Polymer Pinning with Sparse Disorder

Ken Alexander (joint with Quentin Berger)

Univ. of Southern California

June 2014

Usual polymer pinning model: $X = \{X_n\}$ a Markov chain interacting with a quenched random potential (reward/penalty) on the axis in spacetime—potential ω_n at $(n, 0)$, mean-0 i.i.d. r.v.'s. Hamiltonian and Gibbs measure

$$H_{N,\omega}(x) = \sum_{n=1}^{N} (\omega_n + h) \mathbf{1}_{x_n = 0}, \quad \mu_{N,\omega}^{\beta,h}(x) = \frac{1}{Z_{N,\omega}} e^{\beta H_{N,\omega}(x)} P(x).$$

Let $\tau = \{\tau_j\}$ be the return times—a renewal process. We really only need some renewal process τ, not the Markov chain. $\mathbf{1}_{x_n=0}$ becomes $\mathbf{1}_{n \in \tau}$. Assume power-law tails:

$$P(\tau_1 = n) = n^{-(1+\alpha)} \varphi(n), \quad \text{for some } \alpha \geq 0 \text{ and slowly var. } \varphi.$$

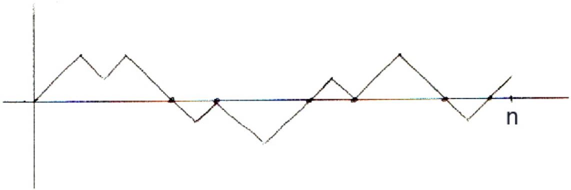

Homogeneous polymer: constant potential $\omega_n \equiv c$.
Annealed polymer: Take mean over ω of each Boltzmann weight, ω_n replaced by $\beta^{-1} \log M(\beta)$ (M=mgf.) Special case of homogeneous.
Free energy $F(\beta, h) = \lim_N \frac{1}{N} \log Z_{N,\omega}^{\beta,h}$.
Let L_n be the number of returns to 0. The polymer is **pinned** if for some $\delta > 0$,

$$\mu_n(L_n \geq \delta n) \to 1 \quad \text{as } n \to \infty.$$

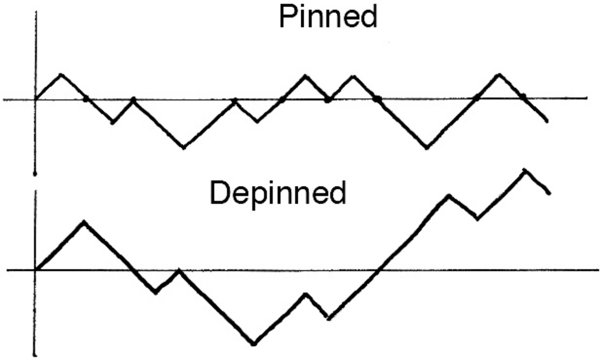

Contact fraction is $C(\beta, h)$ such that $\frac{|\tau|_N}{N} \to C$ in $\mu_{N,\omega}^{\beta,h}$-probability, where $|\tau|_N = |\tau \cap (0, N]|$.

Depinning transition

Critical value $h_c(\beta)$ ($= h_c^{qu}$ or h_c^{ann}) such that

$$h > h_c(\beta) \implies \text{ pinned: } F(\beta, h) > 0, \ C(\beta, h) > 0;$$

$$h < h_c(\beta) \implies \text{ depinned: } F(\beta, h) = 0, \ C(\beta, h) = 0;$$

Jensen's ineq. implies $h_c^{ann} \leq h_c^{qu}$ (quenched is harder to pin.)
Belief: inequality is strict if and only if the overlap is infinite, i.e. $\tau \cap \tau'$ is recurrent for τ' an independent copy. Overlap is infinite for $\alpha > 1/2$; depends on φ for $\alpha = 1/2$. Belief is "almost proved" except that some φ aren't covered for $\alpha = 1/2$. (Giacomin, Lacoin Toninelli 2009, A. 2008.)

