Structured Coordination Games in the Plane

John McAlister
jmcalis6@vols.utk.edu

University of Tennessee, Knoxville Tennessee

February 11th

THE UNIVERSITY OF TENNESSEE KNOXVILLE

DEPARTMENT OF MATHEMATICS

Overview

Game Theoretic Background

For a game with player set V, strategy set C, and payoff functions π_v, here are some definitions which will be important

- A *strategy profile* $u : V \to C$
- A player's *payoff*, $\pi_v(c|u)$ is a measure of how advantageous strategy c is against the strategy profile u for player v
- A player's *best response* to a strategy profile u is $BR_v(u) = \operatorname{argmax}_{c \in C} \pi_v(c|u)$
- A strategy profile, u, is a *Nash equilibrium* is every player is playing a best response to u.

Introduction to the Coordination Game

"Coordination Game" Refers to a wide range of game theoretic interactions which satisfy the *Bandwagon Property*

Bandwagon Property (Kandori et al. 1993)

If the "support" of u is the set of all strategies which are used in u, denoted by $C(u)$ then a game satisfies the Bandwagon property if and only if

$$BR_v(u) \subset C(u)$$

for all $v \in V$.

The easiest example is when the payoff matrix is the identity matrix.

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

Introduction to the Coordination Game

The two player pure coordination game is very simple. If both players play the same strategy they each get a payoff of 1 and if both play different strategies they get a payoff of 0.

Two player coordination

For the strategy profile $u = \{u_1, u_2\}$ if each pure strategy is considered as a standard basis element of $\mathbb{R}^m$ then the payoff function is

$$\pi_1(u_1|u_2) = \pi_2(u_2|u_1) = \langle u_1, u_2 \rangle$$

The two player coordination game is solved, but the multiplayer coordination game has a lot of rich dynamics

Introduction to the Coordination Game

For the Structured Coordination Game played on the graph G with adjacency matrix W, the payoff function is

$$\pi_v(u_v|u) = \sum_{w \in V} W_{v,w} \langle u_v, u_w \rangle$$

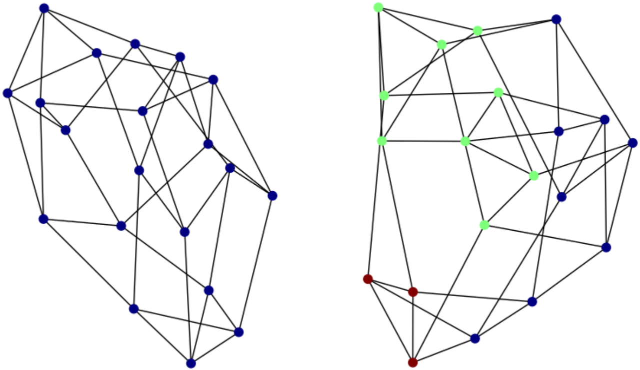

Introduction to the Coordination Game

Strategy Space

Let X be the set of all strategy profiles. If u and v are different strategies in X which differ by a single player changing their strategy they are connected by a single strategic change then u and v are adjacent in X. This adjacency induces a metric on X so we write the strategy space (X, d).

Reduced Strategy Space

Every strategy profile corresponds to a vertex partition. Call the vertex partition corresponding to u, Φu. If two strategy profiles u and v have $\Phi u = \Phi v$ then $u \sim_\Phi v$. Let $\mathcal{A}$ be the set of equivalence classes under $\sim_\Phi$. We construct a metric of $\mathcal{A}$ in the same way as for X.

$$\mathcal{A} = X/S_n$$

Introduction to the Coordination Game

The Coordination game is a potential game meaning there is a function $\mathcal{W}$ such that for two strategy profiles $u = (u_v, u_{-v})$ and $u' = (u'_v, u_{-v})$ then

$$\mathcal{W}(u) - \mathcal{W}(u') = w_v(u_v|u) - w_v(u'_v|u')$$

Lemma 2

$\mathcal{W}(u) = \frac{-1}{2}E(\Phi u)$ is a potential function for the pure coordination game where $E(P)$ is the size of the cutset for the partition P.

If $\Gamma_{\mathcal{A}}(u) = \{a \in \mathcal{A}; d(a, u) \leq 1\}$ then

Corollary 1

If $\mathcal{W}(u) \geq \mathcal{W}(u')$ for all $u' \in \Gamma_{\mathcal{A}}(u)$, then u is a Nash equilibrium.

