Repeated Games and a Peak Beyond Nash Equilibrium

February 10, 2017

Catherine Moon

Repeated Games?

- In systems of multiple self-interested agents, we cannot impose behavior on the agents
- Prisoner's Dilemma, studied in "Dominance" of Vince's Game Theory lecture

	confess	don't confess
confess	-2, -2	0, -3
don't confess	-3, 0	-1, -1

- Cooperation is impossible in one-shot version of this game

REPEATED GAMES?

- In systems of multiple self-interested agents, we cannot impose behavior on the agents
- Prisoner's Dilemma, studied in "Dominance" of Vince's Game Theory lecture

	confess	don't confess
confess	-2, -2	0, -3
don't confess	-3, 0	-1, -1

- Finitely repeated game: unraveling through backwards induction – cooperation is impossible as well

REPEATED GAMES?

- In systems of multiple self-interested agents, we cannot impose behavior on the agents
- Prisoner’s Dilemma, studied in “Dominance” of Vince’s Game Theory lecture

	confess	don’t confess
confess	-2, -2	0, -3
don’t confess	-3, 0	-1, -1

- Infinitely repeated: If I value future enough (discount factor δ), then cooperative action may be sustained
- Folk Theorem

Maximal Cooperation in Repeated Games on Social Networks (IJCAI-15)

Catherine Moon and Vincent Conitzer
Duke University

REPEATED GAME IN SOCIAL NETWORK

- **Common assumption**: an agent's behavior is instantly observable to all other agents
- What if there is a delay in knowledge propagation due to network structure?

- **Question**: Under what conditions can we still sustain cooperation, and can we compute the resulting equilibria?

MOTIVATING EXAMPLE

- Pollution reduction agreement

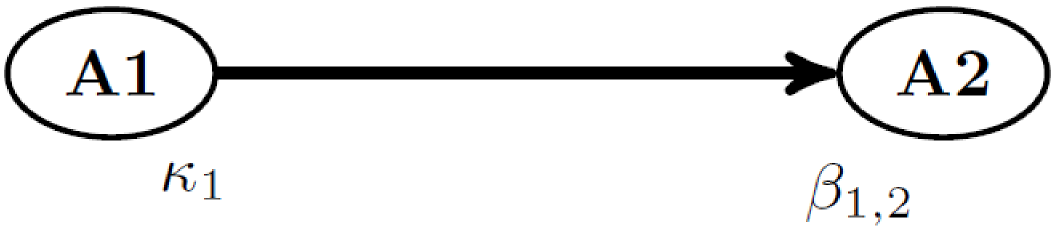

  - Directionality
  - Cooperation involving cost and benefit

- Set of cooperating agents

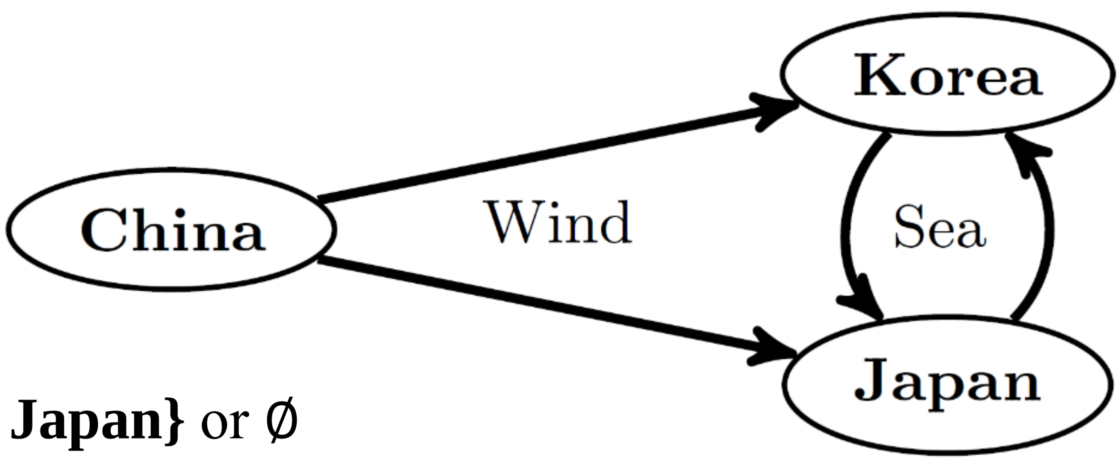

{Korea, Japan} or Ø

Model

- Set of agents A
- Directed graph $G = (V, E)$
- For all $i \in V$, cooperation cost κ_i
- For all $(i, j) \in E$, cooperation benefit $\beta_{i,j}$
- $x_{i,t} \in \{0,1\}$
- Defection is irreversible
- Agent i's round t payoff:

