From Euler to Stellaris: Some of my favorite open problems in probability

David Aldous

6 April 2022

Talk is somewhat "off-topic" for this meeting. I don't try to engage today's mainstream math probability. Most of my research nowadays involves toy models only a few steps away from data, with some plausible story. But for this audience I'm mostly going Old School with just "intellectual curiosity" problems. And I'm hoping to attract some young people to engage them.

- Random Eulerian circuits
- The Nearest Unvisited Vertex walk, and 4X games.
- Covering a compact space by growing random balls.
- A Markov chain-derived mapping of distributions on compact spaces

1. Random Eulerian circuits

A rather obvious observation in introductory graph theory is

Lemma

A finite connected undirected graph has at least one **spanning tree**, *that is a connected edge-subgraph which is a tree.*

Euler proved what's often regarded as "the first theorem in graph theory".

Theorem

A finite, strongly connected, directed graph which is balanced (each vertex has in-degree = out-degree) has at least one **Eulerian circuit**, *that is a tour using each edge exactly once.*

.

Within probability theory, there is a large literature on uniform random spanning trees, because they relate to many other discrete structures – see Lyons - Peres monograph *Probability on Trees and Networks.*

In contrast, there is very little literature on uniform random Eulerian circuits. This is curious because there's **a surprising connection** between the two topics.

In a balanced directed graph, take any spanning tree, with directed edges toward an arbitrary root. From the root do a walk, at each stage arbitrarily choosing an unused edge but saving the spanning-tree-edge until last. This always gives an Eulerian circuit [easy].

True (but **not obvious**) that with a *uniform* random spanning tree and *uniform* random walk-step choices we get a *uniform* random Eulerian circuit.

Fact: It is quite easy to simulate a uniform random spanning tree (of an arbitrary finite connected graph). So we can then simulate a uniform Eulerian circuit on a balanced graph.

Now as a simple example let us consider the discrete torus $\mathbb{Z}_N^d$. Replace each edge by 2 directed edges. So in-degree = out-degree = 2d. Any Eulerian circuit consists of 2d excursions from the origin.

For $d \geq 3$ simple random walk on $\mathbb{Z}^d$ is transient, which strongly suggests the following, which is supported by simulation.

Open problem. Prove (in fixed $d \geq 3$) that of the $2d$ excursions at the origin, each has length $O(1)$ or $\Omega(N^d)$, not of intermediate order (as $N \to \infty$).

I have no idea how one might prove this – can't do theoretical analysis of algorithm output.

Take-away message. There is an unexpected connection between random Eulerian tours and random spanning trees. Known for a long time, but **apparently never exploited**.

2a: The Nearest Unvisited Vertex walk on Random Graphs

Consider a connected undirected graph G on n vertices, where the edges e have positive real lengths $\ell(e)$. Imagine a robot that can move at speed 1 along edges. We need a rule for how the robot chooses which edge to take after reaching a vertex. Most familiar is the "random walk" rule, choose edge e with probability proportional to $\ell(e)$ or $1/\ell(e)$. One well-studied aspect of the random walk is the *cover time*, the time until every vertex has been visited.

Instead of the usual random walk model, let us consider the *nearest unvisited vertex* (NUV) walk

> *after arriving at a vertex, next move at speed* 1 *along the path to the closest unvisited vertex*

and continue until every vertex has been visited. Note this is deterministic and has some length (= time) $L_{NUV}(G, v_0)$ where v_0 is the initial vertex.

Of course *distance* $d(v, v')$ is shortest path length. In informal discussion we imagine lengths are scaled so that distance to closest neighbor is order 1, so L_{NUV} must be at least order n.

Natural first question: when is it $O(n)$ rather than larger order?

There is scattered old "algorithms" literature discussing the NUV walk as heuristics for TSP or as an algorithm for a robot exploring an unknown environment, but that literature quickly moved on to better algorithms.

I will say some results from my preprint *The Nearest Unvisited Vertex Walk on Random Graphs*. Part 2b will explain one motivation.

There is a key starting math observation – implicit but rather obscured in the old literature. For now, we stay with non-random graphs.

Consider **ball-covering**: for $r > 0$ define $N(r) = N(G, r)$ to be the minimal size of a set S of vertices such that every vertex is within distance r from some element of S. In other words, such that the union over $s \in S$ of $\mathrm{Ball}(s, r)$ covers the entire graph.

Proposition

(i) $N(r) \leq 1 + L_{NUV}/r, \ 0 < r < \infty$.
(ii) $L_{NUV} \leq 2 \int_0^{\Delta/2} N(r) \, dr$ *where* $\Delta = \max_{v,w} d(v, w)$ *is the diameter of the graph.*

Note that for continuous spaces, *metric entropy* implies a notion of *dimension* via $N(r) \approx r^{-\mathrm{dim}}$ as $r \downarrow 0$. In our discrete context, if we have *dimension* in the sense

$$N(r) \approx nr^{-\mathrm{dim}}, \quad 1 \ll r \ll \Delta$$

then the Proposition has informal interpretation that L_{NUV} is always $O(n)$ when $\mathrm{dim} > 1$.

