Evolutionary Game Theory, Cultural Modeling, and Third-Party Punishment

Dana Nau

Department of Computer Science
and Institute for Systems Research

University of Maryland

Work done jointly with

- Patrick Roos – one of my PhD graduates
- Michele Gelfand – Dept. of Psychology, U. of Maryland

Introduction

- My main field of research is Artificial Intelligence
  - But a lot of my work has been interdisciplinary

- I'll talk about some collaborative research with
  - Patrick Roos – one of my PhD graduates
  - Michele Gelfand – Dept. of Psychology, U. of Maryland

- Application of evolutionary game theory in cultural psychology
- Example of how an interdisciplinary team can accomplish things that none of us could have done individually

Motivation

- Suppose you're in a store and you see someone shoplifting. What do you do?
  - Try to ignore the incident?
  - Confront the shoplifter?
  - Report it to someone who works in the store?

Motivation

- Suppose you're in a store and you see someone shoplifting. What do you do?
  - Try to ignore the incident?
  - Confront the shoplifter?
  - Report it to someone who works in the store?

- Would you act differently if—
  - it's a big department store in a city that you only visit rarely?
  - it's a small "mom and pop" store in your neighborhood?

Third-Party Punishment (3PP)

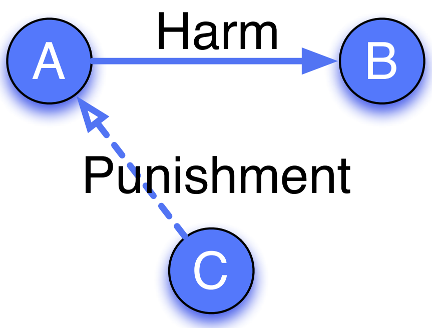

- Individual C (uninvolved third party) punishes individual A for harm that A has caused to B
  - Puzzling both psychologically and game-theoretically
  - Involves a cost to C, and gives C no direct benefit

- Empirical evidence that humans (and some animals) do it
  - Fehr 2003, 2004; Raihani, 2010
- Other empirical studies in which it did not occur
  - Pederson 2013

Objectives and Outline

- Basic Questions
  - What conditions foster the existence of 3PP?
  - How do the dynamics of 3PP relate to other cultural characteristics?

- Investigate these questions using evolutionary game theory

- Outline
  - Background on evolutionary game theory and cultural evolution
  - Our model
  - Results
  - Discussion

Background: Evolutionary Game Theory

- Application of game theory to evolving populations
    - J. Maynard Smith, 1973. The Logic of Animal Conflict. *Nature*.
- Game-theoretic strategies ⇔ different species
    - Each strategy is used by some proportion of the entire population
- Each individual's reproductive success depends on both its strategy and the strategies of others
    - Influences the proportion of each strategy at the next generation

Evolutionary Game Theory and Cultural Evolution

- Use EGT to model evolution of cultural characteristics
    - Axelrod 1986, Binmore & Samuelson, 1994; Ostrum, 2000; Bicchieri, 2006; Chalub *et al*., 2006; Kendal *et al*., 2006; Enquist & Ghirlanda, 2007; Enquist *et al*., 2008
  - Game-theoretic strategies ⇔ possible behaviors
  - Reproduction ⇔ cultural transmission
    - Humans imitate others, learn from others
    - Successful strategies have a higher probability of being adopted by others

Evolutionary Dynamics

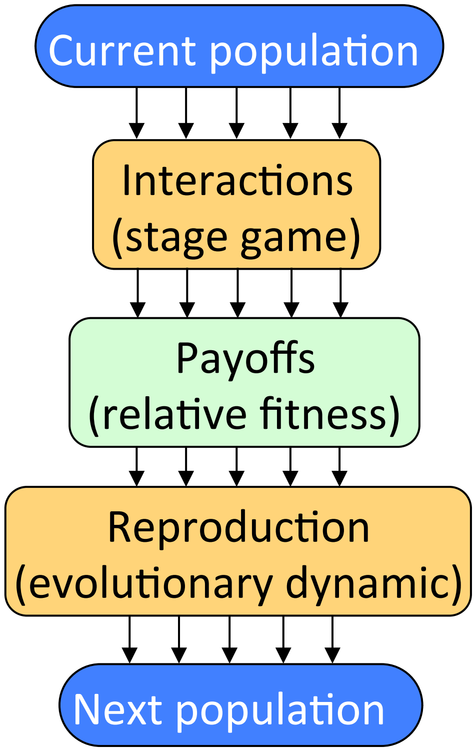

- Interpret game-theoretic payoffs as reproductive fitness
- Several ways; I'll discuss the best one for our purposes
    - Blume 1993; Szabóke 1998; Traulsen 2006, 2007; Roca 2009; Hilbe 2012; Zhang 2012
- Sequence of **stages** / **iterations** / **generations**
    - Game-theoretic interaction at each stage
- **Fermi rule** (from statistical mechanics) to propagate strategies to next iteration
    - Each individual compares its payoff to that of a randomly chosen neighbor
    - Pr[switch to neighbor's strategy] = $1/(1 + e^{s(\pi - \pi')})$
        - π, π' = individual's and neighbor's payoffs
        - $s \geq 0$ is the **selection strength**

