measure on the segment $u_t(x), t \in [-T, T]$ as $T \to \infty$. In this way we are able to construct a u_t–invariant measure on any orbit closure. It is important that the group u_t which is acting is a one parameter subgroup isomorphic to $\mathbb{R}$, and $\mathbb{R}$ is amenable. If the acting group u_t is replaced by a non-amenable group, this averaging technique does not work, and it may not be possible to construct an invariant measure on each orbit closure.

The general rule is that to understand orbit closure of actions of amenable groups such as u_t it often suffices to understand invariant measures, but more work is required for non-amenable groups.

The setting of Benoist-Quint. Here we consider the work of Benoist-Quint only for the torus $X = \mathbb{T}^2 = \mathbb{R}^2/\mathbb{Z}^2$. The work of Benoist-Quint applies in much more general settings, but in this case we see almost all the ideas while avoiding some technical difficulties.

The group $SL(2,\mathbb{Z})$ acts on X linearly, preserving Lebesgue measure. A motivating problem is to understand orbit closures for the action of a subgroup Γ of $SL(2,\mathbb{Z})$ on X.

It is important to avoid bad situations. For example, if

$$\Gamma = \left\langle \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix} \right\rangle,$$

then we are considering a single hyperbolic automorphism of X. The dynamics of such an automorphism are as complicated as possible, and there are a great many closed invariant sets and invariant measures.

Therefor we ask that Γ not be *non-elementary.* By definition this means Γ does not contain a finite index cyclic subgroup. It is equivalent that Γ is Zariski dense in $SL(2,\mathbb{R})$.

Standing assumption: $\Gamma \subset SL(2,\mathbb{Z})$ is non-elementary.

This assumption will be in force for the remainder of this document. The answer to the motivating question is

Theorem 1.2. *All Γ–invariant closed subsets of X are finite or equal to X.*

But how are we to show this? Since Γ is non-amenable, we do not expect to be able to directly construct Γ invariant measures on orbit closures. Instead, we put a probability measure μ on Γ, which allows us to consider *random diffeomorphisms* of $\mathbb{T}^2$ chosen with law μ. It is not hard to show that every Γ–orbit closure admits a measure ν which is, on average, invariant under random elements of Γ. Benoist-Quint classifies such measure μ. We must of course assume that ν "sees" all of Γ, and is not supported on a proper subgroup of Γ.

$$\boxed{\textbf{Standing assumption: } \text{the support of } \mu \text{ generates } \Gamma.}$$

We will also assume that the support of μ is finite. If you'd like, you can think of

$$\mu = \frac{1}{2}\delta_A + \frac{1}{2}\delta_B, \quad \text{where} \quad A = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}.$$

Proposition 1.3 (Kakutani)**.** *Every orbit Γ–orbit closure supports a μ–stationary measure.*

Theorem 1.4 (Benoist-Quint)**.** *All μ–stationary measures ν on X are either finitely supported of Lebesgue measure on X.*

1.2. **Stationary measures.** The idea of a μ–stationary measure formalizes the idea of a measure which is on average invariant under random elements of Γ, when these random elements are chosen with law μ. Formally,

Definition 1.5. A *μ–stationary measure* is a measure ν for which $\mu * \nu = \nu$.

The condition $\mu * \nu = \nu$ is rewritten more explicitly as

$$\int_\Gamma g\nu d\mu(g) = \nu$$

where $(g\nu)(A) = \nu(g^{-1}(A))$. The condition that $g\nu = \nu$ is exactly invariance of ν, so $\mu * \nu = \nu$ really does say that ν is invariant on average.

The problem with the *definition* of a stationary measure is that it seems quite floppy; it leaves the possibility that any particular element of Γ can distort the stationary measure ν in an unconstrained way. This makes the work of Benoist-Quint all the more remarkable. It gives that these measures which might a priori be distorted a lot by Γ are in fact invariant by Γ, and for this reason the result is sometimes called *stiffness*.

Examples of stationary measures. In particular, a Γ–invariant measure is μ–stationary. However, there are many stationary measures which are not invariant in the situation where Γ is non-abelian and features lots of contracting elements.

