Stake-governed random-turn games

Alan Hammond, U.C. Berkeley

Joint work with Gábor Pete, Rényi Institute, BUDAPEST

B.I.R.S. workshop

May 2023

Tug of war ....

.... in ECONOMICS

Harris & Vickers 1987 'Racing with uncertainty...' ≈ 600 citations

.... and in MATHEMATICS

Peres, Schramm, Sheffield & Wilson 2009: 'Tug of war and the infinity Laplacian' ≈ 500 citations

Harmonic functions on a graph

A graph $G = (V, E)$ has boundary $B \subseteq V$ with data $f : B \to \mathbb{R}$.

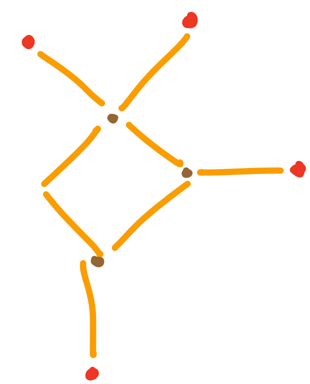

Suppose that $h : V \to \mathbb{R}$, $h|_B = f$, is such that, for each $v \in V \setminus B$, $\omega = h(v) \in \mathbb{R}$ minimizes

$$\sum_{u \sim v} (\omega - h(u))^2 .$$

Then h is harmonic — we may say '2-harmonic' — and

$$h(v) = \frac{1}{\deg(v)} \sum_{u \sim v} h(u) \qquad \forall v \in V \setminus B .$$

2-harmonic functions and simple random walk

Simple random walk $X : \mathbb{N} \to V$, $X(0) = v \in V \setminus B$, jumps at each step to a uniformly chosen neighbour.

Let τ = the time that X reaches B.

Then the 2-harmonic extension $h : V \to \mathbb{R}$, $h|_B = f$, satisfies

$$h(v) = \mathbb{E}\left[f(X(\tau)) \,/\, X(0) = v \right].$$

We could call X the 2-walk!

p-harmonic functions for $p \in (1,\infty)$

Say that $h : V \to \mathbb{R}$, $h|_B = f$, is p-harmonic if,
for each $v \in V \setminus B$, $h(v)$ equals
the minimizer ω of

$$\left(\sum_{u \sim v} |\omega - h(u)|^p \right)^{1/p} .$$

The case $p = \infty$.

Say that $h : V \to \mathbb{R}$, $h|_B = f$,

is *infinity harmonic* if,

for each $v \in V \setminus B$,

$h(v)$ equals the minimizer ω of

$$\max_{u \sim v} |\omega - h(u)|$$

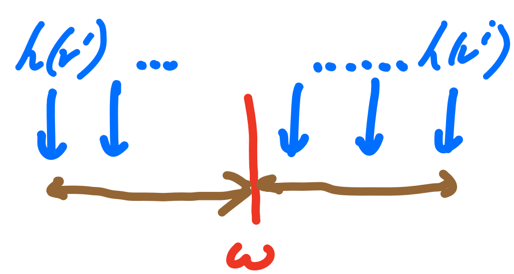

Infinity harmonic functions

In other words, h is ∞-harmonic if

$$h(v) = \frac{1}{2}\left(\max_{u\sim v} h(u) + \min_{u\sim v} h(u)\right)$$

$$\forall v \in V \setminus B.$$

What then is the infinity walk $X : \mathbb{N} \to V$ for which

$$h(v) = \mathbb{E}\left[f(X(\tau)) \,/\, X(0) = v\right] \;?$$

(∞-harmonic) arrival of B

The game RANDOM TUG-OF-WAR

A counter • begins at $v \in V \setminus B$.

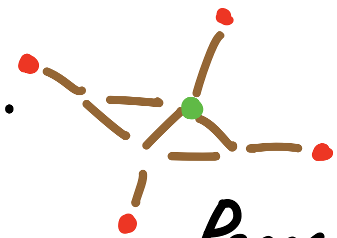

Mina and Maxine win the first turn according to a fair coin flip.

Peres, Schramm, Sheffield, Wilson 2007

The turn victor moves the counter to an adjacent vertex. Similarly for later turns.

The game ends when the counter reaches B, and MINA pays MAXINE

Pay = f(terminal counter location).

Strategies

A strategy for a given player specifies what move to play in any given game position.

Let $\mathcal{S}_- =$ set of strategies for Mina

and $\mathcal{S}_+ =$ set of strategies for Maxine

For $(s_-, s_+) \in \mathcal{S}_- \times \mathcal{S}_+$, set

$$M(s_-, s_+) = \mathbb{E}[\mathrm{Pay}(s_-, s_+)]$$

to be the mean value of the terminal payment Mina $\rightarrow$ Maxine under the strategy pair (s_-, s_+).

Game value

Suppose that Mina must declare to Maxine her strategy at the start of the game.

The BEST she can do is her game value

$$V_- = \inf_{s_- \in S_-} \sup_{s_+ \in S_+} u(s_-, s_+) .$$

Similarly, if it is Maxine who must declare her strategy:

$$V_+ = \sup_{s_+ \in S_+} \inf_{s_- \in S_-} u(s_-, s_+) .$$

$V_+ \leq V_-$ always. If $V_+ = V_-$, then the game has value.

this is the game value

Game value and optimal strategies in tug-of-war

Theorem [PSSW09]

On any finite graph, the game value exists.
It is the infinity harmonic extension

$$h(v) = \frac{1}{2}\left(\max_{u\sim v} h(u) + \min_{u\sim v} h(u)\right)$$

of the boundary data $f : B \to \mathbb{R}$.

Under optimal play, Maxine plays to an h-maximizing neighbour; Mina plays to an h-minimizing one.