$$-x_{i,t}\kappa_i + \sum_{j:(j,i)\in E} x_{j,t}\beta_{j,i}$$

- Each period's payoff is discounted by discount factor δ

OUTLINE

- Analysis in Nash Equilibrium
- Algorithm for finding the unique maximal set of cooperating agents
- Experimental analysis on random graphs and observation of phase transition
- Extending equilibrium notions: Credibility of threats and Credible Equilibrium

DEFINITIONS

Let $S \subseteq A$ for a given network G

- Player i's strategy **grim trigger for *S*** consists of i cooperating until some player $j \in S$ with $(j, i) \in E$ defects, which is followed by i defecting (forever)
- The **grim trigger profile** $T[S]$ consists of all the agents in S playing the grim trigger strategy for S, and all other players always defecting
- For any subset S of the agents and any $i, j \in S$, the **distance from *i* to *j* through *S***, denoted $d(i, j, S)$, is the length of the shortest path from i to j that uses only agents in S. For a set of agents $G \subseteq S$, $d(G, j, S) = \min_{i \in G} d(i, j, S)$

THEORETICAL ANALYSIS: EQUILIBRIUM

Let $S \subseteq A$ for a given network G

Proposition 1: $T[S]$ is an equilibrium if and only if $\forall i \in S$

$$\sum_{j \in S:(j,i) \in E} \delta^{d(i,j,S)} \beta_{j,i} \geq \kappa_i$$

Sketch:

- Every player outside S is best-responding
- For $i \in S$, cooperate if

total utility loss from neighbors eventually defecting

$\geq$

total utility gain from reduced effort

GRIM TRIGGER IS WLOG

Proposition 2: Suppose there exists a pure-strategy equilibrium in which S is the set of players that cooperates forever. Then $T[S]$ is also an equilibrium

Sketch:

- For a given equilibrium, consider some time period τ at which every player outside S has defected (on the path of play)
- If player i considers defecting at this point,

  total utility loss from neighbors eventually defecting is

  at most that from everyone in S playing $T[S]$

Monotonicity

Let $S, S' \subseteq A$ for a given network G

Lemma 3: If $S \subseteq S'$ and the incentive constraint from Proposition 1 holds for i relative to S, then it also holds for i relative to S'.

Sketch:

- We argue that

$$\sum_{j \in S' : (j,i) \in E} \delta^{d(i,j,S')} \beta_{j,i} \geq \sum_{j \in S : (j,i) \in E} \delta^{d(i,j,S)} \beta_{j,i}$$

1. All summands are nonnegative
2. For any i, j, $d(i,j,S') \leq d(i,j,S)$
3. Because $\delta < 1$, $\delta^{d(i,j,S')} \geq \delta^{d(i,j,S)}$

Maximality

Let $S, S' \subseteq A$ for a given network G

Proposition 4: If $T[S]$ and $T[S']$ are both equilibria, then so is $T[S \cup S']$

Sketch:

- Consider some $i \in S \cup S'$; WLOG, suppose $i \in S$
- Need to show: incentive constraint from Proposition 1 holds for i relative to $S \cup S'$
- Follows from Lemma 3, $T[S]$ being an equilibrium, and $S \subseteq S \cup S'$

Algorithm 1

```
D_elimination ← true
while D_elimination = true do
    L_defectors ← ∅
    I ← IncomingEdges(E)
    O ← OutgoingEdges(E)
    for all i ∈ S^current do
        L^i ← ShortestPath(i, j, E) (∀j ∈ S^current)
        C^i ← IncentiveConstraint(L^i, κ_i, {β_{j,i}}_{j:(j,i)∈E}, δ)
        if C^i = false then
            add i to L_defectors
        end if
    end for
    for all i ∈ L_defectors do
        remove (i, j) from E for all j
        remove i from S^current
    end for
end while
```

IncentiveConstraint($L^i, \kappa_i, \{\beta_{j,i}\}_{j:(j,i)\in E}, \delta$) checks the incentive constraint inequality in Proposition 1. If it is satisfied, it returns true, indicating i's willingness to cooperate.