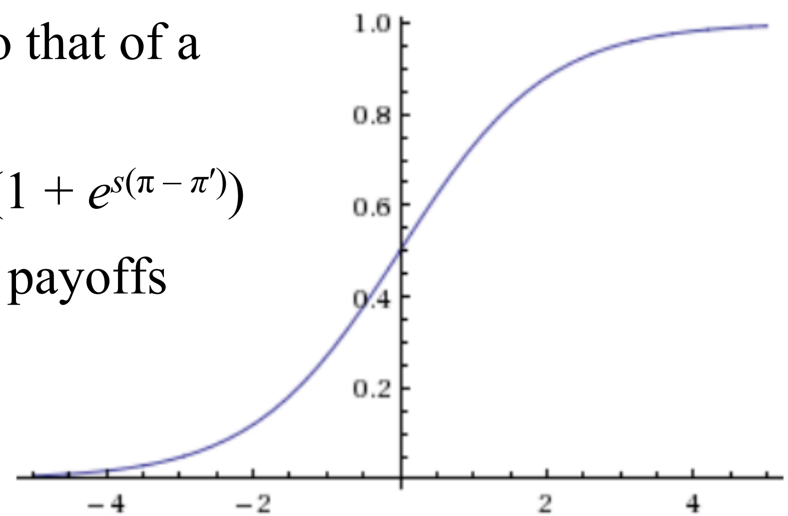

Mutation

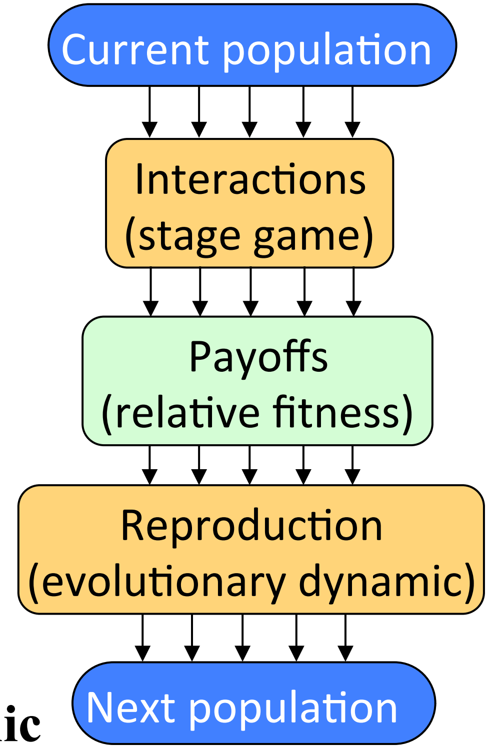

- In biological reproduction, mutation is relatively rare
  - Game-theoretic models often omit it
- In cultural evolution, something analogous to mutation happens more frequently:
  - Individuals try out new behaviors at random
    - Traulsen *et al*, 2009. Exploration dynamics in evolutionary games. *Proc. Nat. Acad. Sci.*
- Modify the Fermi rule to include an **exploration dynamic**
  - Let S = {all available strategies}
  - For each individual, a small probability μ of choosing a strategy s at random from S
    - regardless of how successful s was in the current generation
    - regardless of whether anyone is currently using s

What Can This Accomplish?

- Human interactions are very complicated
  - Evolutionary game-theoretic models omit most of the details
    - What the actions do physically when they are performed
    - All of the factors that might lead one to choose one action rather than another
- Can be difficult to develop a model that accurately reflects the essential nature of the interactions
  - Research papers often devote a lot of space to justifying why a proposed model should be considered a good one
- Can't give exact numeric predictions of what would happen in real life
- But:
  - Can provide explanations of the underlying dynamics
  - Can establish support for causal relationships

Modeling Cooperation among Groups

- **Public Goods Game (PGG)**
  - Henrich 2001, Hauert 2002, Brandt 2003, Henrich 2004, Brandt 2006, Traulsen 2009, Hauert 2010

- *N*-player generalization of the famous Prisoner's Dilemma
  - Each individual is asked to contribute an amount *c*
    - **Cooperator**: contributes
    - **Defector**: doesn't contribute
- The sum of all contributions is multiplied by a factor $b > 1$
  - Represents the benefit that being in a society provides to individuals
- Resulting amount is distributed equally among everyone
  - With full cooperation, all get more than they contributed
  - But defectors get the same amount, without contributing anything
- Evolutionary version: use the PGG as the stage game
  - Evolve to nearly 100% defectors
  - Utility for all individuals is near 0

why not exactly 100%?

The Role of Punishment

- In human societies, punishment of defectors is important in the emergence and maintenance of cooperation
  - Humans are willing to pay a cost to punish deviations from cooperative norms
    - Fehr 2000, Fehr 2002, Fehr 2003, Ostrom 1994, Price 2002, Hammerstein 2003, De Quervain 2004, Nakamaru 2006, Camerer 2006
  - Punishment can establish and maintain cooperative norms in collective action and cooperation games
    - Boehm 1993, Boyd 1992, Henrich 2001, Hauert 2002, Henrich 2006, Boyd 2003, Brandt 2003, Brandt 2006, Hauert 2007
  - Different cultures have different propensities to punish deviations from their societal norms
    - Gelfand *et al*. Differences between tight and loose cultures: A 33-nation study. *Science*, 2011.

Modeling Punishment in the PGG

- At each generation:
  - The contribution phase, then a **punishment phase**:
    - Each individual may pay an amount λ to reduce a defector's payoff by an amount $\rho > \lambda$
  - Each strategy involves two choices:
    - whether to defect, whether to punish defectors
- Problem: it doesn't work
  - Punishing lowers punisher's payoff by λ
    - Evolve to nearly 100% non-punishers
  - Without punishment, defectors have higher payoff then cooperators
    - Evolve to nearly 100% defectors

A More Sophisticated Model

- Hilbe & Traulsen, 2012. Emergence of responsible sanctions without second order free riders, antisocial punishment or spite. *Science*.

- **Punishment reputation**: how likely that others will punish
  - Enables one to infer Cooperating has higher payoff than Defecting
- **Information level**: i = Pr[know what the punishment reputation is]
- Cooperation strategies:

  C: cooperate; $\mathbf{O_C}$: opportunistically cooperate;
  D: defect; $\mathbf{O_D}$: opportunistically defect

  - Opportunistic choices depend on punishment reputation
- Punishment strategies:

  R (Responsible): punish defectors; **S** (Spiteful): punish everyone;
  A (Antisocial): punish cooperators; **N** (Non-punisher): punish no one

- **Result**: evolution toward stable proportions of cooperators/defectors and punishers/nonpunishers

What about Third-Party Punishment?