Modified Model: Pinning-By-Renewals

Sparse disorder: $\sigma = \{\sigma_j\}$ another (quenched) renewal, reward

$$\omega_n = \mathbf{1}_{n\in\sigma}.$$

τ must hit sites $n \in \sigma$ to claim any reward. τ, σ have possibly different tail exponents $\alpha, \tilde{\alpha}$. Gap $W_j = \sigma_j - \sigma_{j-1}$. Disorder is truly "sparse" if $\tilde{\alpha} < 1$ which makes $E^\sigma(W_1) = \infty$. For $\tilde{\alpha} < 1$, typically

$$\sigma_N \asymp N^{1/\tilde{\alpha}}, \quad |\sigma|_N \asymp N^{\tilde{\alpha}},$$

so we can never have free energy $F(\beta) > 0$ by the old definition. Instead:

$$Z_{N,\sigma} = E^\tau\left(e^{\beta|\tau\cap\sigma|_{\sigma_N}}\mathbf{1}_{\sigma_N\in\tau}\right), \quad F(\beta) = \lim_N \frac{1}{N}\log Z_{N,\sigma}.$$

Here $\mathbf{1}_{\sigma_N\in\tau}$ is a convenience which does not alter the free energy.

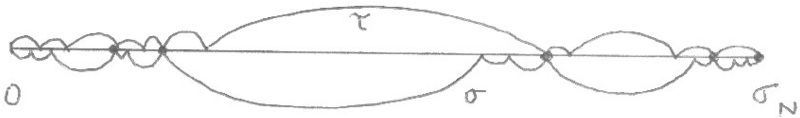

Corresponding annealed model:

$$Z_N^{ann} = E^{\tau\sigma}\left(e^{\beta|\tau\cap\sigma|_{\sigma_N}}\mathbf{1}_{\sigma_N\in\tau}\right).$$

This may be dominated by unusually short trajectories, say σ_N, τ_N both $O(N)$. Related to the "usual" homogeneous model with renewal $\tau\cap\sigma$ and fixed length n:

$$Z_n^{hom} = E^{\tau\sigma}\left(e^{\beta|\tau\cap\sigma|_n}\mathbf{1}_{n\in\sigma}\right).$$

In fact

$$Z_n^{hom} \le \sum_{N=1}^{n} Z_N^{ann}, \quad Z_N^{ann} \le \sum_{n=1}^{\infty} Z_n^{hom}.$$

Can use this to show:

Lemma 1

$\beta_c^{ann} = \beta_c^{hom}$.

Consequence: $\beta_c^{ann} = 0$ if and only if $\tau \cap \sigma$ is recurrent. (Transient renewal must be "bribed" to return to the axis.) Same as

$$\sum_{n=1}^{\infty} P^{\tau}(n \in \tau) P^{\sigma}(n \in \sigma) = \infty.$$

By Doney (1997), for $\alpha, \tilde{\alpha} < 1$,

$$P^{\tau}(n \in \tau) \sim Cn^{-(1-\alpha)} \varphi(n)^{-1}, \quad P^{\sigma}(n \in \sigma) \sim Cn^{-(1-\tilde{\alpha})} \tilde{\varphi}(n)^{-1}.$$

$\tau \cap \sigma$ is always recurrent for $\alpha + \tilde{\alpha} > 1$, depends on φ for $\alpha + \tilde{\alpha} = 1$.

Question: When does $\beta_c^{qu} = \beta_c^{ann}$?

Connection between the usual model and pinning by renewals: Birkner, Greven, den Hollander (2010). In the Gaussian-disorder case (where $h_c^{ann} = -\beta/2$), critical points differ in the usual model if and only if

$$\lim_{T\to\infty} \frac{1}{m_T} \lim_N \frac{1}{N} \log E^\sigma E^\omega \log E^{\tau_T} \left\{ \exp\left(\beta \sum_{n=1}^{\sigma_N} \left[\left(\omega_n - \frac{\beta}{2} \right) \mathbf{1}_{n\in\tau} + 1_{n\in\tau\cap\sigma} \right] \right) \right\} > 0.$$

Here τ_T is the renewal τ with gaps truncated at T, and m_T is the mean of the truncated gap. First term in the sum corresponds to the usual model at the annealed critical point $h_c^{ann} = -\beta/2$. Second term means disorder is supplemented by 1 at times $n \in \sigma$; limit would be 0 without this supplement.