Isolating that Proposition as the starting point, we can easily recover the two classical (1970s) results for non-random graphs.

Corollary

There is a constant A such that, for the complete graph on n arbitrary points in the area-n square, with Euclidean lengths,

$$L_{NUV} \leq An.$$

Note this implies the well known corresponding result $L_{TSP} \leq An$.

Corollary

Let $a(n)$ be the maximum, over all connected n-vertex graphs with edge lengths and all initial vertices, of the ratio L_{NUV}/L_{TSP}. Then $a(n) = \Theta(\log n)$.

The ball-covering relation is not helpful from the algorithms viewpoint. But it is useful for some **random** graph models. In particular, in a model where we take a unweighted graph and then assign random edge-lengths, understanding "balls" is precisely the basic issue in **first passage percolation (FPP)**.

Consider the random graph G_m that is the $m \times m$ grid, that is the subgraph of the Euclidean lattice $\mathbb{Z}^2$, assigned i.i.d. edge-lengths $\ell(e) > 0$. with $\mathbb{E}\ell(e) < \infty$. Because the shortest edge-length at a given vertex is $\Omega(1)$, clearly L_{NUV} is $\Omega(m^2)$. Using the shape theorem for FPP on $\mathbb{Z}^2$ one can show

Corollary

For the 2-dimensional grid model G_m above, the sequence $(m^{-2}L_{NUV}(G_m),\ m \geq 2)$ is tight.

The same techniques would give $O(n)$ upper bounds in other simple models of n-vertex random graphs.

Open problems

- Are there general methods (subadditivity or local weak limits don't seem to work) to prove existence of a limit $c = \lim_n n^{-1} L_{NUV}(G_n)$ for simple models?
- Evaluate c?
- Order of magnitude of $\mathrm{var}(L_{NUV})$ not clear from our small-scale simulations – seems $n^{1\pm\varepsilon}$.

Take-away message. There is an unexpected connection between the NUV walk and FPP. Does this suggest that the variance problem is difficult?

2b: Games people play

I'm interested in probability and graphs; and also games.
Search MathSciNet for "graph and game" in title: get 654.
None are games people actually play.

Are there "graph" games that millions of people do play?

Yes: Go, for instance.

But such traditional board games are closely tied to a fixed graph; I want games that can be played on a random graph, different every time you play. Are there any?

Well . . . . . . yes and no.

4X games

4X (abbreviation of Explore, Expand, Exploit, Exterminate) is a subgenre of strategy-based computer and board games, and include both turn-based and real-time strategy titles. The gameplay involves building an empire. Emphasis is placed upon economic and technological development, as well as a range of non-military routes to supremacy. (Wikipedia).

A representative game is *Stellaris*.

show https://steamdb.info/app/281990/graphs/

show https://stellaris.paradoxwikis.com/Category:Game_concepts

4X games are very complicated in detail. Much over-simplifying, let me invent a simple game which abstracts the common elements of the initial "Explore, Expand" phases, as follows.

My simple game. Copy the background setting of the NUV walk. There is a connected undirected graph G on n vertices, where the edges e have positive real lengths $\ell(e)$. You have a unit that you can move at speed 1 along edges. But you only see a neighborhood of the vertices that you have already visited. The "neighborhood" is defined so that you could (if you choose) implement the NUV walk. Make a game with k players, each with a unit moving simultaneously. A vertex you visit becomes part of your empire; other players cannot visit.

Easy fact: if at least one player is not completely stupid, this simple game will end with the vertices partitioned into the connected empires of the different players.

Goal: form the largest empire.

The graph is different every time you play, a realization of some unknown probability distribution on graphs.

(very vague) Open Problem: What is a good strategy?

- **aggressive:** move away from starting vertex in some direction until meeting an opponent, then attempt to block.
- **defensive:** colonize a growing ball around your starting vertex.
- **NUV:** seems somewhat between.

Intuitively, the best strategy depends on connectivity – for a locally tree-like graph with large visible neighborhood, "aggressive" is clearly better. Fun student project, in progress.

Take-away message? Clearly not do-able as theorem-proof mathematics, but good to "search away from the streetlight" and engage actual 21st century activity.