- What I showed you was a model of **direct** punishment
  - In the PGG, defection reduces everyone's payoff
- Motivation for punishing a defector
  - If punishment makes them stop defecting, it directly benefits you
- In 3rd-party punishment, C punishes A for harming B
  - C incurs a cost, but gets no direct benefit if A stops harming B
- So under what conditions would C do this?

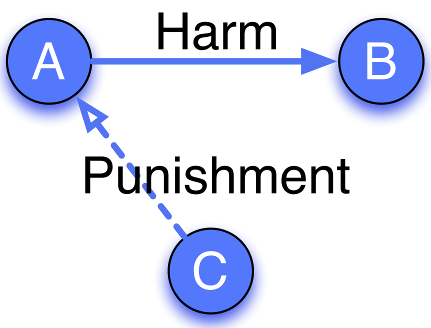

- Create a model similar to Hilbe & Traulsen's, with two main modifications:
  - Third-party punishment
  - Environmental/structural factors related to cultural characteristics

Modeling Third-Party Punishment

- **Interaction phase**: randomly choose pairs of individuals to play a cooperation dilemma
  - $c > 0$: cost of cooperating with other
  - $b > c$: benefit to self if other cooperates

	Cooperate	Defect
Cooperate	$b - c,\ b - c$	$-c,\ b$
Defect	$b,\ -c$	$0,\ 0$

- **Punishment phase**: for each individual that interacted, randomly choose an uninvolved neighbor who may choose to punish the individual
  - Pay λ to reduce their payoff by ρ

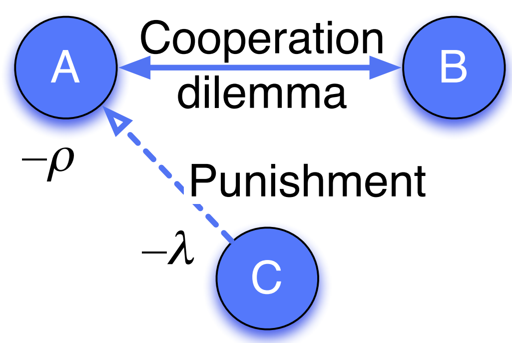

- Punishment reputation, information level, cooperation strategies, punishment strategies
  - similar to Hilbe & Traulsen's

Cultural Characteristics

- **Collectivism vs. Individualism**:
  - One of several cultural scales developed by cultural psychologists
    - Hofstede, et al., 1991. *Cultures and organizations: Software of the mind*. Vol. 2. McGraw-Hill.
  - Individualist cultures (e.g., US, Western Europe)
    - tend to emphasize individual desires and achievements
  - Collectivist cultures (e.g., China, Korea, Japan)
    - tend to emphasize the goals of the family or work group
- Can be applied at different granularities
  - e.g., variations in different parts of the US, or different settings
- Theories of cultural psychology predict 3PP to be more common in collectivist cultures than individualist cultures
  - Can we demonstrate this game-theoretically?

Cultural Characteristics

- How to model collectivism/individualism?
  - No good way to implement them directly in our model
  - Emergent properties of the population, not environmental/structural factors that we can control
- But they correlate with structural factors that we **can** implement
  - **Strength of social ties**
    - generally higher in collectivist cultures than in individualist cultures
  - **Mobility** (ability to leave a social group)
    - generally higher in individualist cultures than in collectivist cultures
- Do this by adding a network structure

Games with Population Structure

- Populations are structured on a network
    - Nowak 1992; Hauert 2002; Nakamaru 2005, 2006; Ohtsuki 2006; Santos 2006; Szabó 2007; Lozano 2008; Roca 2009; Helbing 2010; Perc 2012
  - Each individual is at one of the nodes
  - Edges represent social connections
    - Possibilities for interaction and cultural transmission

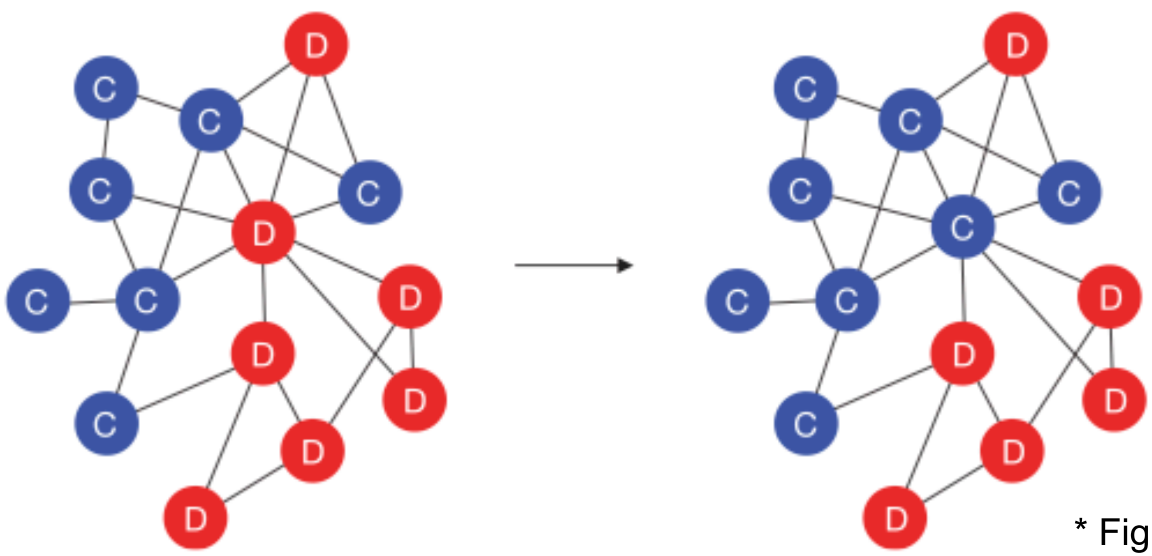

* Figure from Ohtsuki *et al.* 2006

Strength-of-Ties

- **Strength of ties** between humans is measured in terms of how often individuals interact with each other during a period of time.
  - Granovetter 1983 The strength of weak ties: A network theory revisited. *Sociological Theory*, 1(1), 201–233
- If individuals interact with their neighbors and have a limited amount of interactions per time period, then