Aside: *continuum* infinity *harmonic* functions

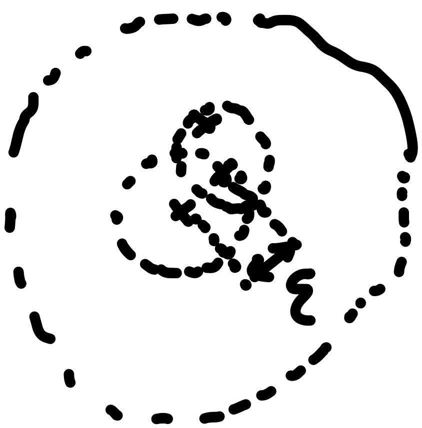

RANDOM TUG-OF-WAR with ε-sized steps in a domain in $\mathbb{R}^d$

$\xrightarrow[2007]{PSSW}$

game value in low ε limit is CONTINUUM ∞-harmonic:

$$0 = \Delta_\infty u := \sum_{i,j} u_{x_i} u_{x_i x_j} u_{x_j}.$$

Question: what is the STRATEGIC IMPORTANCE of a position in a multi-turn game?

the line graph L_{2n}

Suppose that Maxine is offered the right to buy the first move, with the later game running as usual. How much should she be prepared to pay?

If she buys the first move,

her mean payoff under jointly optimal play is $\frac{n+1}{2n} = \frac{1}{2} + \frac{1}{2n}$.

If she does not buy this move, this payoff is $\frac{1}{2}$.

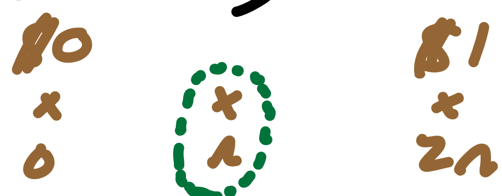

So the fair price of the purchase is $\frac{1}{2n}$.

Equally for any other starting position.

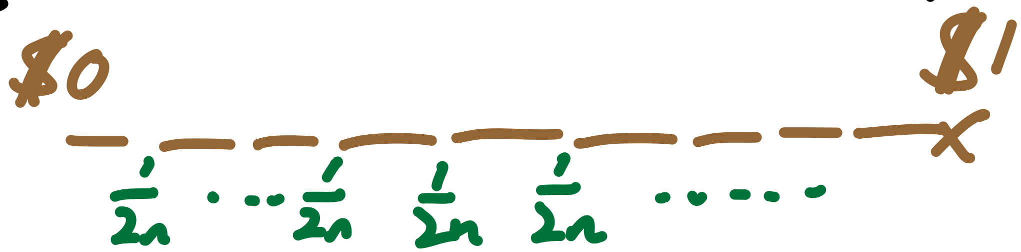

Apparent conclusion:

all positions are EQUALLY important, and each is WORTH $\Theta(1/n)$.

Another view

Under optimal play, the game lasts $\Theta(n^2)$ steps, because the counter X follows SIMPLE RANDOM WALK.

If all STEPS are EQUALLY important, and the total sum at stake is \$1, then the first turn is WORTH $\Theta\left(\frac{1}{n^2}\right)$.

Two differing predictions

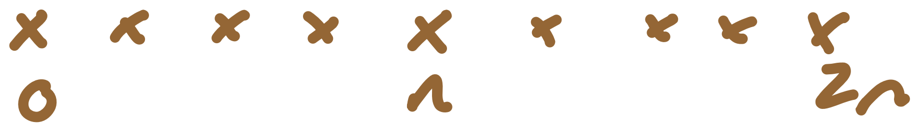

The importance of the first term is

$\Theta\left(\frac{1}{n}\right)$	$\Theta\left(\frac{1}{n^2}\right)$
according to	
the SPACE prediction	the TIME prediction

Stake-governed tug-of-war

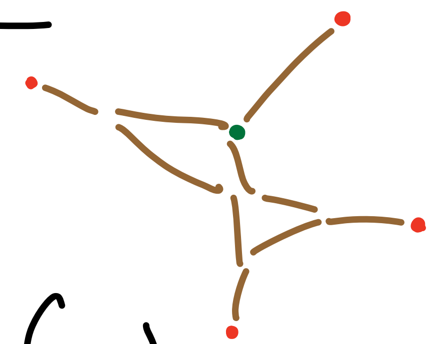

Let $\lambda \in (0,\infty)$. Suppose that Mina and Maxine have respective FORTUNES of one and λ at the OUTSET.

At the start of a given turn, each retains some part of her original fortune. Each is asked to stake some part of this fortune — Maxine stakes a, Mina b.

These sums are deducted from the respective fortunes.

Maxine wins the RIGHT TO MOVE at the turn with probability $\frac{a}{a+b}$; otherwise, Mina does.

The turn victor moves the counter to a neighbouring vertex, as in classical tug-of-war. The next turn proceeds, with the updated fortunes.

At the end of the game, Mina pays Maxine f(terminal counter location).

The players lose all of their remaining fortunes.

Two needed concepts

$\boxed{1}$ λ-biased ∞-harmonic functions $h(\lambda,\cdot)$ satisfy

$$h(\lambda,v) = \frac{\lambda}{1+\lambda}\max_{u\sim v} h(\lambda,u) + \frac{1}{1+\lambda}\min_{u\sim v} h(\lambda,u),$$

$$h(\lambda,\cdot)\big|_{B} = f.$$

These may be computed by Peres-Šunić path decomposition

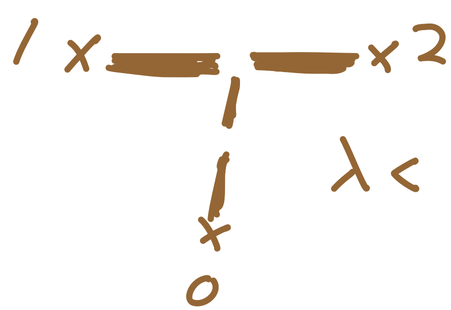

$\lambda < \lambda_c = \frac{\sqrt{5}+1}{2}$

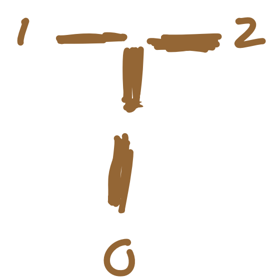

The T-graph

$\lambda > \frac{\sqrt{5}+1}{2}$

Two needed concepts

$\boxed{2}$ Nash equilibrium

$\mathcal{S}_-$ is the space of strategies for Mina ;
$\mathcal{S}_+$ is this space for Maxine .