Tail exponent 0: $\tilde{\alpha} = 0$ means σ is "extremely sparse": $\sigma_N \gg N^k$ for all k (e.g. exponentially large for RW in 2 dimensions.) In the recurrent case, the tail exponent of $\tau \cap \sigma$ is

$$\overline{\alpha} = \alpha + \tilde{\alpha} - 1$$

so in the borderline case $\alpha + \tilde{\alpha} = 1$, $\tau \cap \sigma$ is either "barely transient" or "recurrent but extremely sparse."

In our main theorem we rule out extremely sparse disorder σ.

Main Result:

Theorem 2

For the pinning-by-renewal model:
(i) If $\tilde{\alpha} > 0$ *and* $\alpha + \tilde{\alpha} \geq 1$ *then* $\beta_c^{qu} = \beta_c^{ann}$*. If also* $\tau \cap \sigma$ *is transient (possible only for* $\alpha + \tilde{\alpha} = 1$*), this means* $\beta_c^{qu} = \beta_c^{ann} > 0$*.*
(ii) If

$$\frac{1 - \alpha - \tilde{\alpha}}{\tilde{\alpha}} > \frac{1}{2}$$

(so $\tau \cap \sigma$ *is transient), then* $0 < \beta_c^{ann} < \beta_c^{qu}$*.*

Note that (ignoring marginal cases)

$$\tau \cap \sigma \text{ transient } \leftrightarrow \tilde{\alpha} < 1 - \alpha, \qquad \text{condition in (ii) } \leftrightarrow \frac{3}{2}\tilde{\alpha} < 1 - \alpha,$$

so (ii) says σ is more sparse than is required for transience of $\tau \cap \sigma$. The situation for $0 < \frac{1-\alpha-\tilde{\alpha}}{\tilde{\alpha}} \leq 1/2$ is unclear.

Proof sketch for $\beta_c^{qu} = \beta_c^{ann}$ when $\alpha + \tilde{\alpha} \geq 1$

Strategy for τ to be pinned: find favorable parts of σ, and visit them! But what is "favorable"? The good part is, we don't need to know.

Rate function $I(\delta)$ satisfying

$$P^{\sigma\tau}(L_N \geq \delta N) \approx e^{-NI(\delta)}.$$

Let $\beta > \beta_c^{hom}$ so $F^{hom}(\beta) > 0$. Variational formula for $F^{hom}(\beta)$:

$$F^{hom}(\beta) = \sup_{\delta}(\beta\delta - I(\delta)).$$

Choose $\tilde{\delta}$ with $\beta\tilde{\delta} - I(\tilde{\delta}) > \frac{1}{2}F^{hom}(\beta)$. Look for favorable segments of the quenched σ where τ can achieve contact fraction $\geq \tilde{\delta}$ with a not-too-small probability.

Favorability of σ (or of any length-L segment of σ) for pinning τ is measured by

$$g_L(\sigma) = P^\tau(|\tau \cap \sigma|_L \geq \tilde{\delta} L, \sigma_L \in \tau) \mathbf{1}_{\sigma_L \leq L^q}.$$

(q large, fixed, so $\sigma_L \leq L^q$ just rules out extremely long σ.) Note $\sigma_L \in \tau$ has a cost that is only polynomial in L, so $|\tau \cap \sigma|_L \geq \tilde{\delta} L$ is the main event here. We decompose the space of σ's according to favorability: for $\epsilon > 0$,

$$\begin{aligned} e^{-I(\tilde{\delta})L - o(L)} &\leq P^{\tau\sigma}(|\tau \cap \sigma|_L \geq \tilde{\delta} L, \sigma_L \in \tau, \sigma_L \leq L^q) \\ &= E^\sigma(g_L(\sigma)) \\ &\leq \sum_{0 \leq k \leq 1/\epsilon} P^\sigma\left(g_L(\sigma) \in (e^{-(k+1)\epsilon I(\tilde{\delta})L}, e^{-k\epsilon I(\tilde{\delta})L}] \right) e^{-k\epsilon I(\tilde{\delta})L} \\ &\qquad + e^{-(1+\epsilon)I(\tilde{\delta})L}. \end{aligned}$$

Note $P^\sigma(\cdots)$ is a cost borne by σ, and $e^{-k\epsilon I(\tilde{\delta})L}$ is a cost borne by τ, so each k corresponds to a different cost split.