nodes with many neighbors
→ **low** strength-of-ties

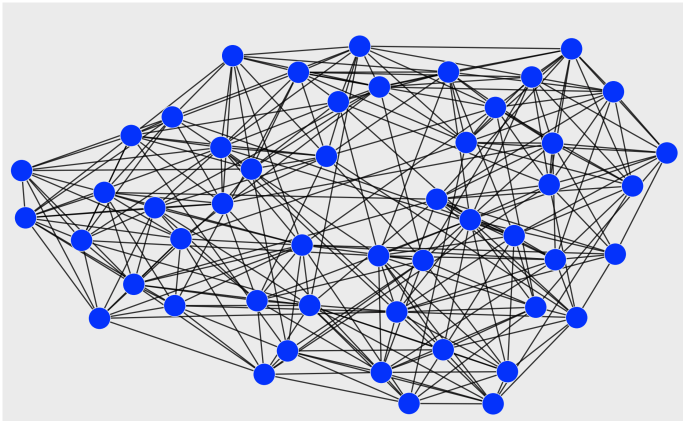

nodes with few neighbors
→ **high** strength-of-ties

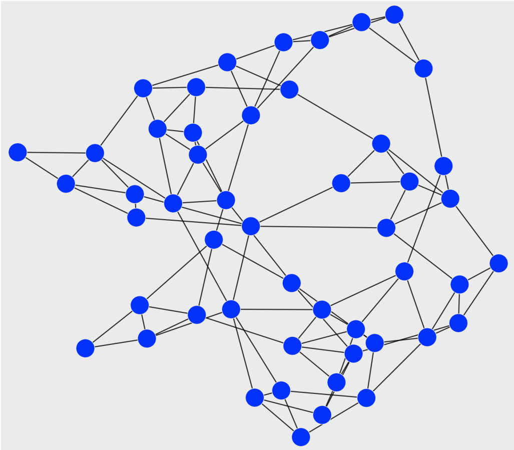

Watts-Strogatz small-world networks

High Strength-of-Ties Enables Evolution of 3PP

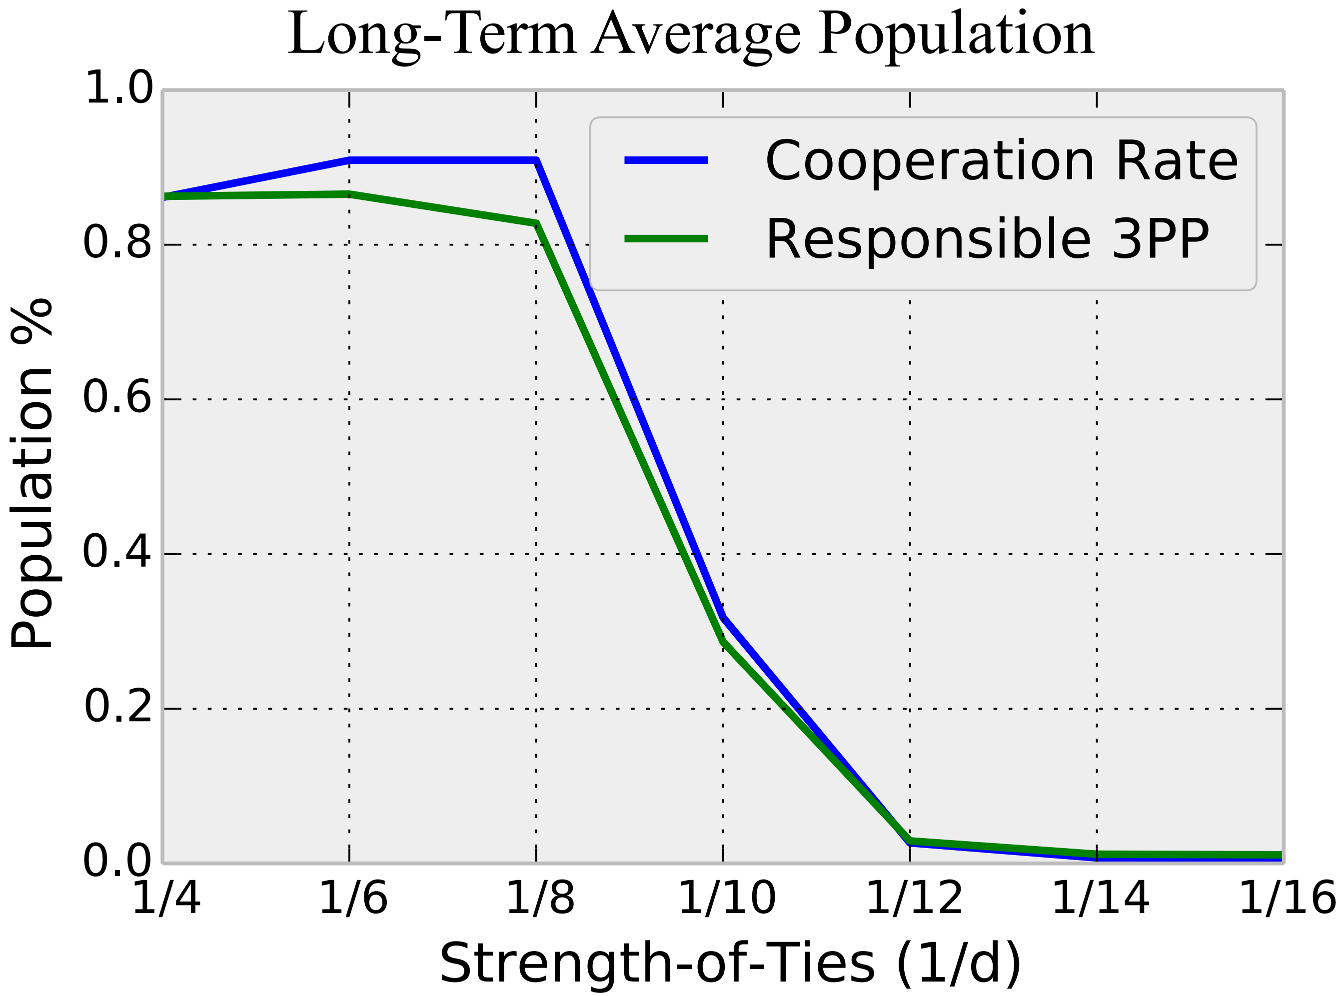

Explanation

- With high strength-of-ties, there can be small groups of agents that interact primarily with each other
  - If their strategies are to cooperate and to punish responsibly, they can
    - achieve high payoffs
    - induce neighbors to adopt their strategy