Introduction to the Coordination Game

Summary

Nash equilibria of the pure coordination game on a graph G are local minimizers of the size of the cutesets they induce.

Structured Coordination in the Plane

The theory to study this game in the dual has been developed form many strategies but here I will discuss it only for 2 strategies. First recall the definition of dual graph

Let G be a 3-edge connected planar graph the dual G^* is a simple planar graph.

1. Every face in G is a vertex of G^*
2. An edge separating two faces in G corresponds to an edge between the two corresponding vertices in G^*
3. Thus, every vertex in G corresponds to a face of G^*.

Structured Coordination in the Plane

Cutspace

For a graph G the cutspace $T(G)$ is the set of all the cuts of the graph G (sets of edges that, if removed, would separate the graph into two parts). This is a vector space over $\mathbb{Z}_2$.

Cyclespace

For a graph G, the cyclespace $C(G)$ is the set of all subgraphs of the graph G in which every vertex has even degree (they admit Eulerian circuits). This is a vector space over $\mathbb{Z}_2$.

If G and G^* are dual to one another then $T(G)$ is isomorphic to $C(G^*)$.

Structured Coordination in the Plane

The cyclespace is a very nice space to imagine and to work with and so we would like to be able to describe the game dynamics in the cyclespace rather than in the strategy space. To do this we need a metric:

Metric on the Cyclespace

Let $s, r \in C(G)$ if $s = r + C_f$ for some face f, where C_f is the cycle of edges incident to the face f, then s and r are adjacent in $(C(G), D)$

It is our goal now to show that $(\mathcal{A}_G, d)$ is isometric to $(C(G^*), D)$.

Structured Coordination in the Plane

This isometry is not hard to prove when there are only two strategies. Note that $\Phi : \mathcal{A} \to T(G^*)$ is a bijection. Call the isomorphism between $T(G)$ and $C(G^*)$ Ψ and we can say the following:

Theorem 3

$\Psi \circ \Phi : \mathcal{A}_G \to C(G^*)$ is an isometry.

The proof requires only to note that if u and u' are a single strategic change away from one another, then there is an f such that $\Psi \circ \Phi u = \Psi \circ \Phi u' + C_f$.

Potential function from $C(G^*)$

Note that $\mathcal{W}(\Phi^{-1} \circ \Psi^{-1} s) = \frac{-1}{2} E(\Psi^{-1} s) = \frac{-1}{2} size(s)$ so minimizing size maximizes the potential function.

Structured Coordination in the Plane

Summary

s is a local minimizer of size in the $(C(G^*), D)$ if an only if $\Phi^{-1} \circ \Psi^{-1} s$ is a Nash equilibrium.

Example

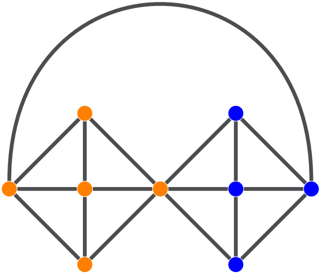

$\approx$

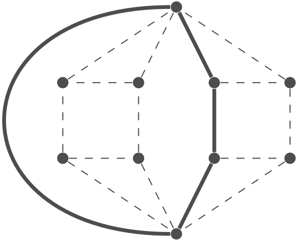

On the left is a Nash equilibrium strategy profile, u, on G, on the right is a locally minimal subgraph, s, in the cycle space of G^*. $\Psi \circ \Phi u = s$.

Dynamics in the Plane

We can also describe the game dynamics through Myopic Best Response

Myopic Best Response (MBR)

MBR is a strategy revision protocol under which every player takes on a best response to the previous strategy profile. For two strategies this is a deterministic process

$$u_v(t+1) = BR_v(u(t))$$

Dynamics in the Plane

We can define a flow through $C(G^*)$ which is equivalent

Cycle Shortening Flow

A sequence of subgraphs satisfies the cycle shortening flow if

$$s(t+1) = s(t) + \sum_{f \in F} H_f(s(t)) C_f$$

where $H_f(s(t)) = 1$ if at more than half of the edges incident to f are included in $s(t)$ and 0 otherwise.