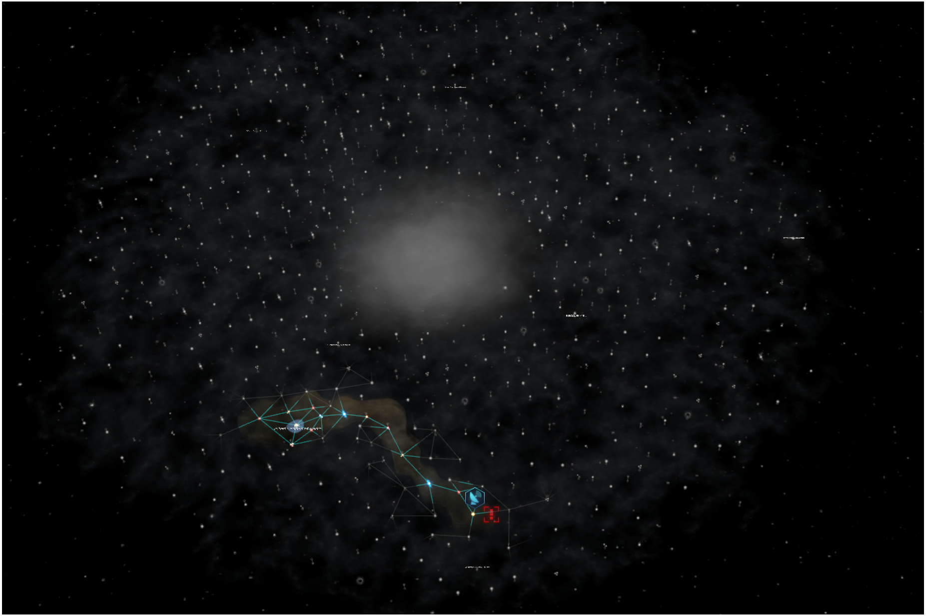

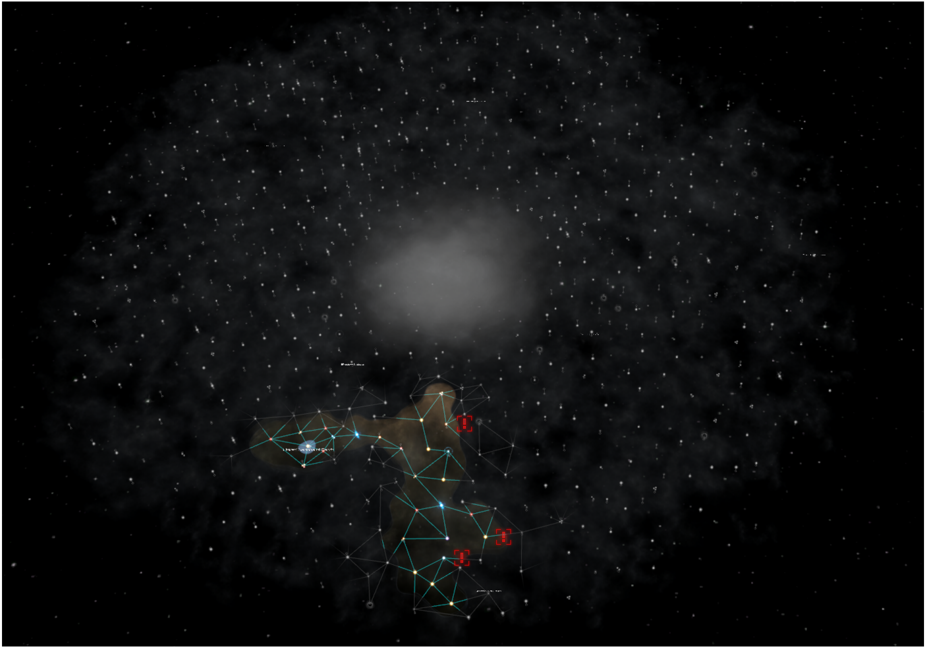

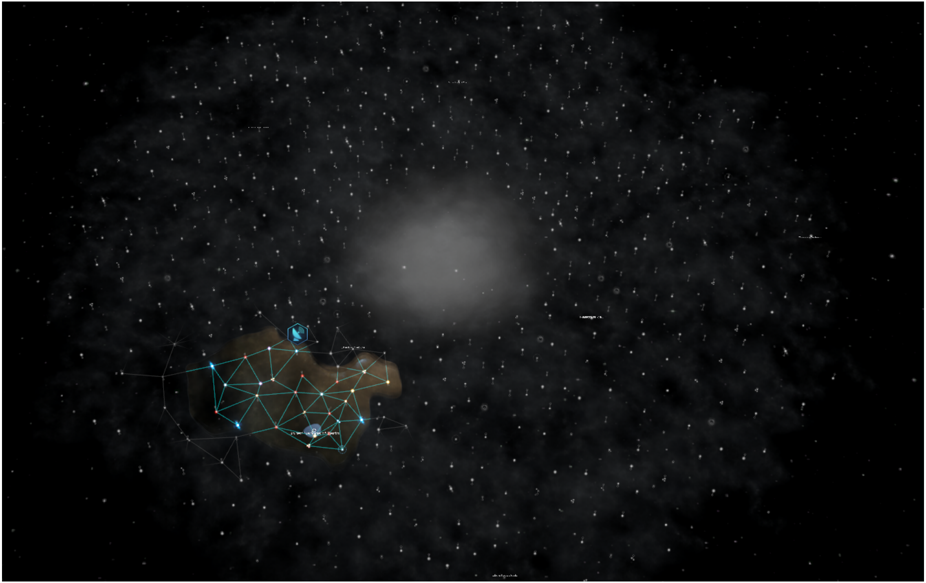

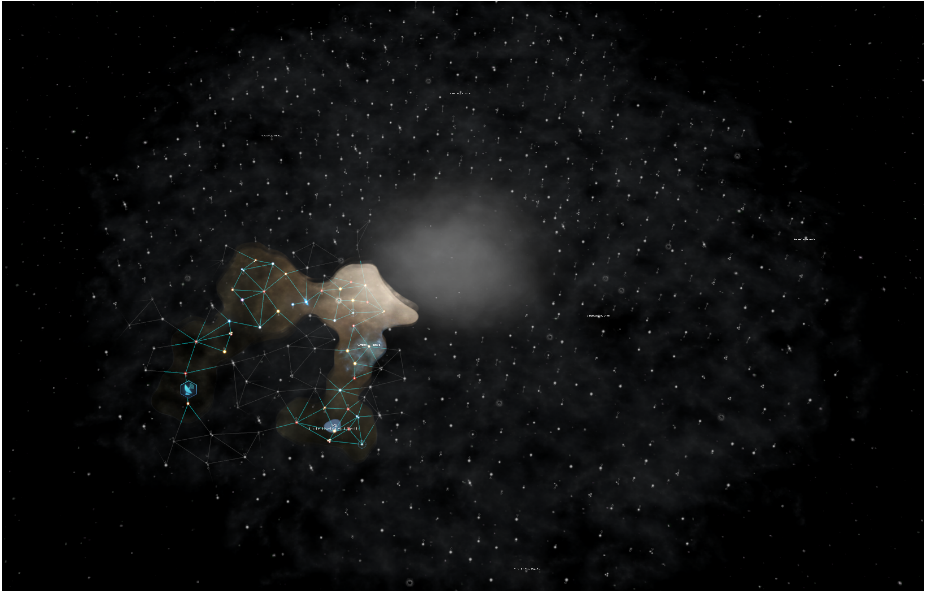

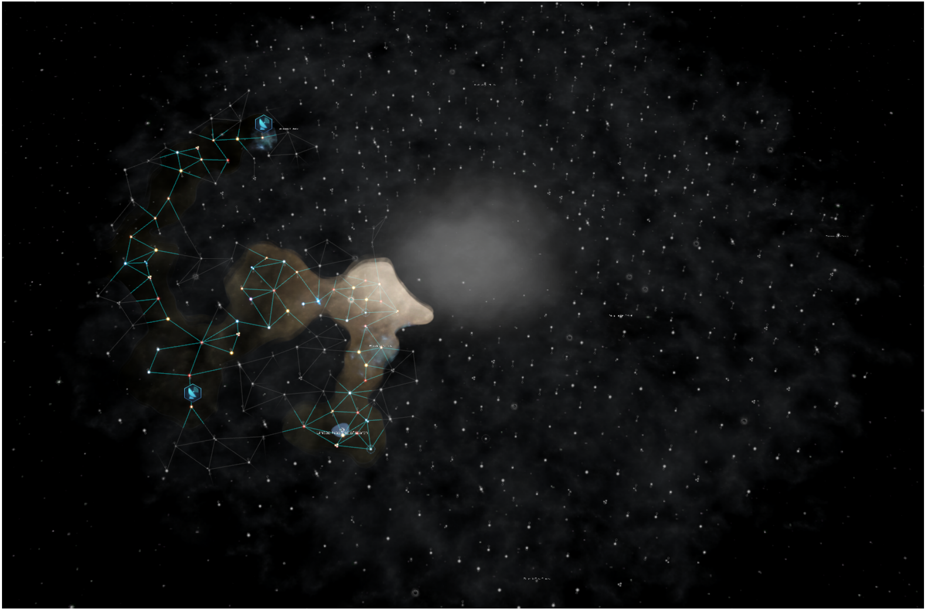

Topics 3 and 4

Write (S, d) for a compact metric space and $\mathcal{P}(S)$ for the space of probability measures on S, with the weak topology. I want to study processes that can be defined on any S and be parametrized by any $\theta \in \mathcal{P}(S)$. One standard example is the i.i.d. empirical process (usually studied in greater generality). Are there other interesting examples?

One reason for studying processes defined on every S is that one can seek both general results and also sharper results for any given S – providing much scope for collaboration with students.

I will discuss two unrelated such processes. The first has been studied (a little) in $\mathbb{R}^d$, the second is apparently novel. Can you think of any others?

3. A random coverage problem

Details in the arXiv preprint

Covering a compact space by fixed-radius or growing random balls .

We have general results (not deep); much scope for more precise analysis on particular S.