  - Create a local environment in which agents are encouraged to cooperate and punish responsibly

- With low strength-of-ties, such agents will interact with each other only occasionally
  - can't maintain high payoffs, eventually switch to other strategies

High Strength-of-Ties Example

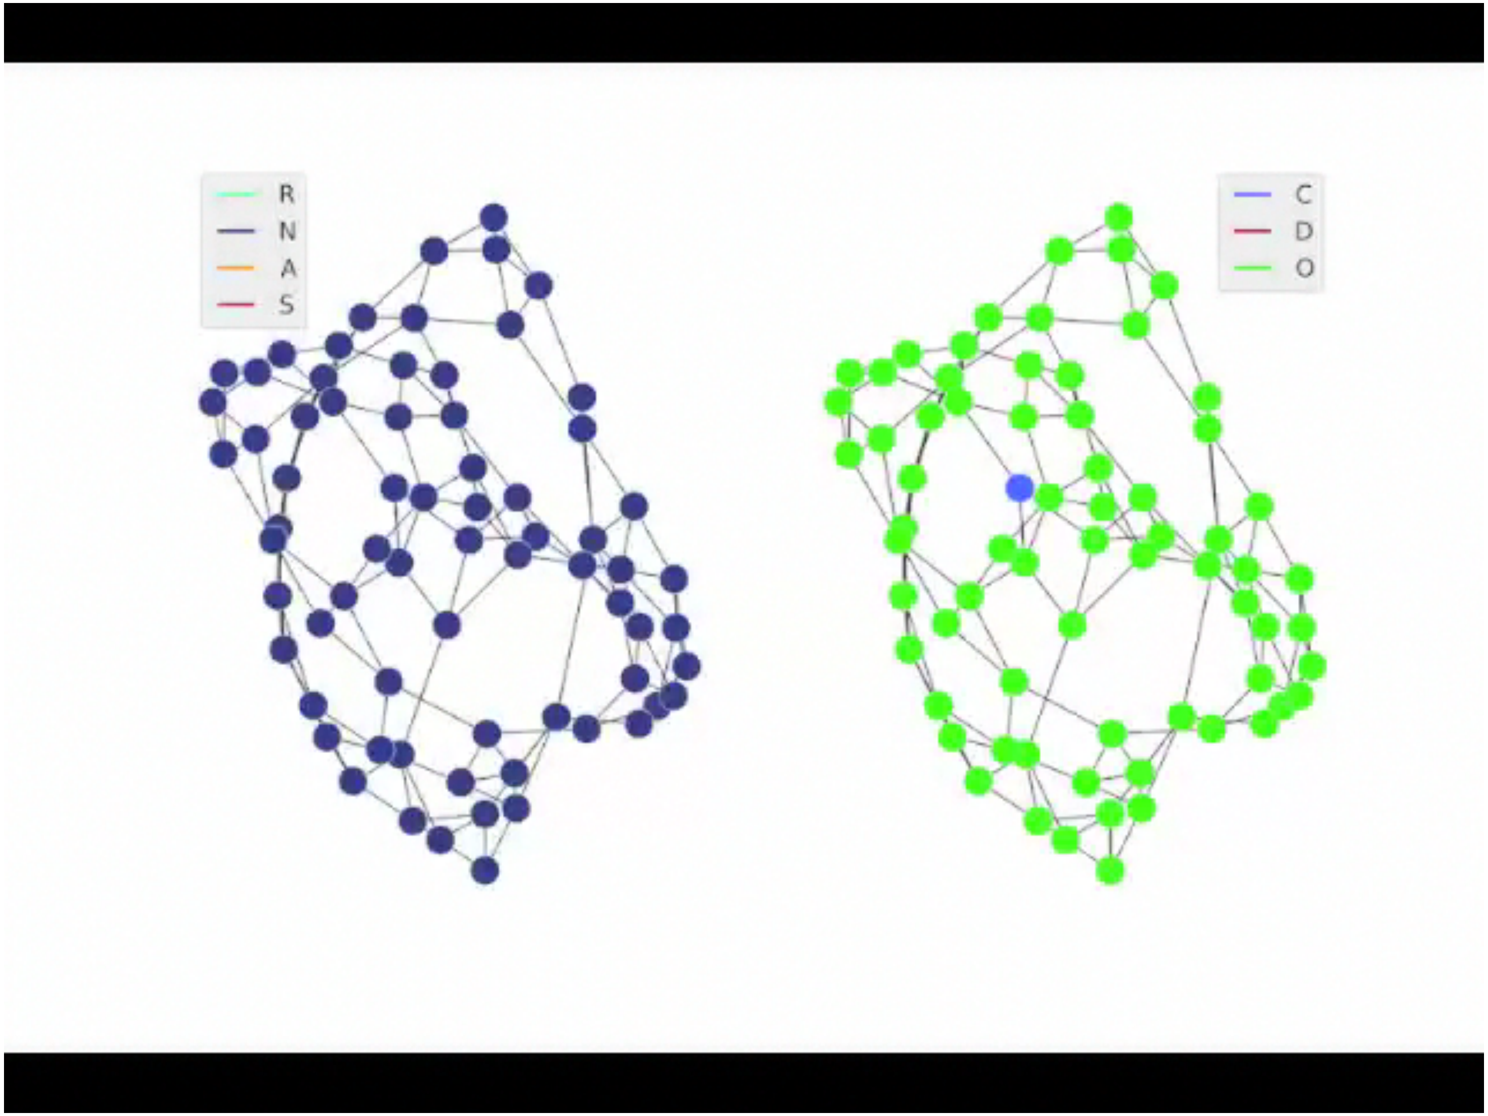

Mobility

- Degree to which humans change their location (social network position within a population)
- Individualistic cultures tend to have very high mobility
  - People can easily exit their social groups
- Collectivistic cultures tend to have low mobility
  - Not as easy to exit social group