A pair $(s_-, s_+) \in \mathcal{S}_- \times \mathcal{S}_+$ is a Nash equilibrium if, for any $s_-' \in \mathcal{S}_-$ and $s_+' \in \mathcal{S}_+$,

$$\mathcal{M}(s_-, s_+') \leq \mathcal{M}(s_-, s_+) \leq \mathcal{M}(s_-', s_+) .$$

At any Nash equilibrium, the payoff is the value of the game.

Dream-world game value

On a finite graph, set

$$\Delta(\lambda, v) = \max_{u \sim v} h(\lambda, u) - \min_{u \sim v} h(\lambda, u)$$

Proposition Suppose that $\Delta(\lambda, v) > 0 \ \forall v \in V \setminus B$.
Suppose that there is a pwc Nash equilibrium in stake-governed tug-of-war on the given graph.
Then the game value exists and equals the λ-biased ∞-harmonic function $h(\lambda, \cdot) : V \to \mathbb{R}$ with $h(\lambda, \cdot)|_B = f$.

<u>Proof</u> (Stake strategy stealing)

Let (s_-^0, s_+^0) be an equilibrium.

Now let s_- be a strategy for Mina under which she stakes the same proportion (of her remaining fortune) as Maxine does under S_+^0, and proposes to move to an h-minimizing neighbour.

Then the resulting $h(\lambda, x_i)$ is a bounded supermartingale whose limiting value is the terminal payoff, because the game almost surely ends in finite time.

Thus,

$$\mu(s_-, s_+^\circ) \leq h(x_0) = h(v)$$

But: (s_-°, s_+°) is an equilibrium $\Longrightarrow$

$$\mu(s_-^\circ, s_+^\circ) \leq \mu(s_-, s_+^\circ).$$

Considering a similar strategy for Maxine leads to

$$\mu(s_-^\circ, s_+^\circ) \geq h(v).$$

$\square$

This conditional result prompts

TWO QUESTIONS :

1. Does a pure Nash equilibrium exist ?
2. If it does, then, when Maxine's fortune is λ times Mina's, and the counter is at $v \in V \setminus B$, the players stake a shared proportion of their present fortunes — call this $S(\lambda, v)$.

What then is the STAKE FUNCTION

$$V \setminus B \to \mathbb{R} : \quad v \to S(\lambda, v) \quad ?$$

A plausible argument that identifies $S(\lambda, v)$

Mina has fortune 1, Maxine $\lambda \in (0,\infty)$.

The counter is at $v \in V \setminus B$.

Write $S = S(\lambda, v)$.

Under optimal play, Maxine stakes λS
and Mina S.

What if Maxine deviates slightly,
to offer $(\lambda + \eta) S$, with $0 < \eta \ll 1$
at the first turn, with optimal later play
?

Change in mean terminal payment

$$= (G - L)\,\eta + O(\eta^2)$$

gain (for Maxine)

loss

The gain term equals

$$G = S^{-1}\frac{\lambda}{(1+\lambda)^2} \times \Delta(\lambda, v) = h(\lambda, v_+) - h(\lambda, v_-)$$

reflects the gain in probability of MAXINE winning the first turn

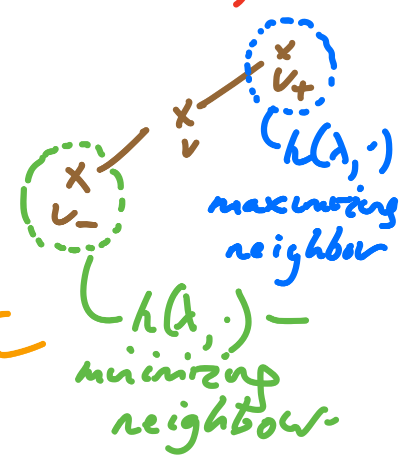

It's tougher for Maxine at turns two and later: her relative fortune is

$$\lambda_{alt} = \frac{\lambda - \lambda(s+\eta)}{1-s} = \lambda - \frac{\lambda\eta}{1-s} + O(\eta^2).$$

Thus, the loss term L takes the form

$$L = (1-s)^{-1}\frac{\lambda}{(1+\lambda)^2}\,\mathbb{E}\sum_{i=1}^{\overset{=\text{game finish}}{\tau}-1}\Delta(\lambda, X_i)$$

Under optimal play,

$$G = L \, .$$

This leads to a formula for the STAKE function $S = S(\lambda, \nu)$:

$$S = \frac{\Delta(\lambda, \nu)}{\mathbb{E} \sum_{i=0}^{\tau-1} \Delta(\lambda, x_i)}$$

the space prediction

similar to the time prediction

Alternatively, we may similarly predict:

$$S = \frac{\Delta(\lambda, \nu)}{(\lambda+1)^2 \frac{\partial}{\partial \lambda} h(\lambda, \nu)}$$

This is space versus money, or

short-term versus long-term.

A plausible picture, but ....

TWO PROBLEMS

1. For some SIMPLE graphs, there is NO pure Nash equilibrium.
2. For some others, $h(\lambda, v)$ fails to be differentiable in λ.