Take the k_0 term corresponding to the optimal cost split (the largest term in the sum) and η small, so that

$$P^\sigma\left(g_L(\sigma) \in (e^{-(k_0+1)\epsilon I(\tilde{\delta})L}, e^{-k_0\epsilon I(\tilde{\delta})L}]\right)e^{-k_0\epsilon I(\tilde{\delta})L}$$

is a positive fraction of the full probability

$$P^{\tau\sigma}(|\tau\cap\sigma|_L \geq \tilde{\delta}L, \sigma_L\in\tau, \sigma_L\leq L^q).$$

Then $\lambda = k_0\epsilon$ represents the fraction of the cost borne by τ in the optimal cost split, and moving $e^{-k\epsilon I(\tilde{\delta})L}$ to the other side we get:

$$P^\sigma\left(g_L(\sigma)\in(e^{-(\lambda+\epsilon)I(\tilde{\delta})L}, e^{-\lambda I(\tilde{\delta})L}]\right) \geq e^{-(1-\lambda+\epsilon)I(\tilde{\delta})L}.$$

Let A denote the event above. Divide σ into blocks of L returns; corresponding gaps in block i are $B_i = (W_{(i-1)L+1}, \ldots, W_{iL})$. Event A is a function of a block so it makes sense to call block i *accepting* (in σ) if $B_i \in A$. Independent from block to block.

On an accepting block, τ can "score big": cost is reduced by factor λ, to hit $\tilde{\delta}L$ of the renewals in σ. In fact the gain for τ is

$$E^\tau\left(e^{\beta|\tau\cap\sigma|_{\sigma_L}}\mathbf{1}_{\sigma_L\in\tau}\right)\mathbf{1}_{B_1(\sigma)\in A} \geq e^{(\beta\tilde{\delta}-(\lambda+\epsilon)I(\tilde{\delta}))L}\mathbf{1}_{B_1(\sigma)\in A}.$$

But there is a cost for τ to find (exp. rare) accepting blocks; frequency

$$p_A := P^\sigma(B_i \in A) \geq e^{-(1-\lambda+\epsilon)I(\tilde{\delta})L}.$$

Let M_i be the index of the ith accepting block; then $M_i - M_{i-1}$ are independent geometric r.v.'s with parameter p_A. We can bound $Z_{M_k L,\sigma}$ below by the contribution from trajectories τ which visit every accepting block, and hit the σ renewals marking the start and end of the block:

$$\log Z_{M_k L,\sigma} \geq \sum_{i=1}^{k} \Big(\log P^\tau(\sigma_{(M_i-1)L} - \sigma_{M_{i-1}L} \in \tau) + \beta\tilde{\delta}L - (\lambda+\epsilon)I(\tilde{\delta})L \Big)$$

Here the log term is the cost to find the ith accepting block from the $(i-1)$st, and the rest is the gain for τ in that block. Log terms are i.i.d. functions of σ.