- Implementation of Mobility:
  - At each iteration, individuals may switch positions with other randomly chosen individuals, with fixed probability *m*
  - High value for *m* ⇔ high mobility

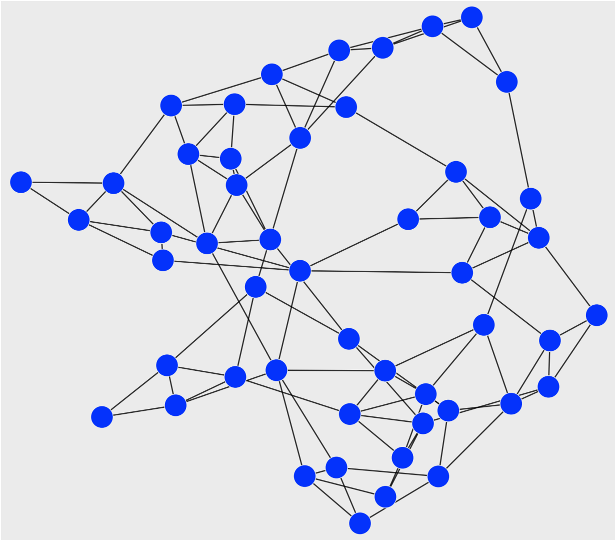

Evolution of 3PP Requires Low Mobility

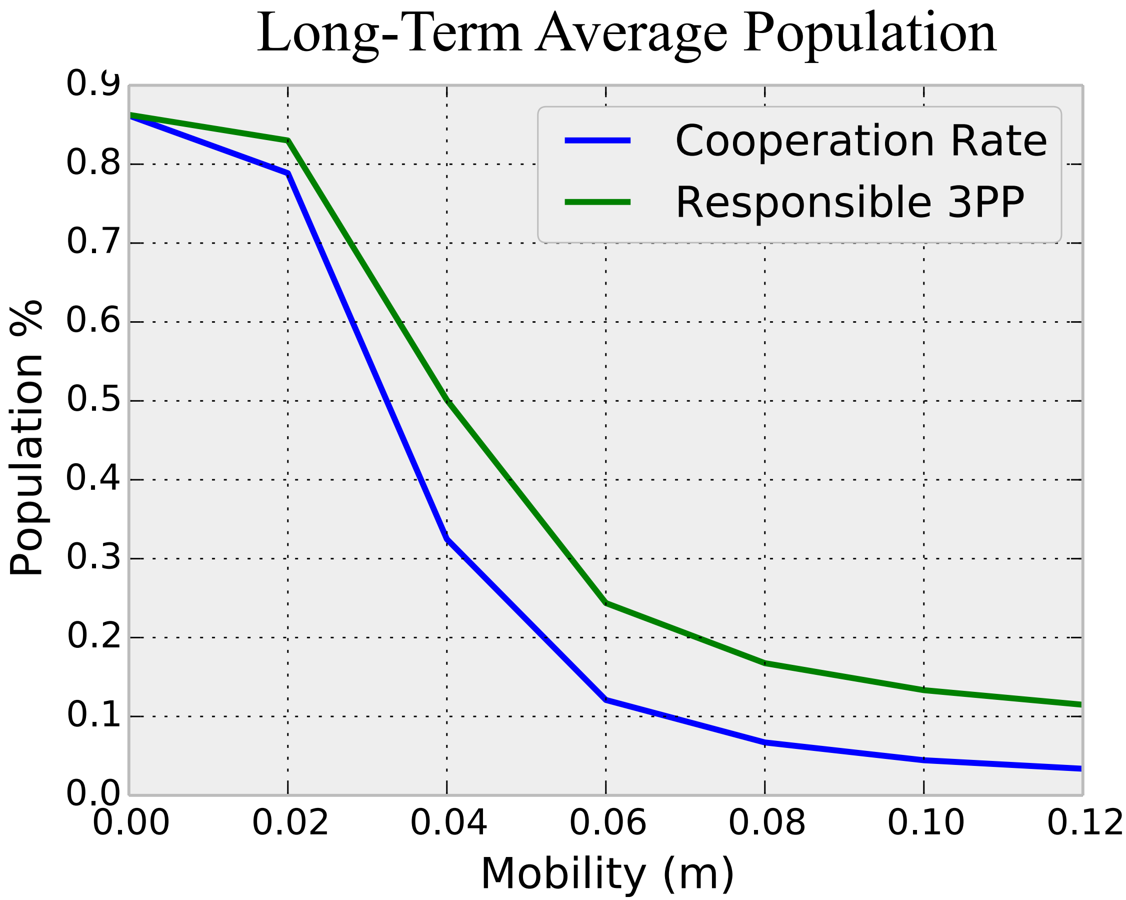

- Successful clusters of 3PP individuals won't last, because the individuals move away

  ➔ 3PP less likely to spread

Discussion

- Example of how an interdisciplinary team can accomplish something that none of us could have done individually

- Results provide support for causal relationship:
  - Evolution of 3PP requires high societal constraint
    - High strength-of-ties, and low mobility
  - Combination of reputation and social structure can lead 3PP to emerge as a trait ultimately beneficial to the individuals carrying it
  - 3PP can't be sustained or uphold cooperation in environments where those factors aren't there
- More generally
  - Individual-level interactions + different structural factors ➔ differences in evolved culture
  - We hope this will help promote cross-cultural understanding
  - Foundation for more complex and ultimately predictive tools

Future Work

- **Culture and Conflict Contagion**:
  - How can we predict when a conflict between two individuals will spread to involve a multitude of others?
  - How do the values and norms in different cultures affect the contagion of conflict?
  - How does conflict contagion relate to factors that influence cultural evolution?