PROBLEM 1 : no pure Nash equilibrium

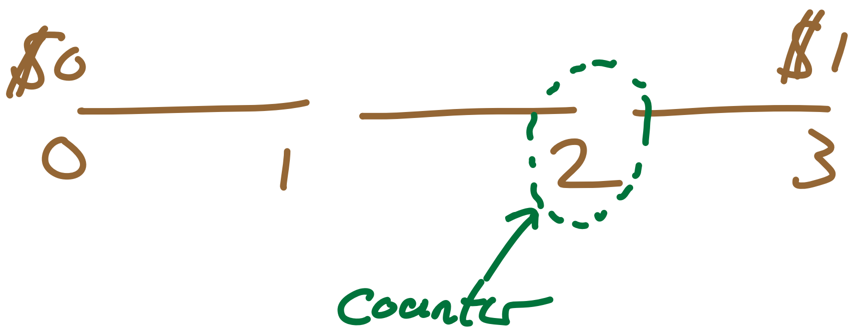

Maxine's relative fortune $= \frac{1}{2}$.

If a pure Nash equilibrium exists, we can plot game value in the case that Maxine stakes $a \in [0, \frac{1}{2}]$, and Mia stakes $b \in [0, 1]$.

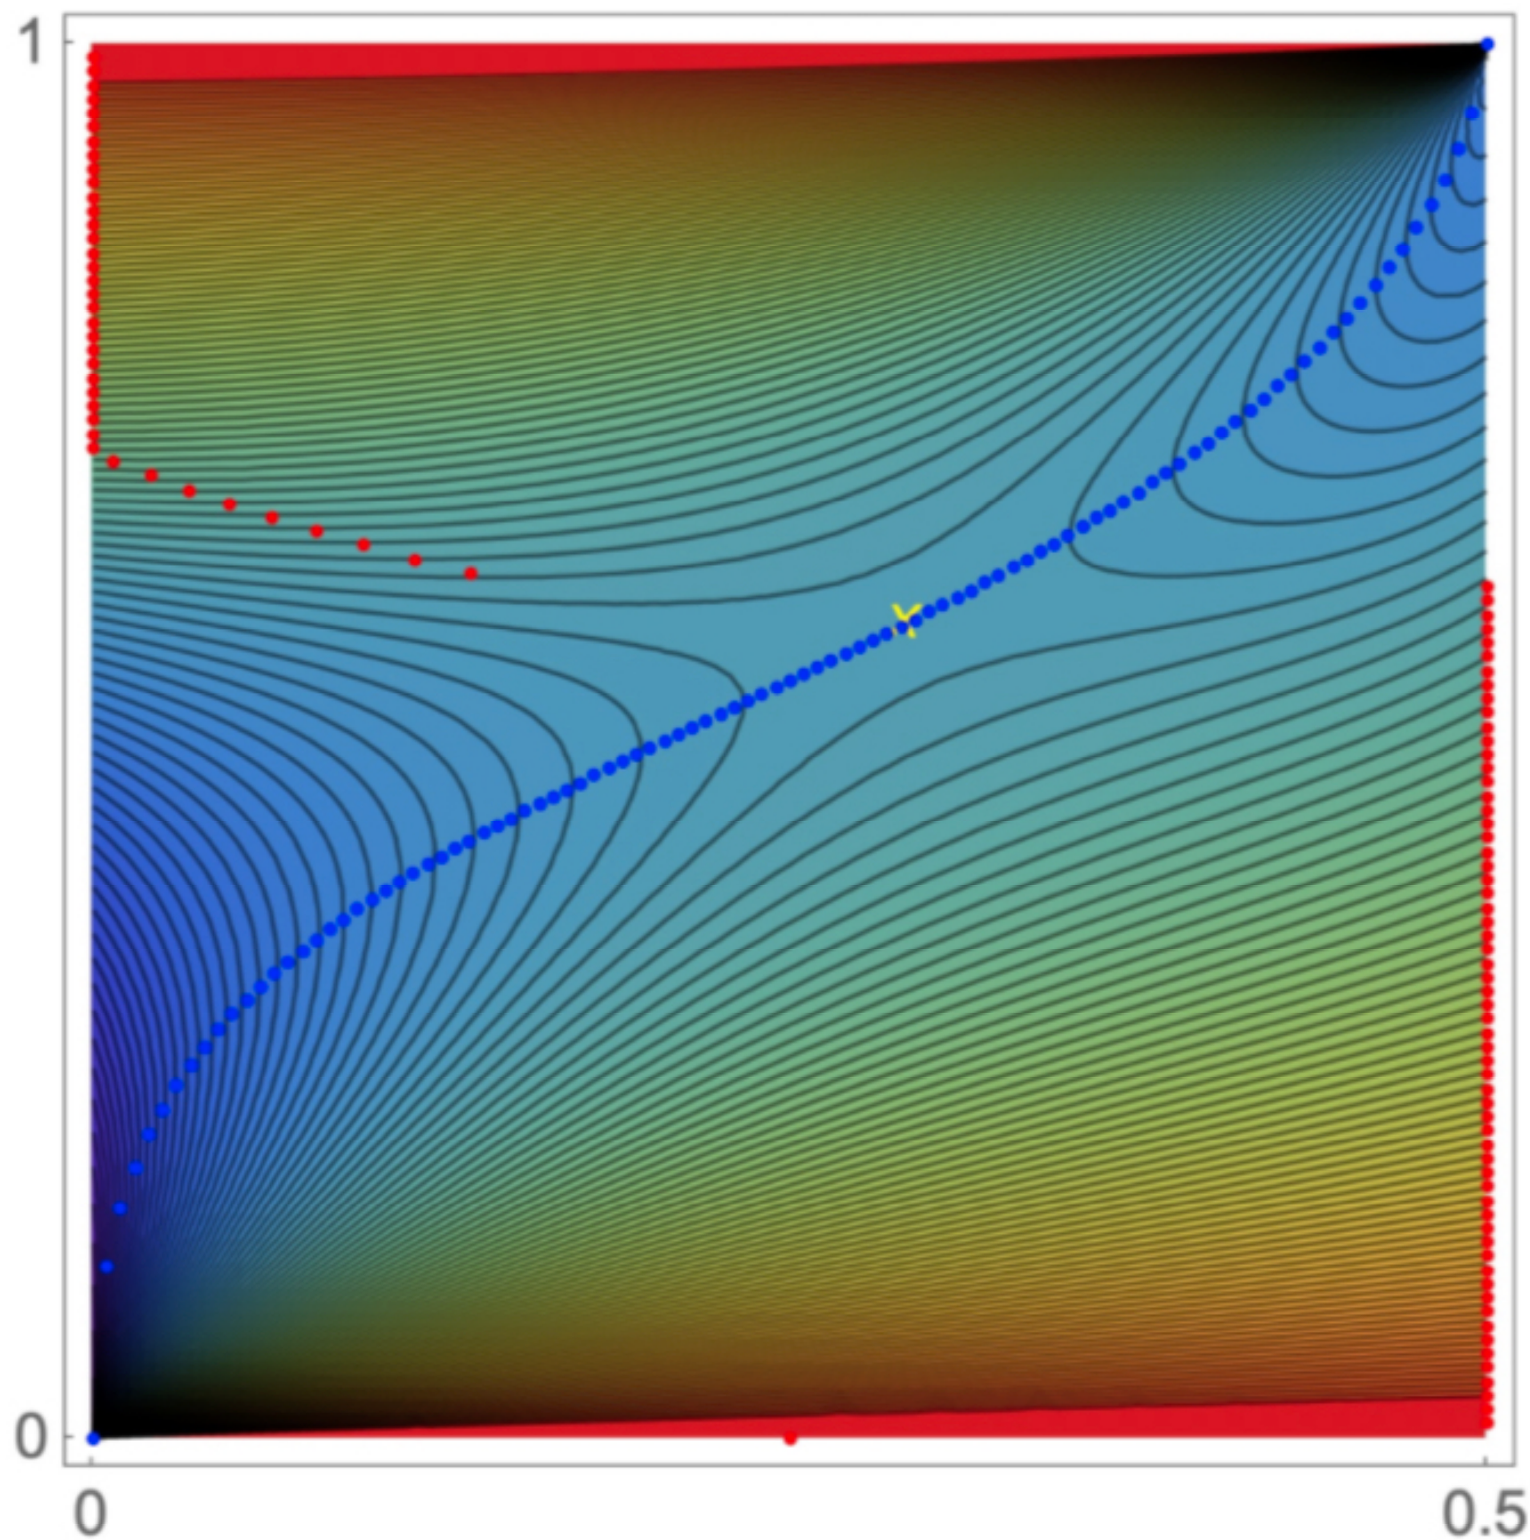
1
0
0
0.5

Problem 2: $\lambda \to h(\lambda, v)$ may not be differentiable:

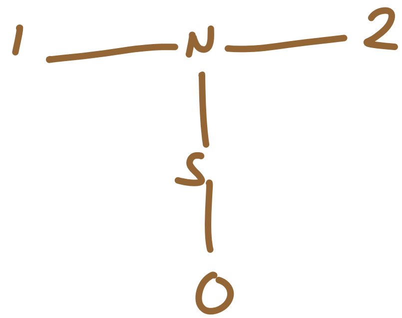

$\lambda < \lambda_c = \frac{1+\sqrt{5}}{2}$

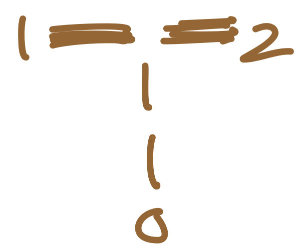

$\lambda > \lambda_c$

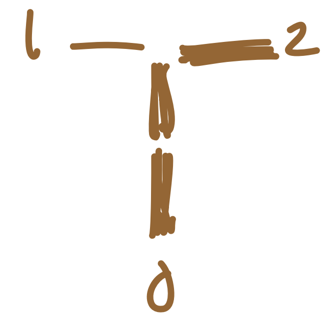

As λ rises through λ_c, the Peres-Šunić path decomposition changes, so that $\lambda \to h(\lambda, v)$ fails to be differentiable at λ_c.

In the proposed STAKE formula,