Consider a compact metric space (S, d), a probability measure θ on S, but now introduce two rates $0 < \lambda < \infty$ and $0 < v < \infty$. Write $0 < \tau_1 < \tau_2 < \dots$ for the times of a rate-λ Poisson process, and write $\sigma_1, \sigma_2, \dots$ for i.i.d. random points of S from distribution θ. The verbal description

> *seeds arrive at times of a Poisson process at i.i.d. random positions, and then create balls whose radius grows at rate v*

is formalized as the set-valued *growth process*

$$\mathcal{X}(t) := \cup_{i:\tau_i \leq t} \, \mathrm{ball}\,(\sigma_i, v(t - \tau_i)). \tag{1}$$

We study the **cover time**

$$C := \min\{t : \ \mathcal{X}(t) = S\}$$

which is finite because $\mathbb{E}\tau_1 = 1/\lambda$ and so (for any θ)

$$1/\lambda \leq \mathbb{E}C \leq 1/\lambda + \Delta/v \tag{2}$$

where Δ is the diameter of S.

We can "standardize" the model by choosing time and distance units to make $\lambda = v = 1$. This is "without loss of generality" as regards explicit inequalities, though does affects asymptotics for a sequence S_n. For the standardized model we can define

$$\chi(S) = \min_\theta \mathbb{E}_\theta C$$

which is just a number associated with S. One **open problem** would be to systematically compare with other numbers associated with compact spaces S. Another **open problem** is that there is no canonical notion of *uniform* distribution on S; to what extent can the minimizing θ play a role as proxy for uniform?

So what **is** done in the preprint? Because S is compact we have ball-covering numbers

$$N(r) := \text{ minimum number of radius } r \text{ balls that cover } S$$

which are finite. It's natural to try to relate the one number $\chi(S)$ to the function $r \to N(r)$.

It is not hard to find some first general upper and lower bounds for $\chi(S)$ in the standardized model.

First, by considering the uniform distribution θ_r on the set of $\mathrm{cov}(r)$ points, we find (cf. coupon-collector)

$$\chi(S) \leq \min_{r>0}[r + N(r) \cdot (1 + \log N(r))].$$

Second, some θ attains $\chi(S)$, so consider the seeds of that process as a set to upper bound $\mathrm{cov}(r)$. We find

$$\chi(S) \geq \sup\{r : N(3r) > 9r\}.$$

How good are these general bounds? There is a notion I'll call *rough dimension*: A space like $[0, L]^d$ has rough dimension d characterized by

$$N(r) \asymp (L/r)^d \text{ for } r << L.$$

Here the general lower and upper bounds, for such a space, are of orders $L^{\frac{d}{d+1}}$ and $L^{\frac{d}{d+1}} \log L$.

For the actual torus $[0, L]^d$ we know sharp asymptotics as $L \to \infty$ as part of extensive historical "applied probability" work on coverage processes in Euclidean space.

Open problem: Study infinite-dimensional examples.

Take-away message?. This particular model is perhaps not well motivated, but "proof of concept" that one **can** devise non-trivial processes that make sense on arbitrary compact spaces.

4. A Markov chain and a mapping $\mathcal{P}(S) \to \mathcal{P}(S)$

Take a pair (j, k) with $k \geq 2$ and $1 \leq j \leq k$. For any probability distribution $\theta \in \mathcal{P}(S)$, define a Markov chain on compact S by:

- from state s, take k i.i.d. (θ) samples, and jump to the j'th closest.

By considering the natural coupling, it is not hard to prove (a good homework problem in a course discussing coupling?) that

Theorem

Every *such chain converges in distribution (and variation distance) to some unique stationary distribution.*

Comment: Model apparently not studied. We mentally envisage S and θ as continuous, but a metric space might have only finitely many points. For this and other reasons, we explicitly specify "break ties uniform randomly". If ties are possible, the chain may not be Feller.

Take a pair (j, k) with $k \geq 2$ and $1 \leq j \leq k$. For any probability distribution $\theta \in \mathcal{P}(S)$, define a Markov chain on S by:

- from state s, take k i.i.d. (θ) samples, and jump to the j'th closest.

Call the stationary distribution $\pi_{j,k}(\theta)$. This defines a mapping $\pi_{j,k} : \mathcal{P}(S) \to \mathcal{P}(S)$. What happens when we iterate this mapping? In particular, what are the fixed points of this mapping?

Our original motivation: Fixed points would have a kind of "self-similarity under sampling" property and might provide interesting examples of specific non-uniform distributions on compact spaces S.

[2021 work with undergraduates Madelyn Cruz and Shi Feng: seeking more collaborators – can share extensive working notes]

The coupling proof tells us nothing explicit about the relation between θ and $\pi_{j,k}(\theta)$. By considering one step of the stationary chain we have, for $\pi = \pi_{j,k}(\theta)$

$$\theta^k(A) \leq \pi(A) \leq k\theta(A), \ A \subseteq S$$

and so π and θ are mutually absolutely continuous.