Conflict Contagion - Example

- Hatfield-McCoy feud (1863–1891)
  - Along the border between Kentucky and West Virginia

Hatfield clan, 1897

Dynamics of Conflict Contagion

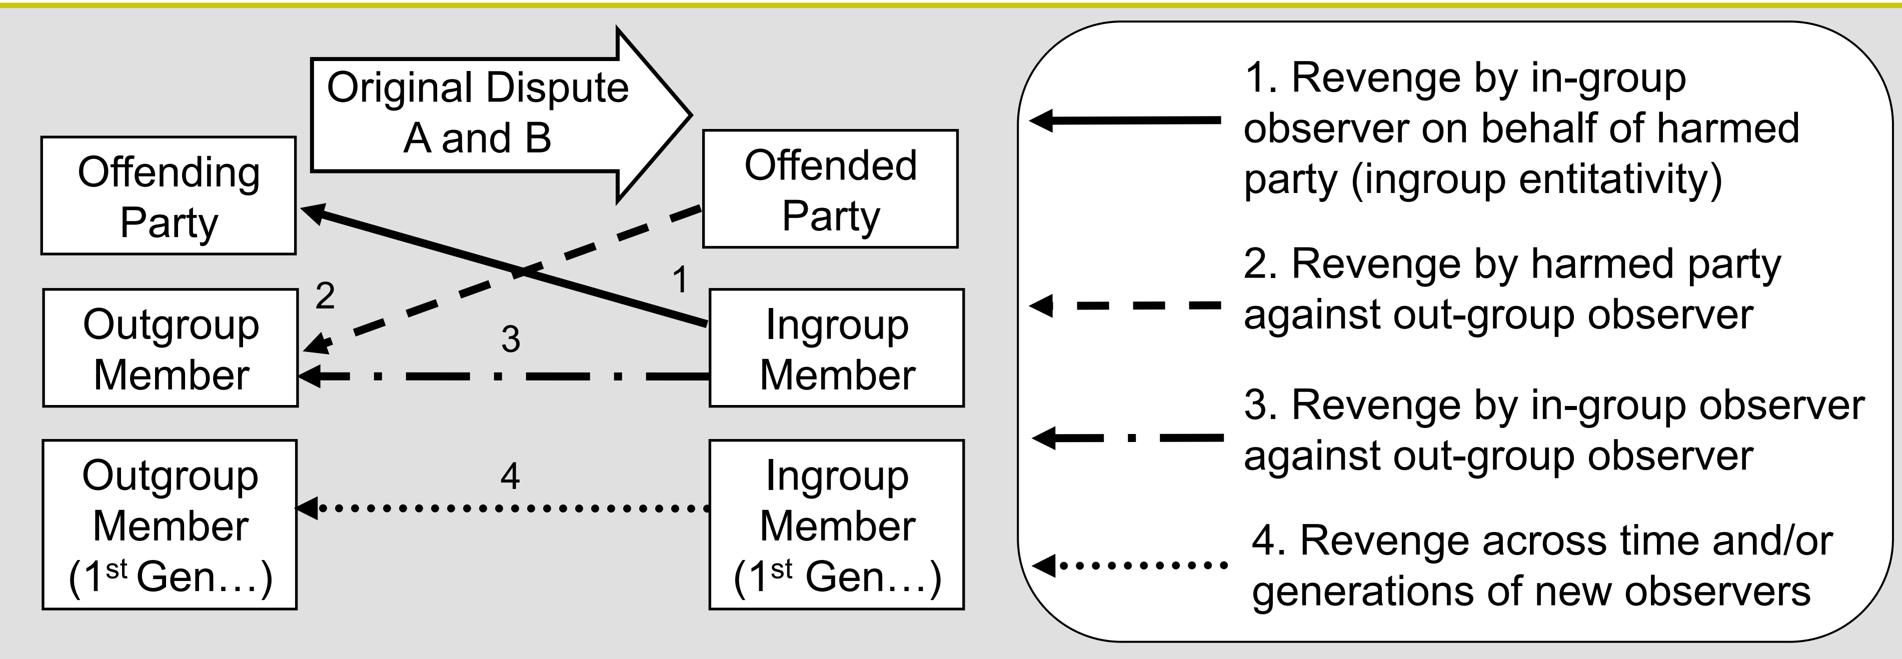

Closely related to **vertical collectivism**

- Lee, Gelfand, & Shteynberg, 2013

- High *ingroup entitativity* – group members interchangeable; depersonalized undifferentiated entities
- High *outgroup entitativity* – outgroups are interchangable
- High *transgenerational entitativity* – ingroup transcends past/future generations