$$S = \frac{\Delta(\lambda, \nu)}{\mathrm{Denom}},$$

Denom equals $(\lambda+1)^2 \frac{\partial}{\partial\lambda} h(\lambda,\nu)$ or $(\lambda+1)^2 \sum_{i=0}^{\tau-1} \Delta(x_i;\lambda)$.

But Denom is sometimes badly defined:

— h is not regular enough; or

— Mira can force different gameplays X by proposing different moves

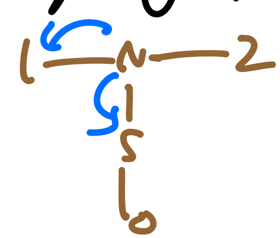

In summary, our TWO PROBLEMS may be construed as:

1. One or other player may be tempted to 'go for broke', so that a local saddle point fails to be global.

2. The λ-biased ∞-harmonic PATH decomposition may switch as λ rises:

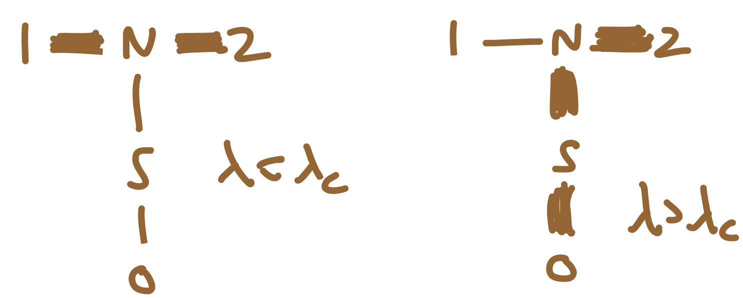

TWO MODIFICATIONS

that we now make in an effort to address the respective difficulties.