Example

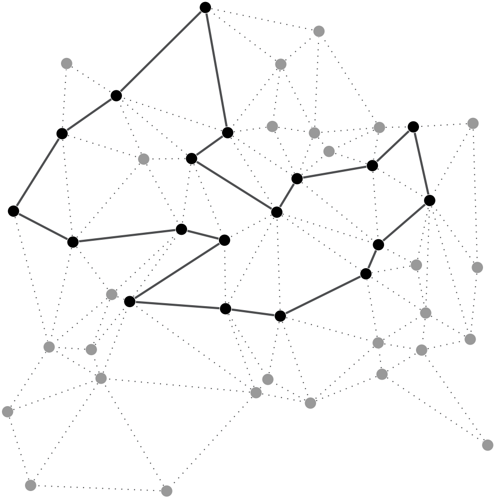

Example

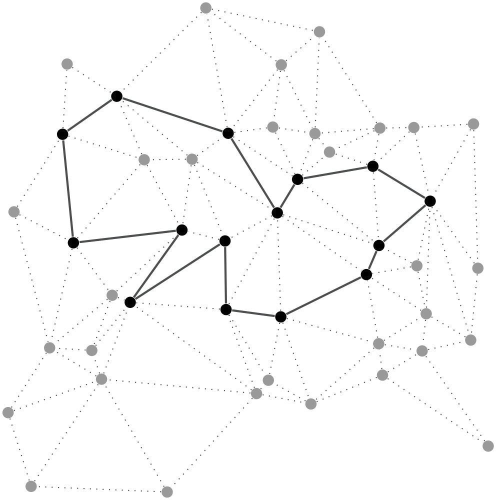

Example

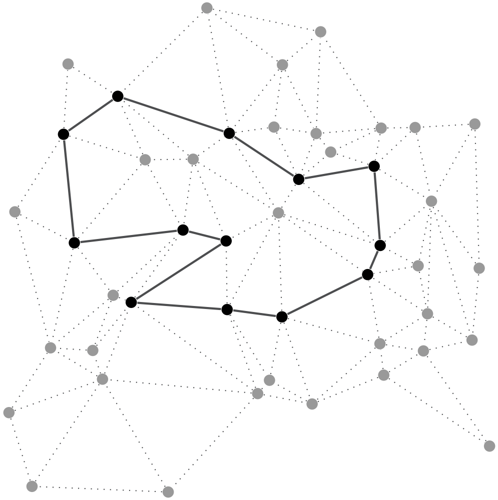

Example

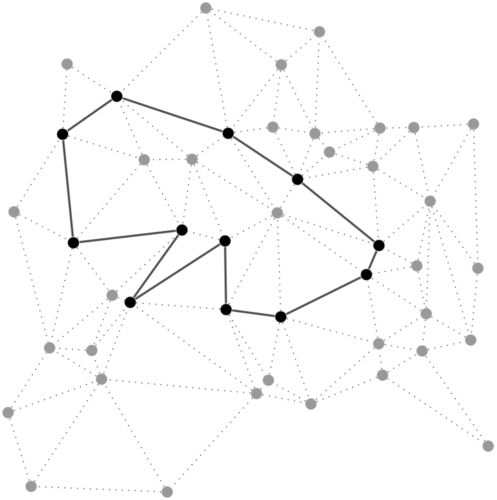

Example

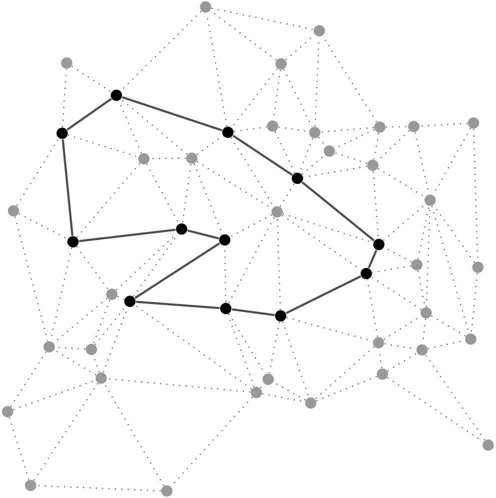

Dynamics in the Plane

Summary

Myopic best response in $\mathcal{A}_G$ is equivalent to the cycle shortening flow in $C(G^*)$.