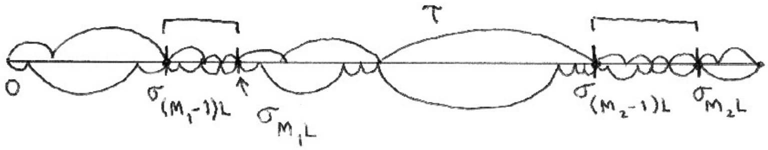

Therefore

$$F(\beta) \geq \liminf_k \frac{1}{M_k L} \log Z_{M_k L,\sigma}$$
$$\geq \frac{1}{E^\sigma(M_1)} \left(\frac{1}{L} E^\sigma \log P^\tau(\sigma_{(M_1-1)L} \in \tau) + \beta\tilde{\delta} - (\lambda + \epsilon) I(\tilde{\delta}) \right).$$

Approximate size of the probability on the right:

$$P^\tau(n \in \tau) = n^{(\alpha\wedge 1)-1}\varphi(n)^{-1}, \quad \sigma_n \approx n^{1/(\tilde{\alpha}\wedge 1)}, \quad (M_1 - 1)L \approx \frac{L}{p_A},$$

which leads to (with η small)

$$E^\sigma \log P^\tau(\sigma_{(M_1-1)L} \in \tau) \geq -\frac{1-(\alpha\wedge 1)+\eta}{\tilde{\alpha}\wedge 1}\left(\log L + \log \frac{1}{p_A}\right).$$

The assumption $\alpha + \tilde{\alpha} \geq 1$ means the fraction here is at most 1+(small). $\log 1/p_A$ is the σ share of the cost, at most about $(1-\lambda-\epsilon)I(\tilde{\delta})L$.

The key is that the cost for τ to find accepting blocks is no worse (up to small ϵ) than the cost of having such blocks occur in the annealed system, since

$$\frac{1-(\alpha \wedge 1)}{\tilde{\alpha} \wedge 1} \leq 1.$$

This leads to

$$F(\beta) > 0 \quad \text{if } \epsilon \text{ is small.}$$

It is essential that when τ moves from one accepting block to the next, it does not have to do it in a single jump to avoid bad regions of disorder, since the disorder is nonnegative. This contrasts with the "usual" model where the ω_n's can be negative. Otherwise the numerator would be bigger than $1-(\alpha \wedge 1)$.

Proof sketch for $0 < \beta_c^{ann} < \beta_c^{qu}$ when $(1 - \alpha - \tilde{\alpha})/\tilde{\alpha} > 1/2$

Use fractional moments. Derrida-Giacomin-Lacoin-Toninelli (2007), plus other papers by these authors. Originally a method for other disordered systems, e.g. Aizenman-Molchanov (1993).

Polymer length $\sigma_N, N = JL$, divide σ (up to σ_N) into J blocks of L renewals. For $\mathcal{I} \subset \{1, \ldots, J\}$ let $Z_{N,\sigma}(\mathcal{I})$ be the contribution to $Z_{N,\sigma}$ from trajectories that visit exactly these blocks. Fix $0 < \gamma < 1$ (close to 1.) Then since

$$Z_{N,\sigma} = \sum_{\mathcal{I}} Z_{N,\sigma}(\mathcal{I}), \quad \text{we have} \quad Z_{N,\sigma}^{\gamma} \leq \sum_{\mathcal{I}} Z_{N,\sigma}(\mathcal{I})^{\gamma}.$$

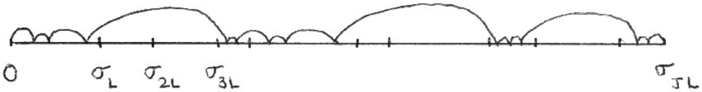

$\mathcal{I} = \{1, 4, 5, 8, \ldots\}$

Interchanging log and E^σ after using this inequality shows

$$\frac{1}{N} E^\sigma \log Z_{N,\sigma} = \frac{1}{\gamma N} E^\sigma \log Z_{N,\sigma}^\gamma \leq \frac{1}{\gamma N} \log \sum_{\mathcal{I}} E^\sigma [Z_{N,\sigma}(\mathcal{I})^\gamma].$$

Consider a change of measure for the disorder, from P^σ to some $\overline{P}^\sigma$:

$$\begin{aligned} E^\sigma [Z_{N,\sigma}(\mathcal{I})^\gamma] &= \overline{E}^\sigma \left[Z_{N,\sigma}(\mathcal{I})^\gamma \frac{dP^\sigma}{d\overline{P}^\sigma} \right] \\ &\leq \left(\overline{E}^\sigma (Z_{N,\sigma}(\mathcal{I})) \right)^\gamma \left(\overline{E}^\sigma \left[\left(\frac{dP^\sigma}{d\overline{P}^\sigma} \right)^{\gamma/(1-\gamma)} \right] \right)^{1-\gamma}. \end{aligned}$$