1. The game will become LEISURELY.
2. We will consider only a class of graphs called ROOT-REWARD TREES.

Modification 1: the leisurely game.

We introduce a parameter $\varepsilon \in (0,1]$

(to join λ, the relative fortune,

and v, the counter location).

RULE CHANGE:

after players stake at a given turn,

the casino flips a coin that

lands HEADS with probability ε.

If the outcome is HEADS, then the turn proceeds as it would have done originally. If it is TAILS, then no move takes place — the submitted stakes are simply lost to the players, and the counter moves nowhere.

IDEA: low choices of ε

disable the go-for-broke strategy.

would you bet your life-savings

when, more likely than not, they will

simply be swept from you ?

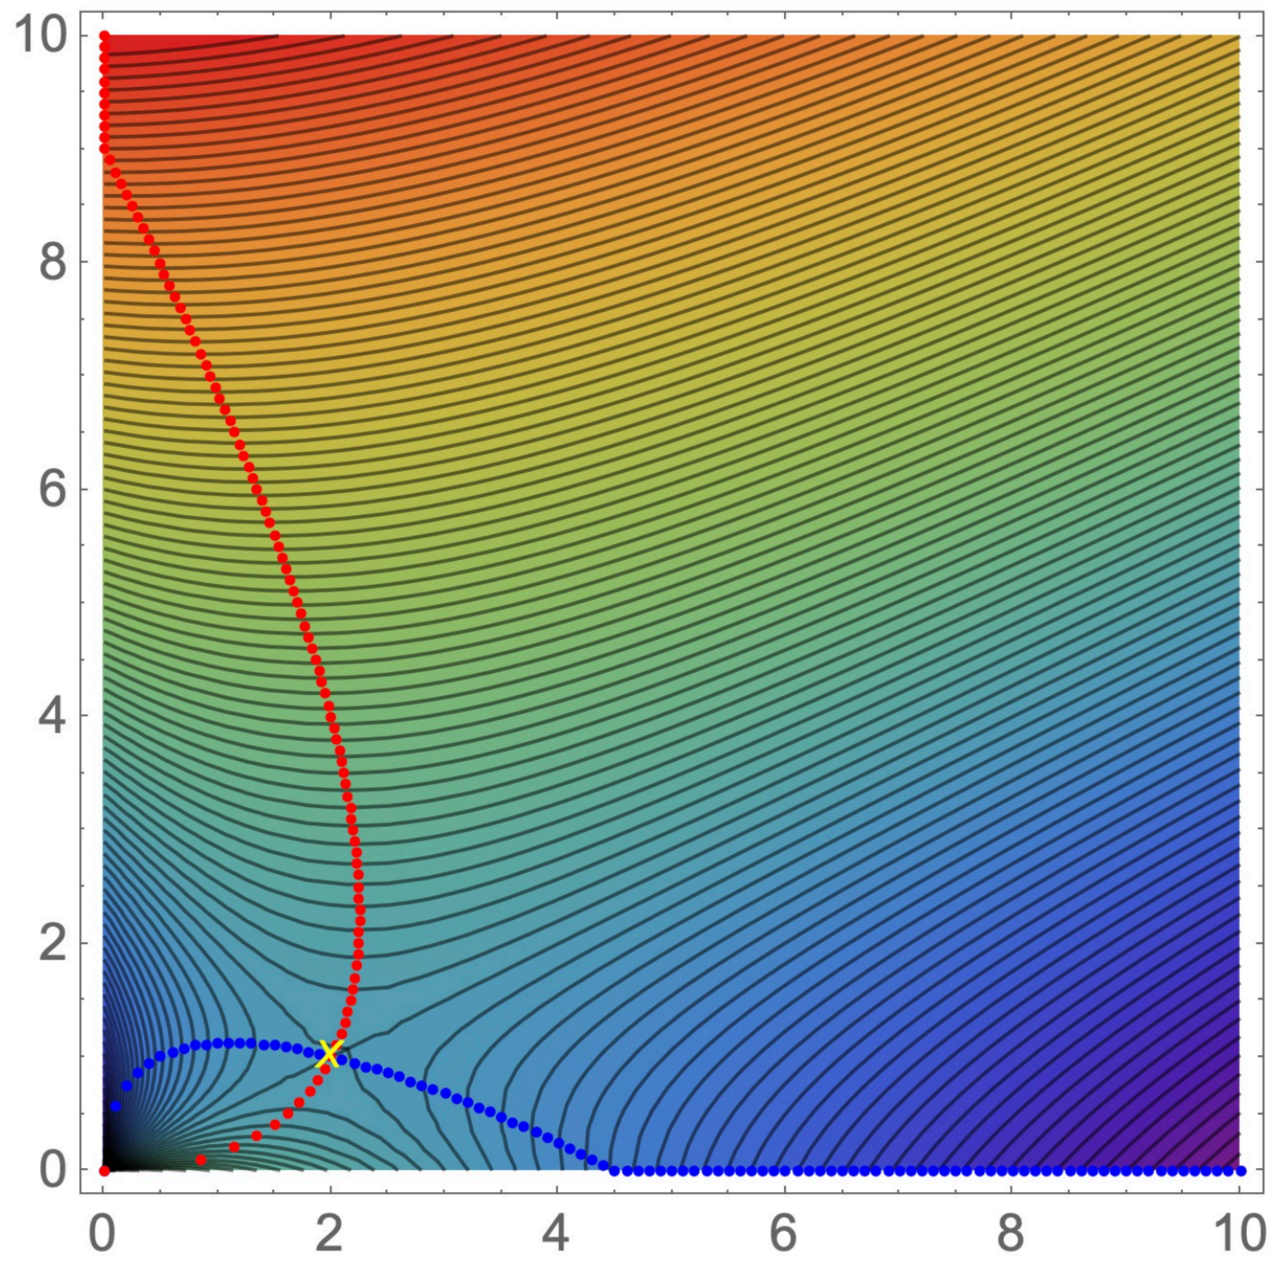

10
8
6
4
2
0
0
2
4
6
8
10

Modification 2: root-reward trees

A root-reward tree is

a finite tree $G=(V,E)$,

with boundary B = set of leaves,

and $f : B \to \{0,1\}$, $f = \mathbb{1}_r$

where r is a distinguished leaf called the root.