Experimental Analysis

Additional assumptions:

- For all $i \in S$, $\kappa_i = \sum_{i \in S:(i,j) \in E} 1$
- For all $(j,i) \in E$, $\beta_{j,i} = \beta$

Random graph models:

- Erdős–Rényi (ER)
- Barabási–Albert preferential-attachment (PA)

Fine Line Between Complete Cooperation and Defection

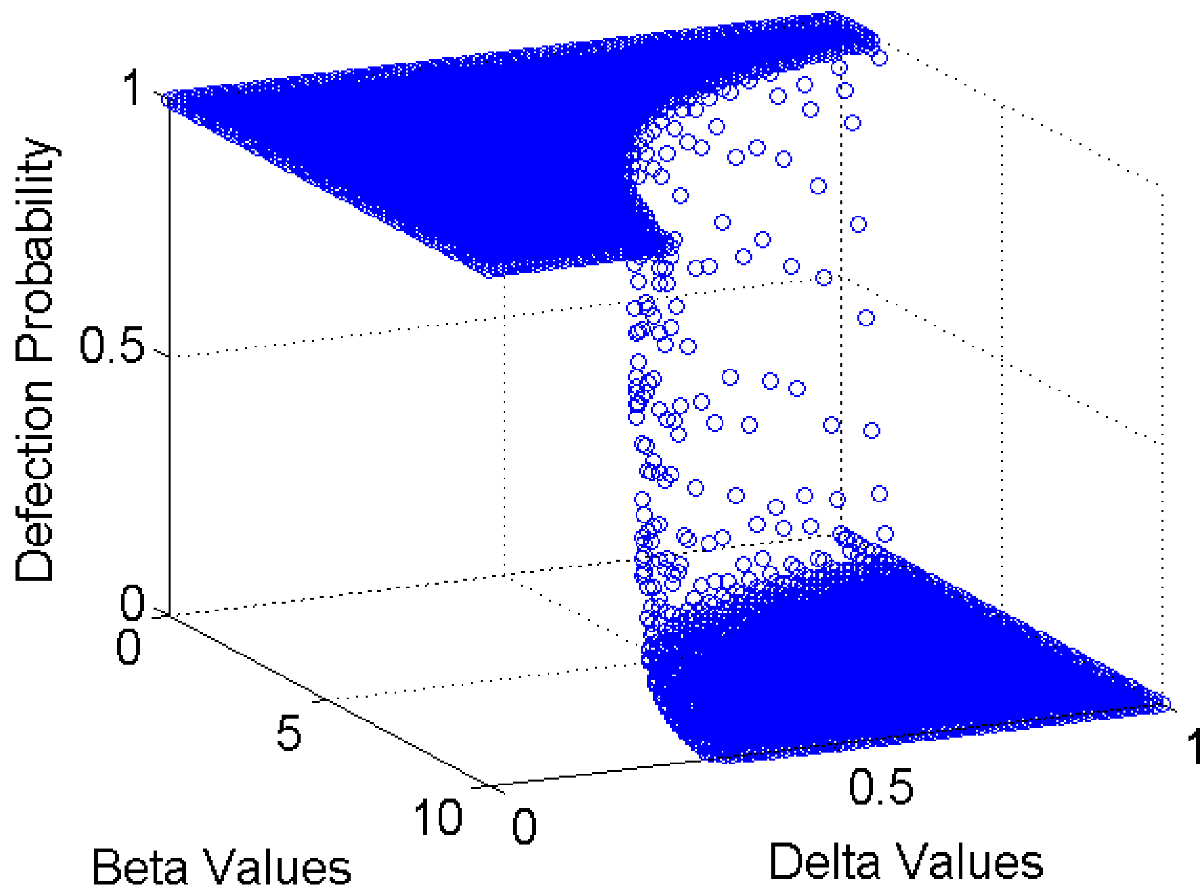

Simulation and Phase Transition

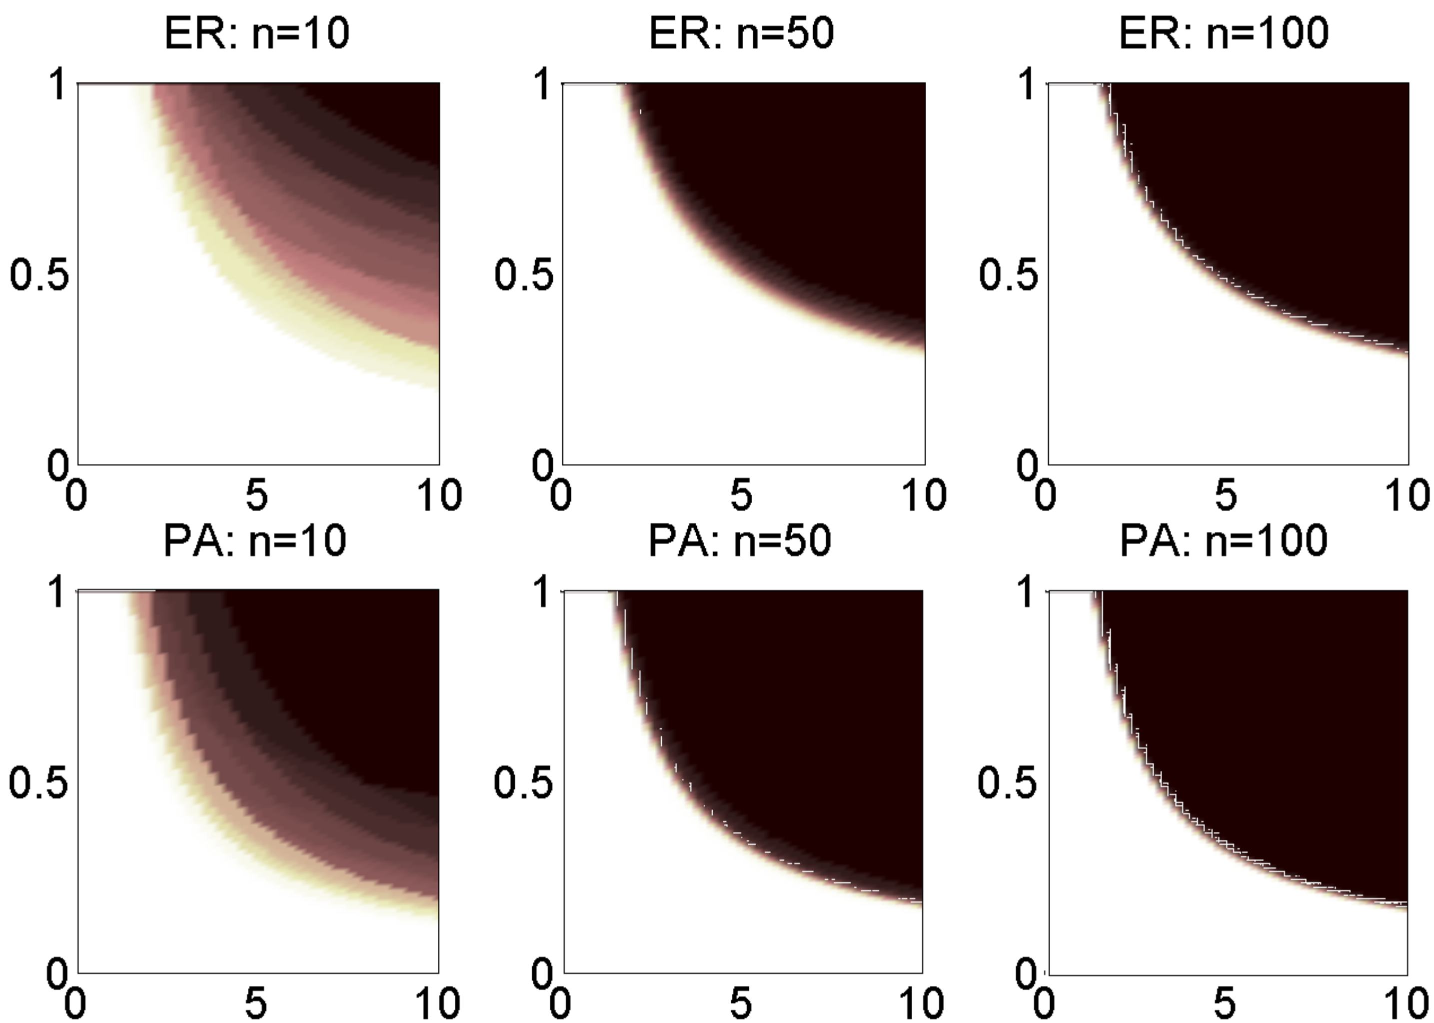