Let us leave the case $\Gamma \subset SL(2,\mathbb{Z})$ for a just moment. A common example (from the theory of Lyapunov exponents) comes from the action of $SL(2,\mathbb{R})$ on $\mathbb{P}^1$. For simplicity let us consider μ a $SO(2)$–invariant absolutely continuous compactly supported measure on $SL(2,\mathbb{R})$. Then the $SO(2)$–invariant Lebesgue measure ν on $\mathbb{P}^1$ is μ–stationary. Although most elements of the support of μ will expand and contract $\mathbb{P}^1$ and thus distort the measure, because μ is $SO(2)$–invariant we see that $\mu * \nu$ is again $SO(2)$–invariant and hence must be equal to ν.

1.3. **The backwards random walk.** Stationary measures are ideally suited to study via random walks, because they are only on average invariant by a random element of Γ. Using a random walk setup will allow us to trade in the stationary measure ν for an *invariant* measure on a much more complicated dynamical system which takes into account all the randomness.

Despite the name, the random walk set up leads to a *non-random* dynamical system, in the same way that a Bernoulli shift is a non-random map modeling the random flipping of coins. In this non-random situation we will eventually be able to understand in which directions different parts of the measure should be invariant, and thus run the additional invariant argument.

The bilateral and forwards random walks. We could consider the space $\Gamma^{\mathbb{Z}} \times X$, with the sift map

$$(\ldots, g_{-1}, g_0, g_1, \ldots, x) \mapsto (\ldots, g_0, g_1, g_2, \ldots, g_1 x).$$

This system (the bilateral random walk) remembers the current position in $\mathbb{T}^2$ of the random walk (x), the past history $(\ldots, g_{-2}, g_{-1}, g_0)$ and the future moves $(g_1, g_2, \ldots)$ of the random walk. The invariant measure is easily seen to be $\mu^{\otimes\mathbb{Z}} \times \nu$, using the fact that ν is μ–stationary.

We could also consider the mild variant of this where the past is forgotten, thus obtaining a unilateral shift. In either case the map represents moving forward one step in time.

However, we will see that the strategy of Benoist-Quint requires a different set up.

The backwards random walk. This set up remembers the current position and past trajectory of the random walk only. That is, it remembers "where you are and how you got there." The map on this space is taking one step into the past. Moving into the future is not well defined.

The benefit of this setup is that, since we have not constrained the future, we are free to move forward into the future randomly and find equidistribution of the resulting points an increasing number of steps into the future.

The cost of the backwards setup is a more complicated invariant measure which is no longer a product.

The formal setup. Set $B = \Gamma^{\otimes \mathbb{Z}_{\geq 0}}$, and set $B^X = B \times X$. A point (b,x) of B^X should be thought of as giving the past history $b = (b_0, b_1, b_2, \ldots)$ of the random walk as well as the current position x in the torus. The unilateral shift is denoted

$$T : B \to B, \qquad T(b_0, b_1, b_2, \ldots) = (b_1, b_2, \ldots),$$

and the map which "takes one step backwards in time" is denoted

$$T^X : B^X \to B^X, \qquad T^X(b,x) = (Tb, b_0^{-1}x).$$

We will be required to understand measure ν_b on X which give the distribution of the point x on the torus given that a specific past b has occurred. But first a warm up.

Lemma 1.6. *The distribution of the final position $x \in X$ of the random walk, conditioned on the the last two steps of the random walk being b_0, b_1, is exactly $(b_0 b_1)_* \nu$.*

Explanation. The distribution of the end point $x \in X$ of the random walk is ν. Similarly this is the distribution for the position of the random walk at two steps into the past. That is, ν is also the distribution for where the point was two steps ago. This can be seen using that ν is μ stationary. Convolving μ by ν gives the distribution of a point one step into the future. But this is again ν, because ν is stationary. Thus the distribution for the position in the torus is ν at every stage of the random walk.

In particular, the distribution of the point y two steps before the random walk ends is ν. Now we condition for these final two steps to be b_1 and then b_0. So after the first step the distribution is $(b_1)_*\nu$, and after the second step it is $(b_0 b_1)_* \nu$. $\square$

There is a natural measure on the set B of possible past histories, given as $\beta = \mu^{\otimes \mathbb{Z}_{\geq 0}}$.