The change of measure must be chosen so that the Radon-Nikodyn factor here is at most a constant (say e) for each block visited; then

$$E^\sigma [Z_{N,\sigma}(\mathcal{I})^\gamma] \leq e^{|\mathcal{I}|} \left(\overline{E}^\sigma (Z_{N,\sigma}(\mathcal{I})) \right)^\gamma.$$

The factor $\overline{E}^{\sigma}(Z_{N,\sigma}(\mathcal{I})$ in this bound can be viewed as one term of the partition function ($=$ sum over $\mathcal{I}$) for a renormalized annealed system. Power $\gamma \leftrightarrow$ change exponent $1+\alpha$ to $\gamma(1+\alpha)$, in the gap distribution $P^{\sigma}(\sigma_1 = n) = n^{-(1+\alpha)}\varphi(n)$. Power still > 1 if γ near 1.

The renormalized annealed system serves as an upper bound and has an effective reward (potential) of order $\gamma \log \overline{E}^{\sigma}(Z_{L,\sigma})$ for each length-L block (i.e. renormalized site) visited. So we need to choose $\overline{P}^{\sigma}$ to make $\overline{E}^{\sigma}(Z_{L,\sigma})$ small, so that its log is $\ll 0$ and the renormalized process is depinned, free energy 0. But we must do this without making the Radon-Nikodym factor large.

$$f_L(\sigma) = \frac{d\overline{P}^{\sigma}}{dP^{\sigma}} \quad \text{(function of } (\sigma_1, \ldots, \sigma_L))$$

should be small for those rare σ blocks that are favorable to being hit by τ, and larger on unfavorable blocks.

Idea from Birkner, Greven, den Hollander (2011), Berger, Toninelli (2010), Birkner, Sun (2009): re-express the partition function via $e^\beta = 1 + z$:

$$\begin{aligned} Z_{L,\sigma} &= E^\tau \left[(1+z)^{\sum_{n=1}^{L} \mathbf{1}_{\sigma_n \in \tau}} \mathbf{1}_{\sigma_L \in \tau} \right] \\ &= E^\tau \left[\sum_{m=1}^{L} z^m \sum_{1 \le i_1 < \cdots < i_m = L} \prod_{k=1}^{m} \mathbf{1}_{\sigma_{i_k} \in \tau} \right] \\ &= \sum_{m=1}^{L} \sum_{1 \le i_1 < \cdots < i_m = L} \prod_{k=1}^{m} z P^\tau(\sigma_{i_k} - \sigma_{i_{k-1}} \in \tau) \\ &= \sum_{m=1}^{L} \sum_{1 \le i_1 < \cdots < i_m = L} \prod_{k=1}^{m} z P^\tau(\sigma_{i_k - i_{k-1}} \in \tau). \end{aligned}$$

Now use transience to rewrite the last probability:

$$e^{\beta_c^{ann}} = e^{\beta_c^{hom}} = \frac{1}{P^{\sigma\tau}((\tau \cap \sigma)_1 < \infty)}$$

and

$$\frac{1}{z_c^{ann}} = \frac{1}{e^{\beta_c^{ann}} - 1} = E^{\sigma\tau}(|\tau \cap \sigma|_\infty) = \sum_j P^{\sigma\tau}(\sigma_j \in \tau)$$

yield

$$zP^\tau(\sigma_{i_k - i_{k-1}} \in \tau) = \frac{z}{z_c^{ann}} K^*(i_k - i_{k-1}) \frac{P^\tau(\sigma_{i_k} - \sigma_{i_{k-1}} \in \tau)}{E^\sigma P^\tau(\sigma_{i_k} - \sigma_{i_{k-1}} \in \tau)},$$

where

$$K^*(n) = \frac{P^{\sigma\tau}(\sigma_n \in \tau)}{\sum_j P^{\sigma\tau}(\sigma_j \in \tau)}.$$

K^* can be viewed as a distribution of gaps in a renewal τ^*, with a renewal at time n corresponding to $\sigma_n \in \tau$ for the corresponding trajectories.