We study the **iterative process** which iterates the map $\pi_{j,k} : \mathcal{P}(S) \to \mathcal{P}(S)$. This does not have a simple "process" interpretation. And this project is **maybe crazy** because we don't know explicitly what the map $\pi_{j,k}$ actually is. However, for any given S and (j,k) there is an explicit equation determining fixed points θ so (in principle) one can try to solve to find all the fixed points.

The bottom line is:

- Simulations and conjectures reveal very counter-intuitive behavior.
- We have only some fragments of rigorous proofs.
- Proving anything substantial seems beyond the authors' capabilities ......

so the remainder of this talk is also rather fragmentary.

First minor observations

Consider $\phi \in \mathcal{P}(S)$ which is invariant (that is, a fixed point) under $\pi_{j,k}$ for given (j,k). If the support of ϕ is smaller than S then it is more natural to consider ϕ as an invariant measure on the support. So our basic question can better be phrased as

- Given S and (j,k), what are all the invariant measures **with full support** on S?

On every compact metric space S we have an obvious "preservation of symmetry" result for the action of $\pi_{j,k}$.

Lemma

If $\theta \in \mathcal{P}(S)$ is invariant under an isometry ι of S then $\pi_{j,k}(\theta)$ is also invariant under ι.

Fragment 1: Fixed points existing by symmetry

In some cases there are distributions $\phi \in \mathcal{P}(S)$ which are invariant (that is, fixed points) "by symmetry" for all $\pi_{j,k}$. In particular

(i) The distribution δ_s degenerate at one point s;

(ii) The uniform two-point distribution $\delta_{s_1,s_2} = \frac{1}{2}(\delta_{s_1} + \delta_{s_2})$;

(iii) The Haar probability measure on a compact group S with a metric invariant under the group action.

(iv) On a finite space S, a sufficient condition for the uniform distribution to be invariant is that S is *transitive*, that is if for each pair s, s' there is an isometry taking s to s'. This is equivalent to the finite case of Haar measure. But for finite S a weaker condition suffices, because all that matters is the *rank matrix* – see later.

In those cases the distribution is invariant for all $\pi_{j,k}$. So the question becomes:

for a particular S and (j, k), are there invariant distributions with full support, other than those "forced by symmetry" as above?

Fragment 2: The case $S = \{a, b\}$ is not trivial

One might suppose that the case of a 2-element set $S = \{a, b\}$ would be trivial, but it is not. Parametrizing a distribution θ on S by $p := \theta(a)$, we view the mapping $\pi_{j,k} : \mathcal{P}(S) \to \mathcal{P}(S)$ as a mapping $\pi_{j,k} : [0, 1] \to [0, 1]$ defined as follows. In the associated 2-state Markov chain, the transition probabilities are

$$\text{prob}(a \to b) = \mathbb{P}(\mathrm{Bin}(k, p) < j); \quad \text{prob}(b \to a) = \mathbb{P}(\mathrm{Bin}(k, p) > k - j)$$

for Binomial random variables. From the stationary distribution we find

$$\pi_{j,k}(p) = \frac{\mathbb{P}(\mathrm{Bin}(k, p) > k - j)}{\mathbb{P}(\mathrm{Bin}(k, p) > k - j) + \mathbb{P}(\mathrm{Bin}(k, p) < j)}.$$

So a fixed point is a solution of the equation

$$\pi_{j,k}(p) = p. \tag{3}$$

We know by symmetry that $p = 0, p = 1/2, p = 1$ are fixed points; are there others? By symmetry it is enough to consider $0 < p < 1/2$.

We have not tried to find solutions analytically, but we will show results of numerical calculations of the iterates $\pi^n_{j,k}(p), n = 1, 2, 3, \dots$. For a given (k, j), we observe three possible types of qualitative behavior:

1. $\pi^n_{j,k}(p) \to 0$ as $n \to \infty$, for all $0 < p < 1/2$.
2. $\pi^n_{j,k}(p) \to 1/2$ as $n \to \infty$, for all $0 < p < 1/2$.
3. There exists a critical value $p_{crit} \in (0, 1/2)$ such that p_{crit} is invariant : $\pi_{j,k}(p_{crit}) = p_{crit}$
   and $\pi^n_{j,k}(p) \to 0$ as $n \to \infty$, for all $0 < p < p_{crit}$
   and $\pi^n_{j,k}(p) \to 1/2$ as $n \to \infty$, for all $p_{crit} < p < 1/2$.

For us, (3) is the interesting case: there is a non-obvious fixed point, but it is unstable. It first arises with $k = 5, j = 4$, as shown in the Figure. We see the critical value $p_{crit} = 0.17267\dots$.

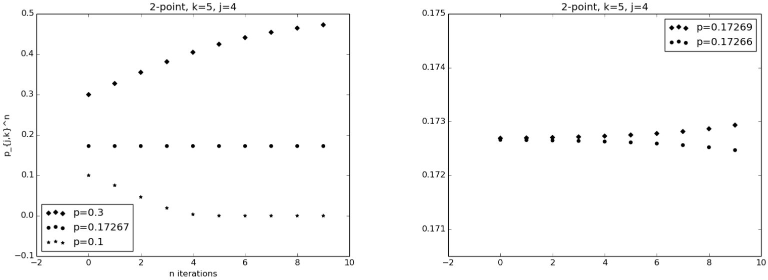

Figure: $S = \{a, b\}$; $k = 5, j = 4$. Iterates $n = 0, 1, 2, \ldots, 10$. Left panel shows type (3) behavior, Right panel shows the unstable fixed point at 0.17267.