Proposition 1.7 (Furstenberg)**.** *Let ν be a μ–stationary probability measure on X. Then the limit*

$$\nu_b := \lim_{n \to \infty} (b_0 \cdots b_{n-1})_* \nu$$

exists for β almost every $b \in B$, and satisfies

$$\nu_b = (b_0)_* \nu_{Tb}.$$

Moreover, ν is the average of the ν_b:

$$\nu = \int_B \nu_b d\beta(b).$$

The existence of the limit follows directly from the Martingale Convergence Theorem. We repeat the interpretation of ν_b: it is the distribution for the final position in the torus of the random walk, conditioned on the fact that the past history is b.

In other situation we'd never think of conditioning on the entire past of the random walk, because often random walks diverge to infinity; it is relevant that the torus is compact.

The ν_b are the main objects of the proof. It is the ν_b on which we will run the additional invariance argument, concluding that almost every ν_b is Lebesgue. This makes a certain amount of sense: to understand which directions the measure sure be invariant under, one has to remember the maps b which have been applied. In particular, ν_b will be shown to be invariant by translations in some direction *depending on* b. But this direction will almost always be irrational, so almost every ν_b will be Lebesgue, so ν will be Lebesgue.

The construction of the invariant measure on B^X is stilled owed to the reader. But first let us pause for a diversion.

An aside on stationary measures for abelian group actions. No one ever talks about stationary measure for $\mathbb{Z}$. This is because of the following result, where we again temporarily suspend our usual notational assumptions.

Proposition 1.8 (Choquet, Deny)**.** *Let Γ be an abelian group acting on a space X. Let μ be a probability measure on Γ whose support generates Γ, and let ν be a μ–stationary probability measure on X. Then μ is Γ–invariant*

The same statement holds when Γ is nilpotent (Guivarc'h-Raugi) but not when Γ is solvable.

Lemma 1.9 (Hewitt-Savage zero-one law)**.** *Let $B = \Gamma^{\mathbb{Z}}_{\geq 0}$, where now we can consider Γ as any countable set. Let Σ be the group of permutations of $\mathbb{Z}_{\geq 0}$ which fix all but finitely many numbers. Then the action of Σ on B given by*

$$\sigma(b) = (b_{\sigma(0)}, b_{\sigma(1)}, b_{\sigma(2)}, \ldots)$$

is ergodic.

Proof of Proposition. We consider the measurable function $b \to \nu_b$ defined on B. Since

$$\nu_b = \lim_{n\to\infty} (b_0 \cdots b_{n-1})_* \nu$$

and Γ is commutative, we see that ν_b is constant on orbits of Σ. By ergodicity of the Σ action, ν_b is almost everywhere constant. Since $\nu = \int_B \nu_b d\beta(b)$, we get $\nu = \nu_b$ for almost every b. Since

$$\nu = \nu_b = (b_0)_* \nu_{Tb} = \nu$$

we get that ν is Γ invariant. □

The invariant measure. We know what the invariant measure on B is ($\beta = \mu^{\otimes \mathbb{Z}_{\geq 0}}$). The invariant measure on B^X is more complicated.

Lemma 1.10. *There is a unique measure β^X on $B^X = B \times X$ such that if ϕ is a function on $B \times X$ which only depends on the first n coordinates of b and the x, then*

$$\int_{B\times X} \phi(b,x) d\beta^X(b,x) = \int_{B\times X} \phi(b, b_0 \cdots b_{n-1} y) d\beta(b) d\nu(y).$$

The proof is immediate from the Caratheodory Theorem. Here you should think of y is the position in the torus of the random walk n steps into the past (whereas x is the position at present). Intuitively, if you have a function which only depends on the last n steps of the walk, you should go back in time n steps to y, and then average over all ways $(b_0, \ldots, b_{n-1})$ of going back to the present, and this will give you the integral. This is exactly what the above lemma says.

Lemma 1.11. *β^X is T^X invariant, and the pushforward of β^X to B under the projection is equal to β.*

Thus β^X is the desired invariant measure for the backwards random walk setup B^X.

2. Going backwards and forwards in time, and conditional measures on stable leaves

3. Time change and the horocyclic flow

Math Department, University of Chicago, 5734 South University Avenue, Chicago, IL 60637

E-mail address: alexmwright@gmail.com