Then

$$Z_{L,\sigma} = E^{\tau^*}\left[\left(\prod_{k=1}^{|\tau^*|_L} \frac{z}{z_c^{ann}} \frac{P^\tau(\sigma_{i_k} - \sigma_{i_{k-1}} \in \tau)}{E^\sigma P^\tau(\sigma_{i_k} - \sigma_{i_{k-1}} \in \tau)}\right) \mathbf{1}_{L\in\tau^*}\right]$$

so we can also view $\frac{P^\tau(\sigma_{i_k}-\sigma_{i_{k-1}}\in\tau)}{E^\sigma P^\tau(\sigma_{i_k}-\sigma_{i_{k-1}}\in\tau)}$ as the reward associated with the segment of σ between the returns at i_{k-1} and i_k. Note this partition function corresponds to a partition function of a polymer of fixed length L.

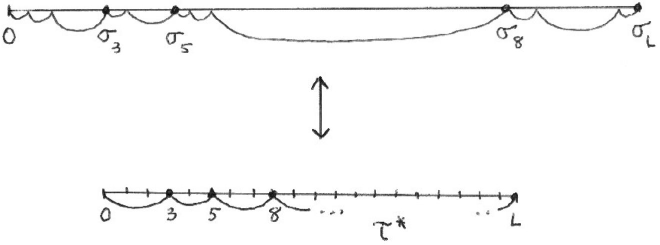

We want a change of measure $d\overline{P}^{\sigma} = f(\sigma)dP^{\sigma}$ for which

$$\overline{E}^{\sigma}(Z_{L,\sigma}) = E^{\tau^*}\left[\left(\prod_{k=1}^{|\tau^*|_L} \frac{z}{z_c^{ann}} \frac{\overline{P}^{\sigma\tau}(\sigma_{\tau_k^*} - \sigma_{\tau_{k-1}^*} \in \tau)}{P^{\sigma\tau}(\sigma_{\tau_k^*} - \sigma_{\tau_{k-1}^*} \in \tau)}\right) \mathbf{1}_{L\in\tau^*}\right] \qquad (1)$$

is small. (Note τ distribution is unchanged in $\overline{P}^{\sigma\tau}$.) But the Radon-Nikodym factor

$$\overline{E}^{\sigma}\left[\left(\frac{1}{f(\sigma)}\right)^{\gamma/(1-\gamma)}\right]$$

must be near 1, cannot be large. So $f(\sigma)$ should only be small on a small set of σ's—those which make the main contribution to (1).

Consider the events

$$D_L(\tau^*) = \{(\sigma,\tau) : \sigma_{\tau_k^*} \in \tau \text{ for all } k \leq |\tau^*|_L\}$$

satisfying (for $f_L(\sigma) = d\overline{P}^{\sigma}/dP^{\sigma}$)

$$E^{\sigma\tau}[f_L(\sigma)\mathbf{1}_{D_L(\tau^*)}] = \overline{P}^{\sigma\tau}(D_L(\tau^*)) = \prod_{k=1}^{|\tau^*|_L} \overline{P}^{\sigma\tau}(\sigma_{\tau_k^*} - \sigma_{\tau_{k-1}^*} \in \tau)$$

This means that we can rewrite further:

$$\overline{E}^{\sigma}(Z_{L,\sigma}) = E^{\tau^*}\left[E^{\sigma\tau}\left(f_L(\sigma) \mid D_L(\tau^*)\right)\left(\frac{z}{z_c^{ann}}\right)^{|\tau^*|_L}\mathbf{1}_{L\in\tau^*}\right]$$

while by comparison

$$E^{\sigma}(Z_{L,\sigma}) = E^{\tau^*}\left[\left(\frac{z}{z_c^{ann}}\right)^{|\tau^*|_L}\mathbf{1}_{L\in\tau^*}\right].$$

Thus we need that for "most" τ^*, $E^{\sigma\tau}(f_L(\sigma) \mid D_L(\tau^*))$ is small, though $E^{\sigma\tau}(f_L(\sigma)) = 1$. $f_L(\sigma)$ must reflect what the information "τ hit all the renewals σ_j with $j \in \tau^*$" tells us about σ, not for a specific τ^* but averaged over all τ^*.