Continuous Extension

This poses the question: What if do this same process in $\Omega \subset \mathbb{R}^2$.

- Every player is a point in Ω
- the strategy profile is $u : \Omega \to \mathbb{R}^m$ so that $u(x) = \hat{e}^i$ for some i.
- Every player interacts with those around it according to a integrable kernel $K(x, y) \in C_b^0(\Omega; L^1(\Omega))$
- The payoff that a player receives for playing a strategy $i \in C$ against u is $w_x(\hat{e}^i|u) = g^i[u](x) = \int_\Omega K(x, y)u^i(y)dy$
- $BR_x(u) = \mathrm{argmax}_{i\in\{1,\dots,m\}} g^i[u](x)$

Continuous Extension

We will start with the easy case: Assume that $u : \Omega \to \{[1,0],[0,1]\}$ is such that the strategic communities are separated by a smooth curve

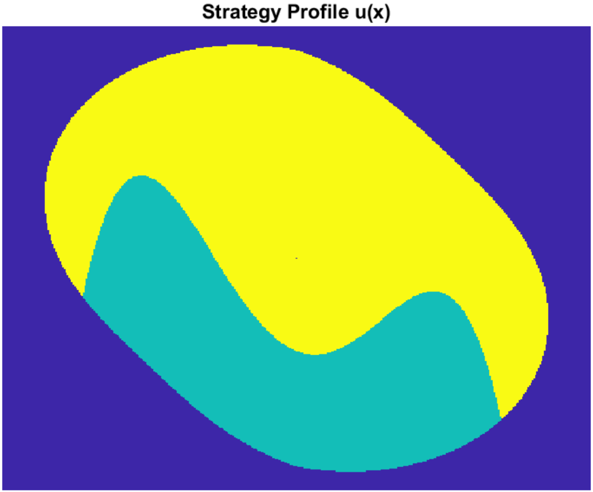

Continuous Extension

Compute the payoff for strategy 1 as $g^1[u](x) = \int_\Omega K(x,y)u^1(y)dy$. Here $K(x,y) = \chi_{B_\epsilon(x)}(y)$.

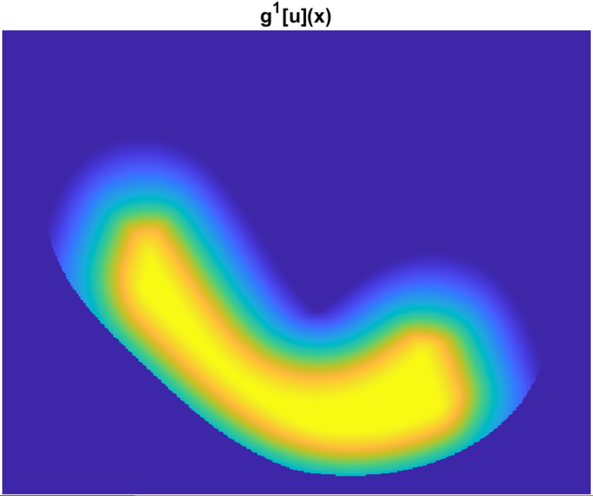

Continuous Extension

Compute the payoff for strategy 2 as $g^2[u](x) = \int_\Omega K(x,y)u^2(y)dy$. Here $K(x,y) = \chi_{B_\epsilon(x)}(y)$.

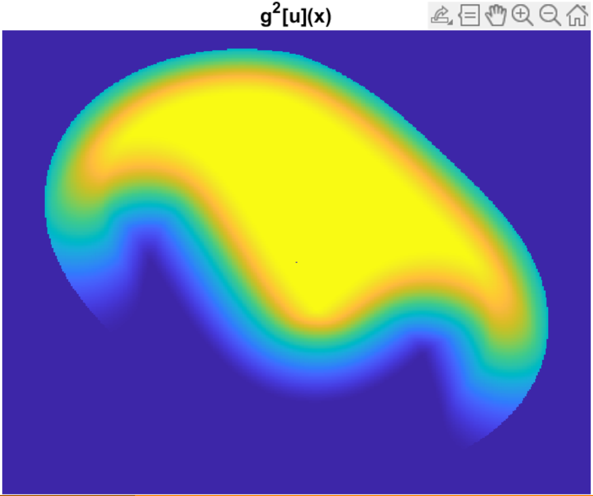

Continuous Extension

The players whose best response is strategy 1 are in the superlevel set $\{x \in \Omega; g^1[u](x) - g^2[u](x) > 0\}$