Agents’ defection probabilities for different values of β and δ

SIMULATION AND PHASE TRANSITION

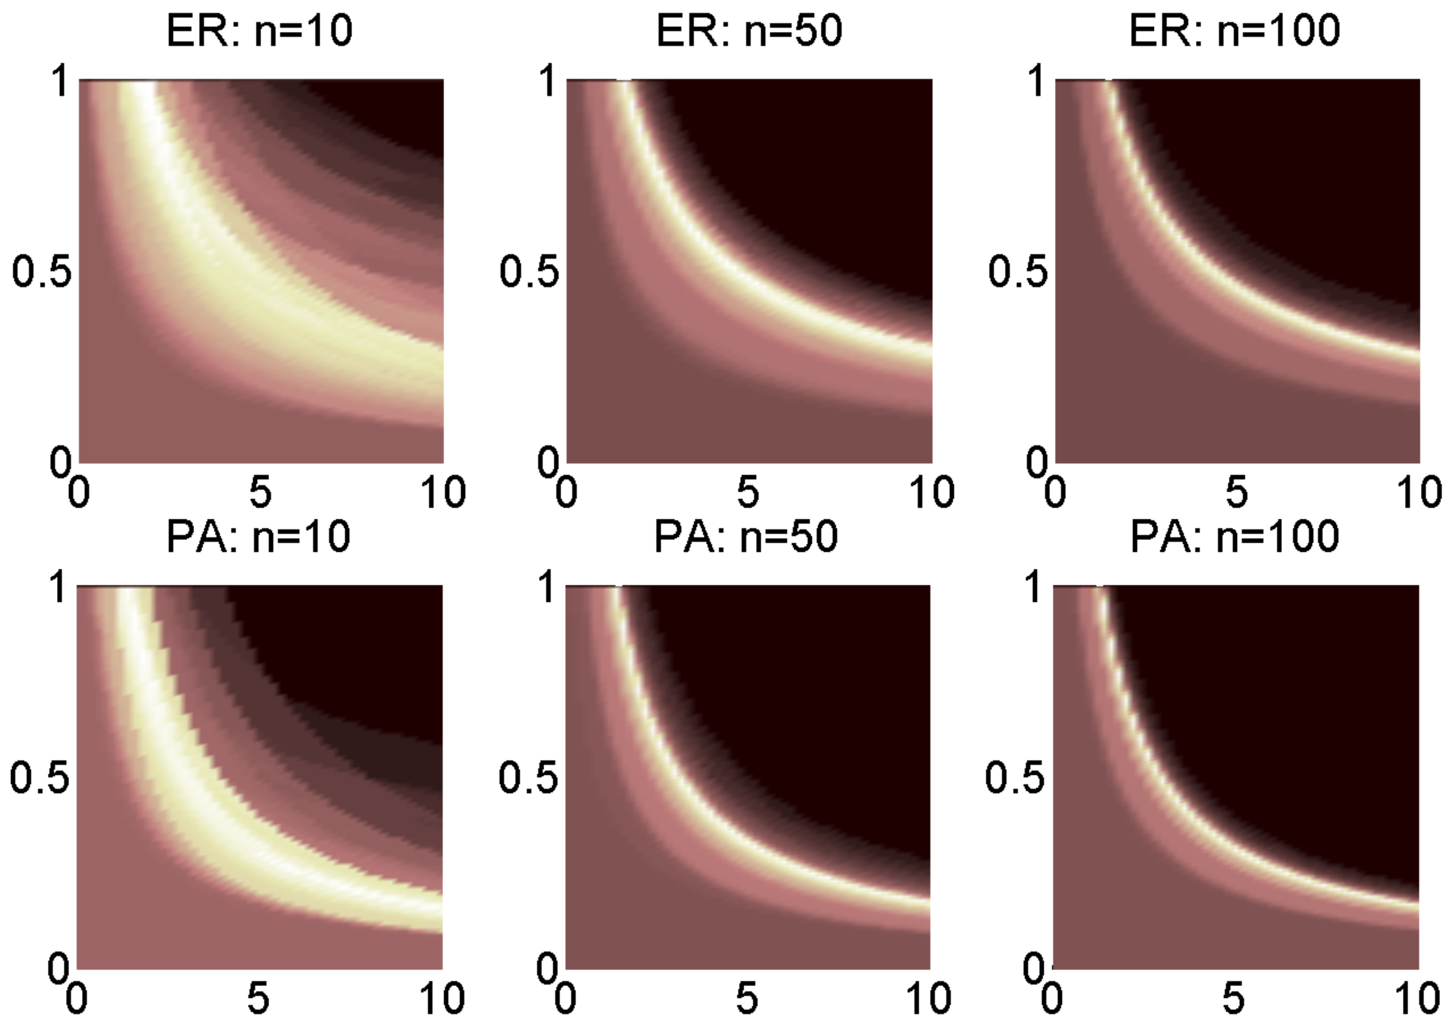

Average number of iterations needed until convergence for different values of β and δ

Credibility of Threats in Equilibrium

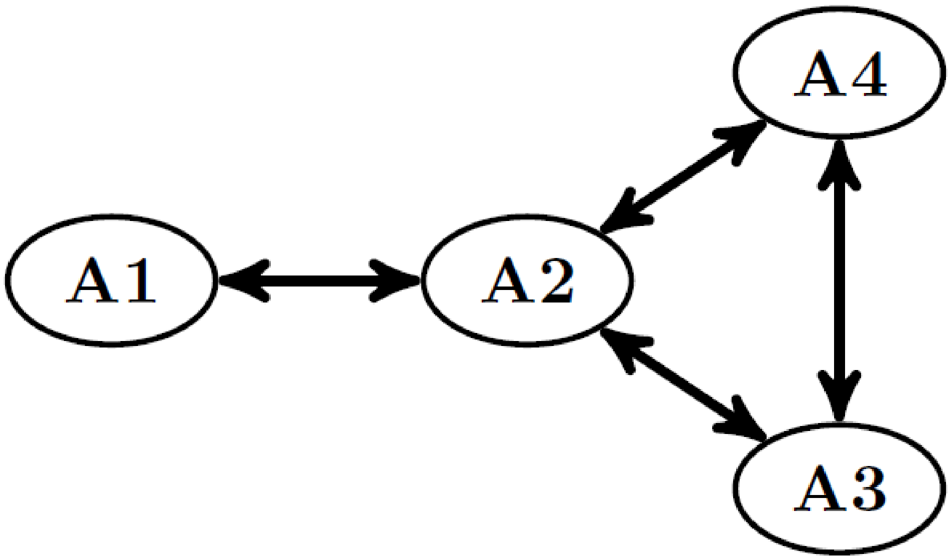

- The maximal set of cooperating agents, S^*, in Nash Equilibrium may involve threats of grim-trigger defection that are not credible

CREDIBLE EQUILIBRIUM

- Equilibrium extensions beyond NE:

  Subgame-perfect? Perfect Bayesian? Sequential?

- **Credible equilibrium**: if an agent learns that some deviation from the equilibrium has taken place, then she will be best off following her equilibrium strategy *regardless of her beliefs about which deviations have taken place*, assuming others also follow their equilibrium strategies from this round on

CREDIBLE EQUILIBRIUM

Lemma 7: Sufficient to check singleton deviations

Lemma 8: Sufficient to check one-round postponement of punishment

Theorem 9: Suppose that for some set $S \subseteq A$, $T[S]$ is a Nash equilibrium. Then $T[S]$ is a credible equilibrium if and only if for any $k, i \in S$ with $(k, i) \in E$, it holds that $\kappa_i - \delta^{d(k,j,S)} \sum_{j \in D} \beta_{ji} \geq 0$, where