Maybe excessive to claim 0.17267... is *interesting* but encouraging that there exist non-obvious fixed points (for certain (j, k)).

Table: $S = \{a, b\}$ and $2 \le k \le 9$. The values of j with each type of behavior, and (critical values) of critical points.

k	$(0 \leftarrow)$	$(critical)$	$(\rightarrow 1/2)$
2	1		2
3	$[1, 2]$		3
4	$[1-3]$		4
5	$[1-3]$	4 (0.17267)	5
6	$[1-4]$	5 (0.09558)	6
7	$[1-5]$	6 (0.06276)	7
8	$[1-5]$	6 (0.26405)	$[7, 8]$
9	$[1-6]$	7 (0.18884); 8 (0.03364)	9

The Table shows the type of behavior – types (i) or (ii) or (iii) above – for all pairs (j, k) with $k \le 9$. One take-away message is that for $S = \{a, b\}$ there exist some (j, k) for which $\pi_{j,k}$ has fixed points in addition to those existing by symmetry, but these fixed points are unstable.

Of course the 2-point space may be very special. What properties extend to other S? Let's look at the unit interval.

Fragment 3: The case $S = [0, 1]$

We have studied, by simulation, iterates starting from the uniform distribution $U[0, 1]$. Because $U[0, 1]$ is symmetric about $1/2$, all iterates must be symmetric about $1/2$.

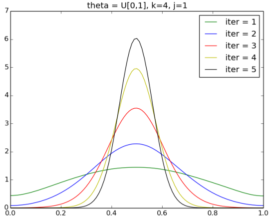

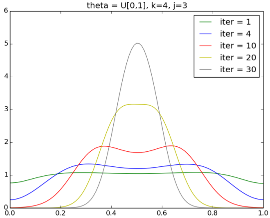

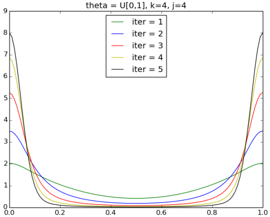

The Figure above shows the case $k = 4$ and the first few iterates of $U[0,1]$ for $j = 1, 2, 3$. Note that, here and throughout, the vertical scale and the numbers of iterations shown may not be the same from one panel to the next. What we see strongly suggests that the iterates are converging, quickly for $j = 1$ but rather slowly for $j = 3$, toward the degenerate distribution $\delta_{1/2}$. This is strongly supported by examining the standard deviations of the iterates, shown on log scale in the Figure below, and suggesting a scaling limit distribution.

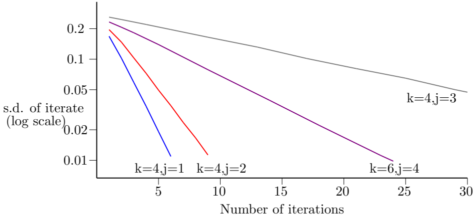

In contrast, the Figure for $j = 4$ strongly suggests that the iterates are converging quickly toward the mixture $\delta_{0,1}$. These two "extreme" behaviors – convergence to $\delta_{1/2}$ for smaller j or to $\delta_{0,1}$ for larger j – appear to hold for all k. The Table shows which behavior appears to hold in simulations for each pair (j, k) with $k \leq 9$.

Table: Conjectured limits of iterates from $U[0, 1]$; the values of j with each type of behavior.

k	$\to \delta_{1/2}$	$\to \delta_{0,1}$
2	1	2
3	$[1, 2]$	3
4	$[1 - 3]$	4
5	$[1 - 4]$	5
6	$[1 - 4]$	$[5, 6]$
7	$[1 - 5]$	$[6, 7]$
8	$[1 - 6]$	$[7, 8]$
9	$[1 - 7(?)]$	$[8, 9]$

As mentioned earlier, one can always write down an equation for a fixed point. On $[0,1]$ a density function $f(t)$ is a fixed point for $\pi_{j,k}$ iff