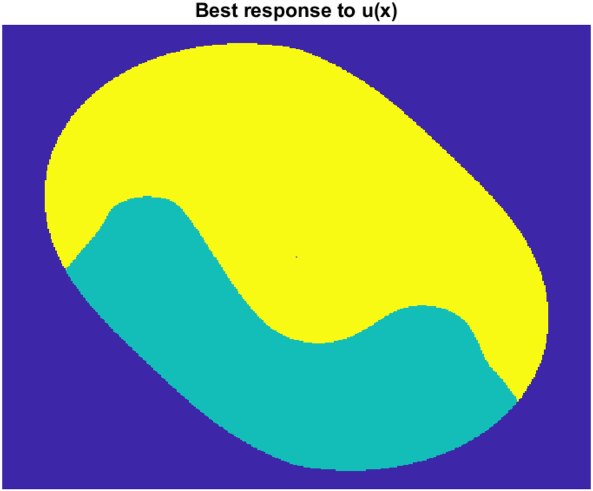

Continuous Extension

If we overlay the strategic boundary of the best response to u over the strategy profile u we see the boundary moves familiarly

Boundary of the Best Response to u(x) superimposed on u(x)

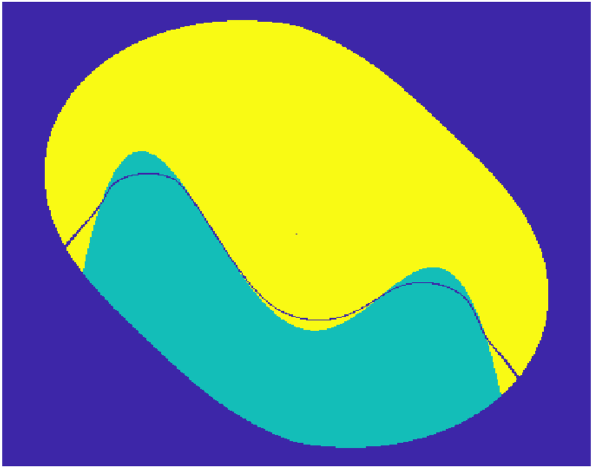

Continuous Extension

For every strategic profile u, $g^i[u]$ is uniformly continuous. This can give us some nice equilibrium results.

Lemma 9

If u is a Nash equilibrium strategy profile with multiple strategic communities C^1 and C^2, and if x is on the boundary of C^1 and C^2 then $g^1[u](x) = g^2[u](x)$

If we relax the assumptions on the strategy profile so that the strategic boundaries form a complex of curves(c^i) and junction points (p^i) where two curves do not come within ϵ of each other away from junction points we can get the following corollaries

Continuous Extension

Corollary 2

For a sufficiently well behaved network of strategic boundaries, if the detection radius is ϵ, then a boundary curve with endpoints *a* and *b* must lie in the following rectangle:

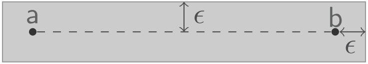

Corollary 3

In a Nash equilibrium, there are no closed isolated boundary curves.

Continuous Extension

The idea of this proof is like a maximum principle argument.

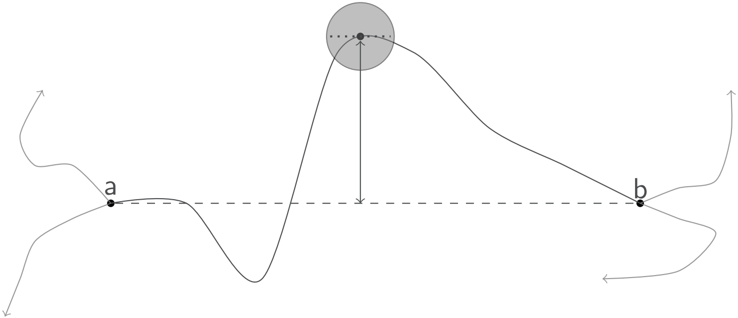

pick the point with the greatest orthogonal displacement from the line $\overline{ab}$. The curve must bisect it but the curve does not go above the diameter shown the the curve must be the diameter.

Continuous Extension

The idea of this proof is like a maximum principle argument.