KEY FACT: the Peres-Šunić path decomposition of a λ-biased ∞-harmonic function on a ROOT-REWARD tree is INDEPENDENT of $\lambda \in (0,\infty)$.

Root-reward tree ...... and its 'essence' tree.

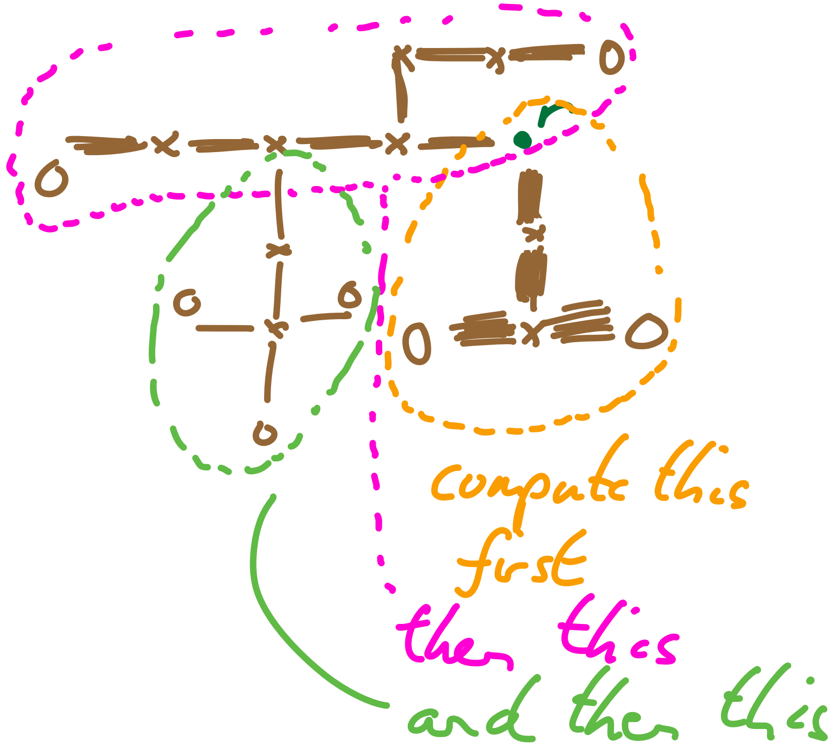

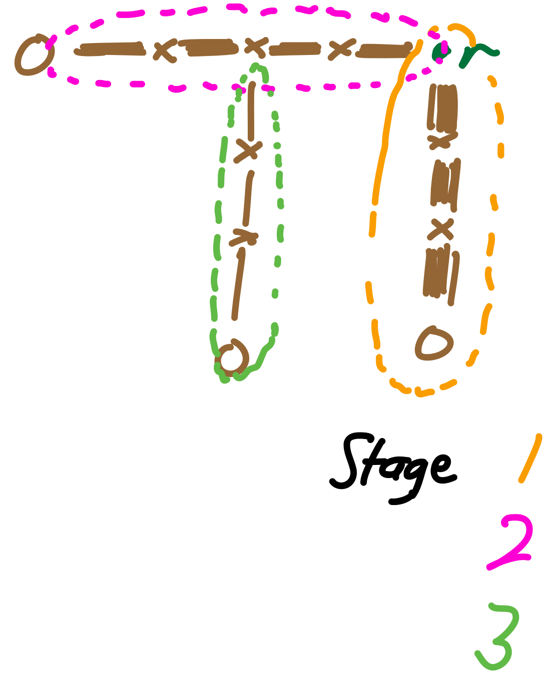

This decomposition is determined by subtree diameters and is independent of $\lambda \in (0, \infty)$.

Game value and Nash equilibria for the leisurely stake game on root-reward trees

Theorem [H, Pete]

On any root-reward tree, and for any compact $K \subset (0,\infty)$, there is $\varepsilon_K > 0$ such that, for any relative fortune $\lambda \in K$, and for $0 < \varepsilon < \varepsilon_K$, in the leisurely game with parameter ε starting at any vertex $v \in V \setminus B$,

game value equals $h(\lambda, v)$;

every Nash equilibrium consists of $h(\lambda,\cdot)$-maximizing/minimizing moves;

and the stakes are $(S, \lambda S)$, with

$$S = \frac{\varepsilon \Delta(\lambda, v)}{(\lambda+1)^2 \frac{d}{d\lambda} h(\lambda, v)} = \frac{\Delta(\lambda, v)}{\mathbb{E}\sum_{i=0}^{\tau-1} \Delta(\lambda, X_i)}, \quad X_0 = v.$$

Caveat:

The payment Mina → Maxine when play NEVER ENDS must be the FULL AMOUNT $1 in this theorem.

Otherwise, Mina can fight Maxine's supposedly optimal play on a line graph

(Mina stakes much of her fortune

(Mina stakes nothing here

By so doing, Mina can force infinite play.

So the FULL PAYMENT

$$\text{Mina} \xrightarrow{\$1} \text{Maxine}$$

for infinite games is needed to discourage Mina from doing this.

Self-funded stake-governed tug-of-war.