One natural guess: τ's success tells us that the length σ_L of the polymer is shorter than usual. This turns out to be incorrect. Even for a single gap σ_1, the information that τ hit σ_1 does not shorten σ_1 (much) on average. But it does shrink the tail, i.e. it reduces the probability that σ_1 is exceptionally large. This suggests the right approach: using a statistic $f_L(\sigma)$ based on the number Y_L of gaps $\sigma_j - \sigma_{j-1}, j \leq L$, exceeding some large K_L.

It is necessary to verify that the typical resulting change in Y_L can be seen above the noise of the random fluctuations in Y_L. This turns out to be a consequence of the fact that when the tail exponent $\alpha^* = (1 - \alpha - \tilde{\alpha})/\tilde{\alpha}$ of τ^* exceeds $1/2$, there is infinite overlap in τ^*.

Thus we have found our desired $f_L(\sigma)$ to make the change of measure.

Alternate definitions of pinning: In the pinning-by-renewals model, one has

$$|\tau|_{\sigma_N} \gg N$$

when $\alpha \wedge 1 > \tilde{\alpha}$, that is, the fraction of τ renewals that hit σ renewals is approaching 0. Should this really count as pinning? We can restrict to trajectories with $|\tau|_{\sigma_N} \leq bN$ for some $b \geq 1$. Let

$$F_2(\beta, b) = \lim_N \frac{1}{N} \log Z_{N,\sigma}(|\tau|_{\sigma_N} \leq bN).$$

Theorem 3

If $\tilde{\alpha} > 0$ and $\alpha + \tilde{\alpha} \geq 1$, then for every $\beta > \beta_c^{ann}$ there exists $b_0(\beta)$ such that $F_2(\beta, b) > 0$ for all $b \geq b_0(\beta)$.

Proof Idea: As a lower bound for $Z_{N,\sigma}(|\tau|_{\sigma_N} \leq bN)$, consider only trajectories τ which "skip over" very long gaps between accepting blocks.

Another twist: What if the rewards for hitting σ_j's are random?

$$H_{N,\sigma,\omega}(\tau) = \sum_{j=0}^{N} (\beta\omega_j + h)\mathbf{1}_{\sigma_j \in \tau} \quad \text{or} \quad \sum_{i=0}^{\infty} (\beta\omega_i + h)\mathbf{1}_{i\in\sigma}\mathbf{1}_{i\in\tau}$$

ω_j tied to σ_j or ω_i tied to site i give the same model. But it is convenient to use the second form:

$$E^{\omega}E^{\sigma}(\log Z_{N,\omega,\sigma}) \le E^{\omega}(\log E^{\sigma} Z_{N,\omega,\sigma}) \le \sum_{n=1}^{\infty} Z_{n,\omega}^{usual}$$

for all N, where on the right the renewal sequence is $\tau \cap \sigma$. Recall in the recurrent case the tail exponent of $\tau \cap \sigma$ is $\alpha + \tilde{\alpha} - 1$. If this is $> 1/2$ then for $h_c^{ann}(usual) < h < h_c^{qu}(usual)$ we have the right side finite (Mourrat 2012), so

$$\lim_N \frac{1}{N} E^{\omega}E^{\sigma}(\log Z_{N,\omega,\sigma}) = 0.$$

The usual model and the above pin-to-renewal model with random rewards have the same h_c^{ann}. This proves:

Theorem 4

Suppose $\alpha + \tilde{\alpha} > 3/2$ and ω_1 has a finite exponential moment. Then for the pin-to-renewal model with random rewards, $h_c^{ann} < h_c^{qu}$.