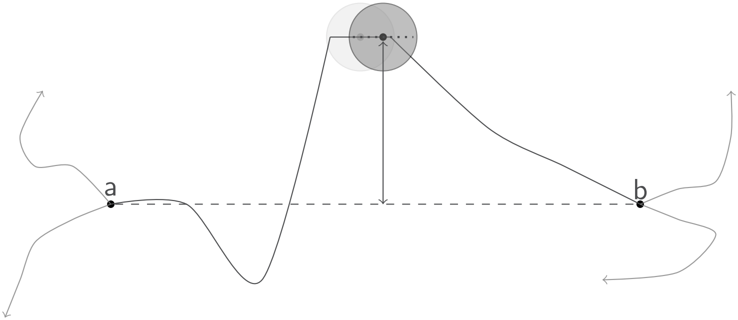

Repeat the same process on a point on the original diameter. Again the curve must be below the diameter parallel to $\overline{ab}$ and bisect the sensing area so it must be exactly that diameter.

Continuous Extension

The idea of this proof is like a maximum principle argument.

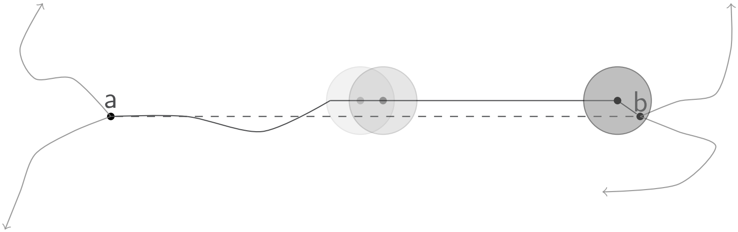

We can repeat this process until the sensing area includes point *a* or *b* at which point we cannot say anything. This process must terminate. We use this to prove the result.

Continuous Extension

What do the game dynamics look like under MBR?

Myopic Best Response

In this setting the Myopic Best response is quite the same. For the strategic profile u compute $g^i[u]$ for all $i \in C$. The region of Ω playing strategy 1 will be defined by a super level set.

$$g^1[u](x) - g^2[u](x) > 0$$

Likewise for strategy 2.

We want to consider how this changes the boundary. At any point on strategic boundary between strategies 1 and 2, consider the players in the normal direction.

Continuous Extension

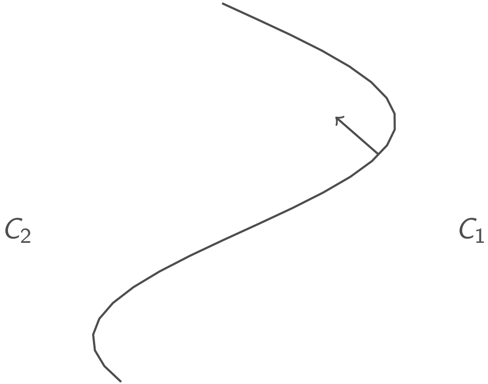

The boundary between strategic communities 1 and 2

Continuous Extension

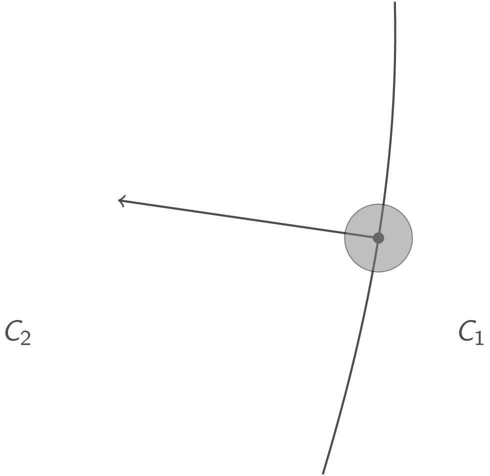

If $\kappa > 0$ then clearly this this player's best response is strategy 1.

Continuous Extension

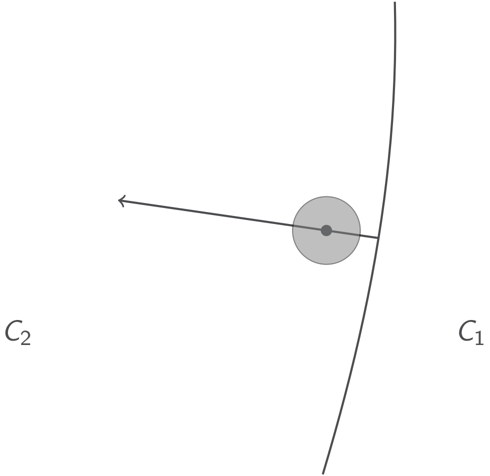

This this player's best response is strategy 2.