In the game that we have considered, FORTUNES were limited, making them a PRECIOUS resource that players must spend — and conserve — over the likely lifetime of the game.

What about another means of making a resource PRECIOUS?

Both PLAYERS can spend whatever they please at any given turn, but now it is THEIR OWN MONEY that they spend — so that the stakes of a given player sum up to act as a running cost, to be deducted from the terminal payment.

Self-funded stake-governed tug-of-war,
on finite line graphs and on $\mathbb{Z}$,
is the subject of

'On the Trail of Lost Pennies',

arXiv: 2209.07451

— non-unique Nash equilibria on L_6

— countably many equilibria on $\mathbb{Z}$

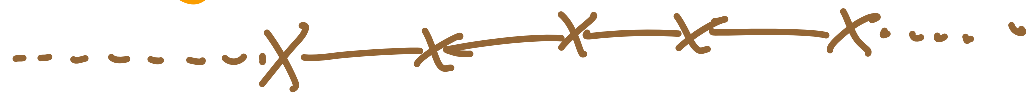

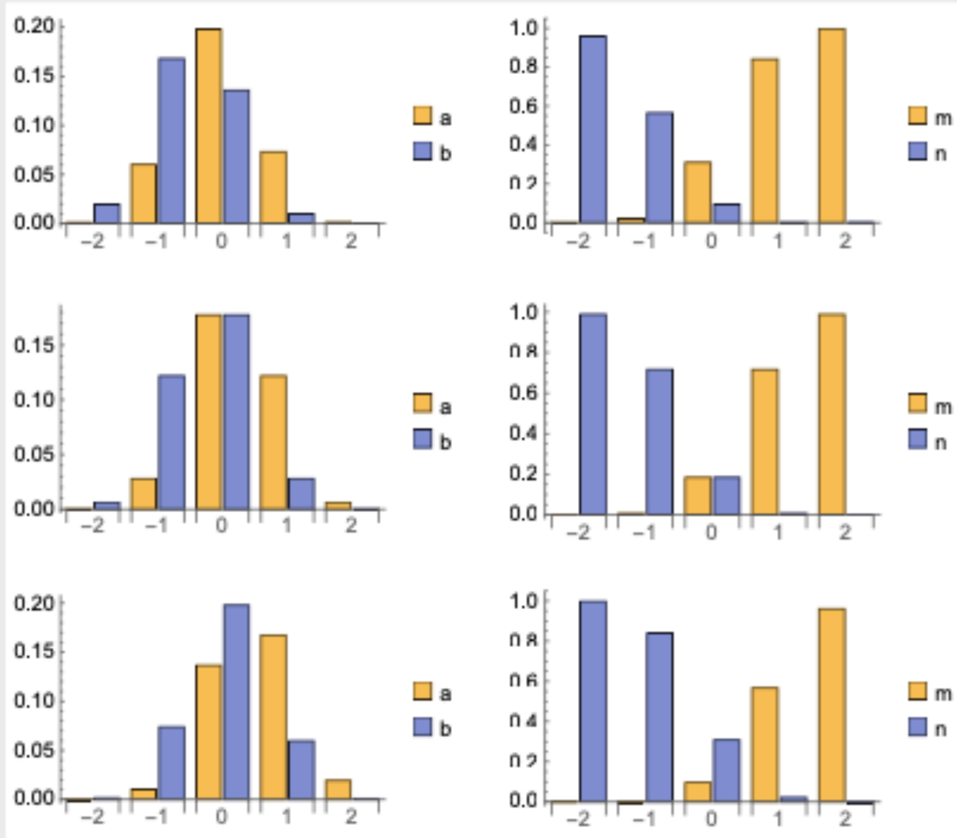
a
b
m
n
0.20
0.15
0.10
0.05
0.00
1.0
0.8
0.6
0.4
0.2
0.0
−2
−1
0
1
2

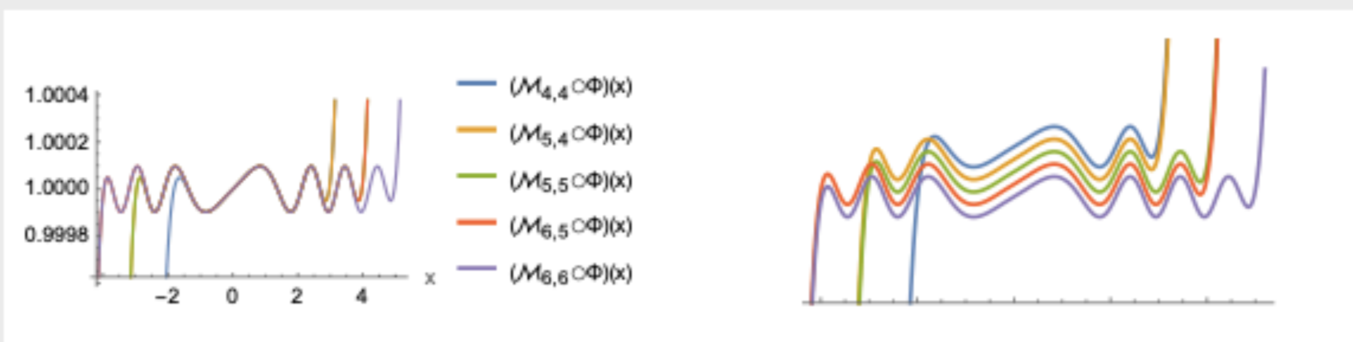
1.0004
1.0002
1.0000
0.9998
x
$(\mathcal{M}_{4,4}\circ\Phi)(x)$
$(\mathcal{M}_{5,4}\circ\Phi)(x)$
$(\mathcal{M}_{5,5}\circ\Phi)(x)$
$(\mathcal{M}_{6,5}\circ\Phi)(x)$
$(\mathcal{M}_{6,6}\circ\Phi)(x)$


Picasso