$$D = \{j \in I: d(k,j,S) + 1 \leq d(k,j,S_{-i})\}$$

Credibility of Maximal Eq

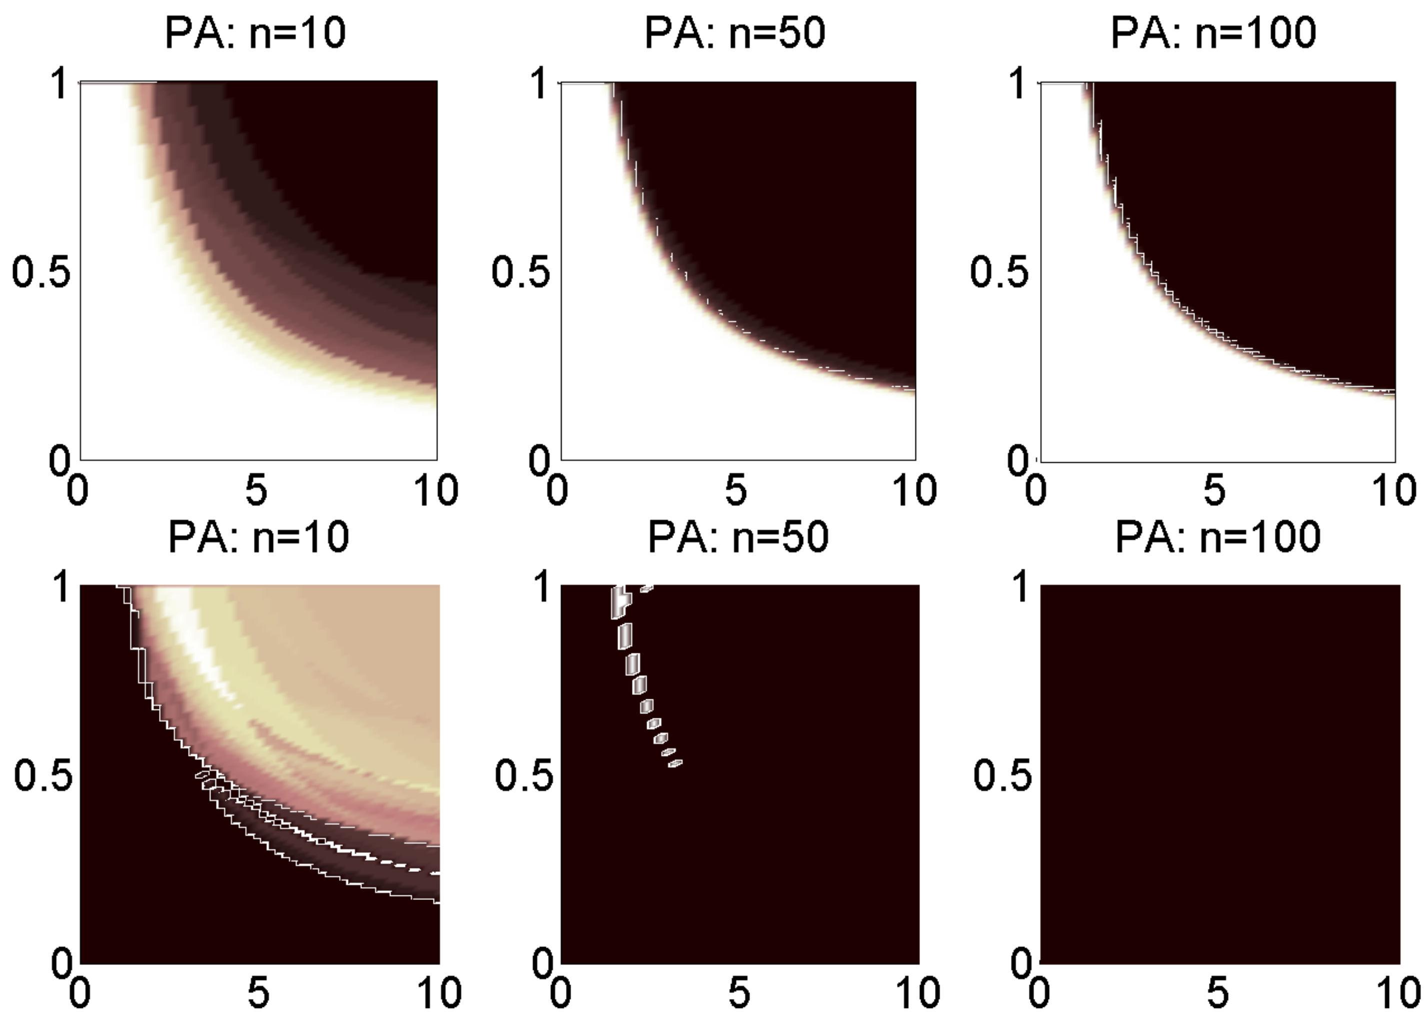

Top: Gradient graph showing agents' defection probabilities;
Bottom: gradient graph showing the fraction of cases where T[S*] is a CE

Thank You!