Continuous Extension

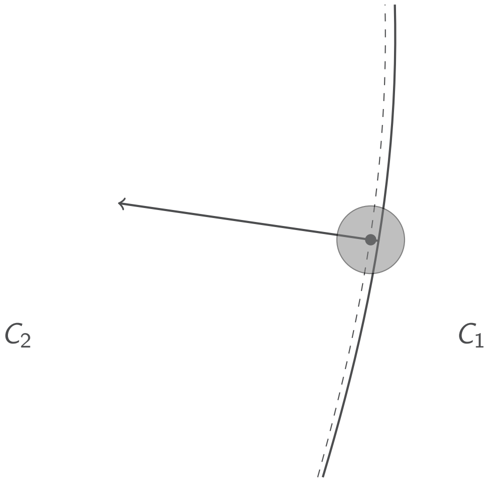

There must be a distance $d(s, x)$ at which a player's sensing area is bisected. This is where the boundary will be after one round of MBR.

Continuous Extension

The discrete time evolution equation is

$$s(t + \Delta t) - s(t) = d(s, x)\vec{n}$$

If we take the scaling the the time step is proportional to the sensing area $\Delta t = \epsilon^2$. Then in the limit as $\epsilon \to 0$ we see

$$\lim_{\Delta t \to 0} \frac{s(t + \delta t) - s(t)}{\Delta t} = \lim_{\epsilon \to 0} \frac{d(s, x)}{\epsilon^2}\vec{n}$$

Continuous Extension

Numerical experiments where the boundary is a circle with curvature κ show the following convergence

$$\lim_{\epsilon \to 0} \frac{d(s,x)}{\epsilon^2} = \frac{\kappa}{2\pi}$$

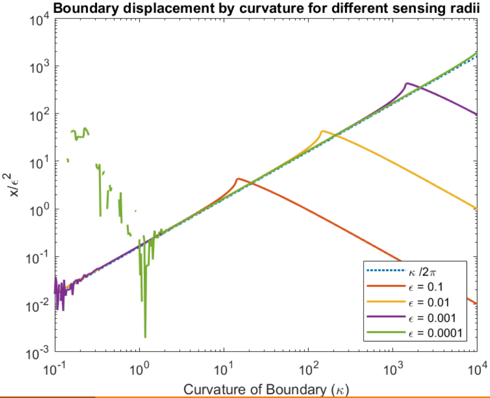

Continuous Extension

In more detail, consider the circle of radius R and the circle of radius ϵ a distance d from the edge of the larger circle. Let $A(R, \epsilon, d)$ be the area of intersection for those two circles

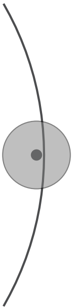

Continuous Extension

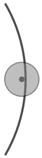

The distance at which a players sensing area is bisected by the curve $(d(s,x))$ should satisfy

$$\frac{A(R,\epsilon,d(s,x))}{\pi\epsilon^2} = \frac{1}{2}$$

And from the numerical results we will attempt to show that

$$\lim_{\epsilon\to 0} \frac{A(R,\epsilon,\frac{\epsilon^2}{2\pi R})}{\pi\epsilon^2} = \frac{1}{2}$$

Continuous Extension

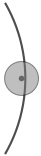

Asymptotic Analysis

$$\frac{A(R, \epsilon, \frac{\epsilon^2}{2\pi R})}{\pi \epsilon^2} = A_1(\epsilon, R) + A_2(\epsilon, R) + A_3(\epsilon, R)$$

First we can show that $A_1 \to \frac{1}{2}$ with elementary methods. A_2 and A_3 diverge but we can show that they diverge in the opposite as

$$A_2(\epsilon, R) \approx \frac{-\epsilon}{R}$$
$$A_3(\epsilon, R) \approx \frac{\epsilon}{R}$$