Proof. According to (26), invariant distribution must satisfies the following equation for all t in $[0,1]$

$$\begin{aligned}1 = \binom{k}{j-1,1,k-j} \cdot \Bigg\{ & \left[\int_0^{\frac{t}{2}} f_n(y) \left(\int_0^t f_n(x)\,dx\right)^{j-1} \cdot \left(\int_t^1 f_n(x)\,dx\right)^{k-j} dy\right] \\ & + \left[\int_{\frac{t}{2}}^t f_n(y) \left(\int_{2y-t}^t f_n(x)\,dx\right)^{j-1} \cdot \left(\int_0^{2y-t} f_n(x)\,dx + \int_t^1 f_n(x)\,dx\right)^{k-j} dy\right] \\ & + \left[\int_t^{\frac{1+t}{2}} f_n(y) \left(\int_t^{2y-t} f_n(x)\,dx\right)^{j-1} \cdot \left(\int_0^t f_n(x)\,dx + \int_{2y-t}^1 f_n(x)\,dx\right)^{k-j} dy\right] \\ & + \left[\int_{\frac{1+t}{2}}^1 f_n(y) \left(\int_t^1 f_n(x)\,dx\right)^{j-1} \cdot \left(\int_0^t f_n(x)\,dx\right)^{k-j} dy\right]\Bigg\}\end{aligned}$$

Shi Feng (undergrad) studied this by careful and elaborate calculus, initially in the case $j=2, k=2$. From the "$t=0$" identity one can argue to a contradiction, and this can be made into a rigorous proof of

Theorem

There are no $\pi_{2,2}$-invariant distributions on $[0,1]$ other than those of the form δ_s or δ_{s_1,s_2}.

The argument extends to some, but not all, pairs (j,k).

Simulations of the iterative process on $[0, 1]$ **starting with a non-uniform distribution** show analogous behavior: either convergence to $\delta_{0,1}$ or to δ_s for some s depending on the initial distribution.

At a rigorous level, the key open questions for $S = [0, 1]$ are

- Does there exist (for any (j, k)) any invariant distribution with full support?
- Does there exist (for any (j, k)) any distribution other than δ_s or $\delta_{0,1}$ that occurs as a limit of iterates from some initial distribution with full support?

We suspect the answer to each is "no". Of course, "no" to the second question would imply "no" to the first question.

Fragment 5: Finite S.

A finite metric space can be represented by the matrix D of distances $d(i,j)$. By taking all the non-zero distances to be between 1 and 2, the triangle inequality is automatically satisfied. Consider the example of a 5-element space with distance matrix

$$D = \begin{pmatrix} 0 & 1.714 & 1.341 & 1.656 & 1.74 \\ 1.714 & 0 & 1.298 & 1.794 & 1.03 \\ 1.341 & 1.298 & 0 & 1.715 & 1.844 \\ 1.656 & 1.794 & 1.715 & 0 & 1.524 \\ 1.74 & 1.03 & 1.844 & 1.524 & 0 \end{pmatrix}$$

What matters for our purposes, assuming as in this example that all non-zero distances are distinct, is the *rank matrix* R, where $r(i,j) = 4$ means that $d(i,j)$ is the 4'th smallest of $\{d(i,1), d(i,2), \ldots, d(i,|S|)\}$. For the distance matrix D above, the rank matrix is

$$R = \begin{pmatrix} 1 & 4 & 2 & 3 & 5 \\ 4 & 1 & 3 & 5 & 2 \\ 3 & 2 & 1 & 4 & 5 \\ 3 & 5 & 4 & 1 & 2 \\ 4 & 2 & 5 & 3 & 1 \end{pmatrix}$$

By numerical calculation, for $\pi_{1,2}$ on this space there is an invariant distribution

$$\theta \approx (0.149\ 0.188\ 0.203\ 0.298\ 0.162)$$

for which the transition matrix is

$$K \approx \begin{pmatrix} 0.276 & 0.097 & 0.304 & 0.297 & 0.026 \\ 0.111 & 0.341 & 0.222 & 0.089 & 0.237 \\ 0.159 & 0.265 & 0.365 & 0.185 & 0.026 \\ 0.139 & 0.036 & 0.118 & 0.507 & 0.201 \\ 0.083 & 0.28 & 0.041 & 0.298 & 0.298 \end{pmatrix}$$

This example was found (by Shi Feng) by simulating random distance matrices D, obtaining the rank matrix R, and then numerically solving for invariant distributions θ until finding a solution with full support. Note this involved non-linear equations: we need to solve $\theta K = \theta$ but here K depends on θ, for instance for $\pi_{1,2}$

$$k(i,i) = 1 - (1 - \theta(i))^2$$

$$\text{if } r(i,j) = 5 \text{ then } k(i,j) = \theta^2(j).$$

Note also that for $|S| = 5$ there are only a finite number of possible rank matrices R, so this counter-example is not like a counter-example depending on a real parameter taking a specific value.

State of this project

Disappointing: we have not found interesting distributions on particular S.

Instead we have a range of open problems about ways in which the behavior is non-interesting.

- For which S and (j, k) are there invariant distributions other than those "forced by symmetry"?
- True or false: For every S and every (j, k), every invariant distribution except δ_s and δ_{s_1,s_2} is unstable (to a generic perturbation).
- True or false: For every S, the iterative process for $\pi_{1,2}$ from almost all initial $\phi \in \mathcal{P}(S)$ converges to some δ_s (depending on ϕ).

If not true in general, is it true for $S \subset \mathbb{R}^d$?