Continuous Extension

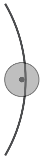

This is a way of backing into the limit

$$\lim_{\epsilon \to 0} \frac{d(s, x, \epsilon)}{\epsilon^2} = \frac{1}{2\pi R}$$

Which would get us a step closer to showing that

$$\frac{\partial s}{\partial t} = \lim_{\Delta t \to 0} \frac{s(t + \delta t) - s(t)}{\Delta t} = \lim_{\epsilon \to 0} \frac{d(s, x)}{\epsilon^2} \vec{n} = \frac{1}{2\pi} \kappa \vec{n}$$

Numerical Examples

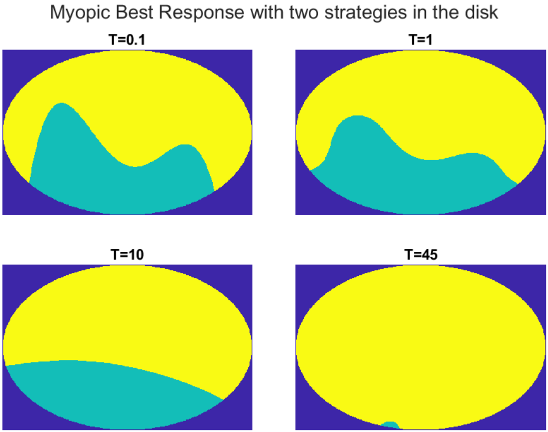

With a small enough sensing radius, the boundary curves evolve as we expect

Numerical Examples

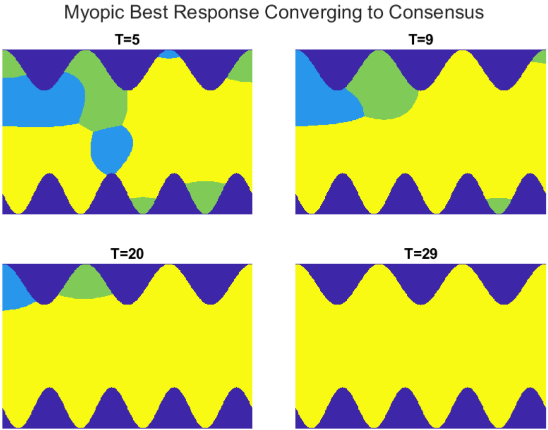

With three strategies the boundary curves evolve in a way similar to the network flow (except when two curves come withing ϵ of each other).

Numerical Examples

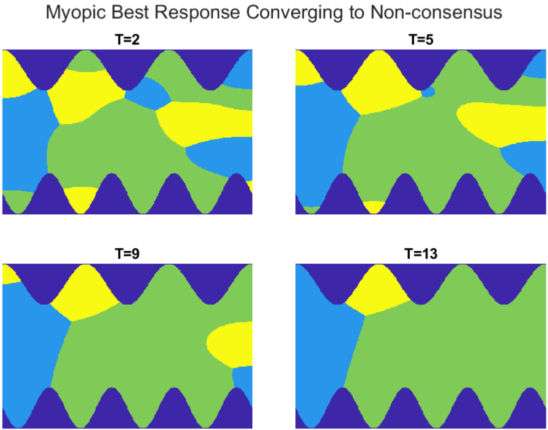

In non-convex domains it is easy to find stationary solutions which are *non-consensus equilibria*.

Conjecture

conjecture

Let $u_\epsilon(t)$ be a sequence of strategy profiles evolving under myopic best response with $supp(K(x,\cdot)) \subset B_\epsilon(x)$ and with time step Δt. Let $s_u^\epsilon(t)$ be the sequence of associated strategic boundary complexes. If $\Delta t = \epsilon^2 \to 0$ then $s_u^\epsilon(t) \to S(t)$ a smooth continuous one parameter family of curves (networks) which satisfy the free boundary curve shortening (network) flow.

Concerns

1. What if the strategy profile is not well behaved to begin with?
   The first step of the MBR has a "smoothening" property. Because $g^i[u]$ is uniformly continuous regardless of the characteristics of u
2. What happens when two boundary curves meet an annihilate or merge?
   This can happen for any positive sensing radius but with and arbitrarily small sensing radius, the curves must be arbitrarily close to one another to annihilate or merge (avoidance principle?).

Live Demo

I will now demonstrate the flows through the MATLAB code which is available on my github (github.com/jmcalis) and take any questions.