Relation to Third-Party Punishment

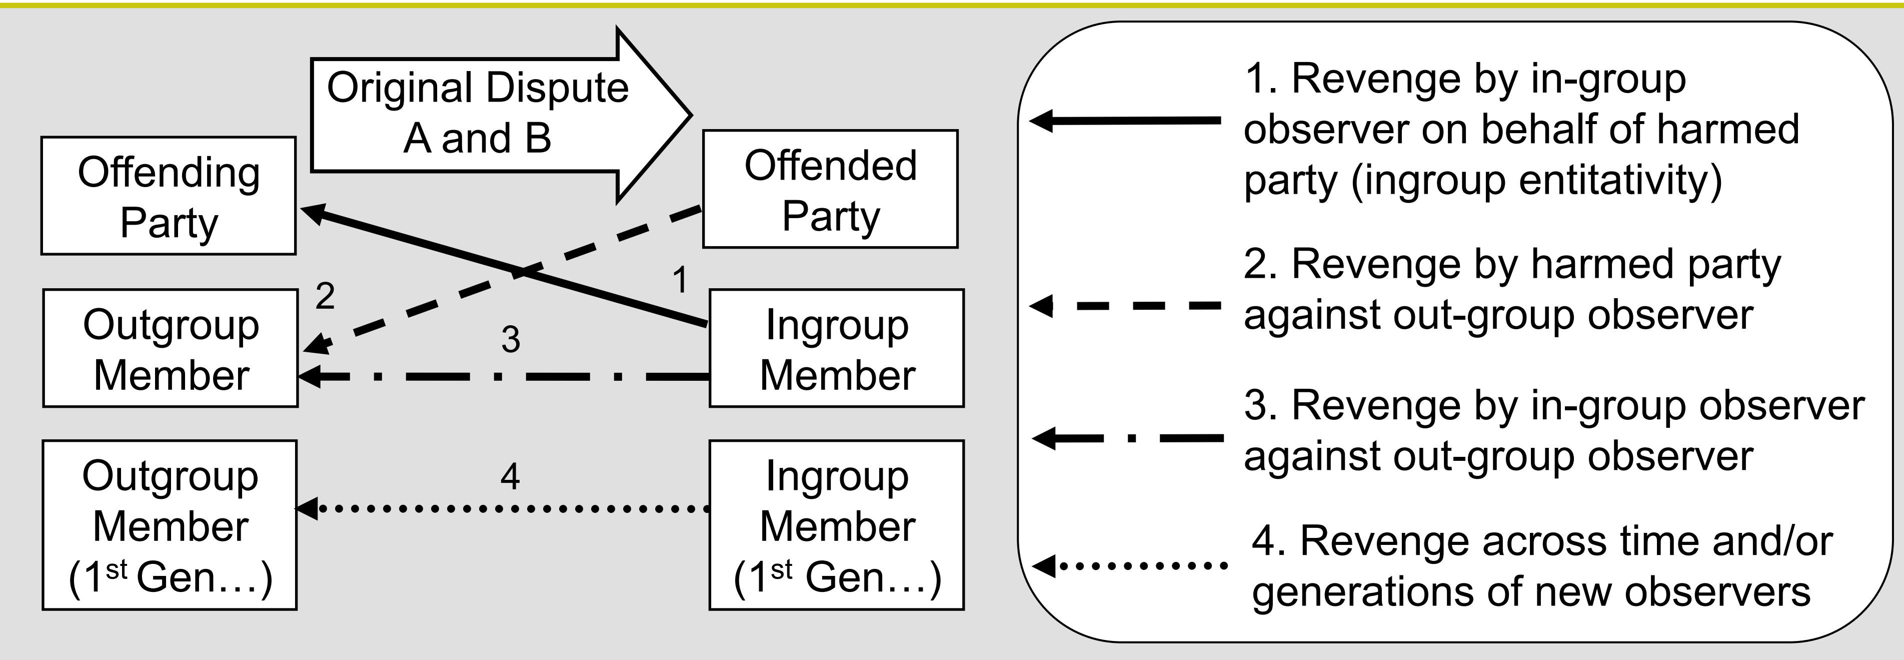

- Combination of 3PP and entitativity
  - Punish others on behalf of someone else in your group
- Central mechanism by which conflict can spread across individuals

Publication of our Current Work

PROCEEDINGS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

High strength-of-ties and low mobility enable the evolution of third-party punishment

Patrick Roos[1,3], **Michele Gelfand**[2], **Dana Nau**[1,3,4] **and Ryan Carr**[1,3]

- Author Affiliations

[1]Department of Computer Science, University of Maryland, College Park, MD, USA
[2]Department of Psychology, University of Maryland, College Park, MD, USA
[3]Institute for Advanced Computer Studies, University of Maryland, College Park, MD, USA
[4]Institute for Systems Research, University of Maryland, College Park, MD, USA
e-mail: roos@cs.umd.edu

Abstract

As punishment can be essential to cooperation and norm maintenance but costly to the punisher, many evolutionary game-theoretic studies have explored how direct punishment can evolve in populations. Compared to direct punishment, in which an agent acts to punish another for an interaction in which both parties were involved, the evolution of third-party punishment (3PP) is even more puzzling, because the punishing agent itself was not involved in the original interaction. Despite significant empirical studies of 3PP, little is known about the conditions under which it can evolve. We find that punishment reputation is not, by itself, sufficient for the evolution of 3PP. Drawing on research streams in sociology and psychology, we implement a

COLLEGE OF COMPUTER, MATHEMATICAL, & NATURAL SCIENCES

INFORMATION FOR... SUPPORT CMNS

SEARCH

About CMNS | Departments | Research | News & Events

News & Events

SHARE

- News
  - Game theory shows why some communities punish bad neighbors
- Events
- Features
- Newsletters

Home / News & Events / News / Game theory shows why some communities punish bad neighbors

Game theory shows why some communities punish bad neighbors

Friday, December 6, 2013

COLLEGE PARK, MD—You're shopping for holiday gifts when you spot someone pocketing a nice pair of leather gloves. What do you do?

A new study by University of Maryland researchers appearing this week in the *Proceedings of the Royal Society B* predicts

The Atlantic
CITIES
PLACE MATTERS

Why Some Communities Police Themselves, While Others Don't

MIKE RIGGS DEC 16, 2013 5 COMMENTS

Greater Greater Washington

The Washington, DC area is great. But it could be **greater**.

LINKS

Breakfast links: Saving the city

by Chad Maddox • December 17, 2013

Self-policing DC: UMD researchers suggest that people with strong ties to a neighborhood are more likely to intervene on another's behalf as a "responsible third party punisher", which may mean less self-policing in transient cities like DC. (Atlantic Cities)

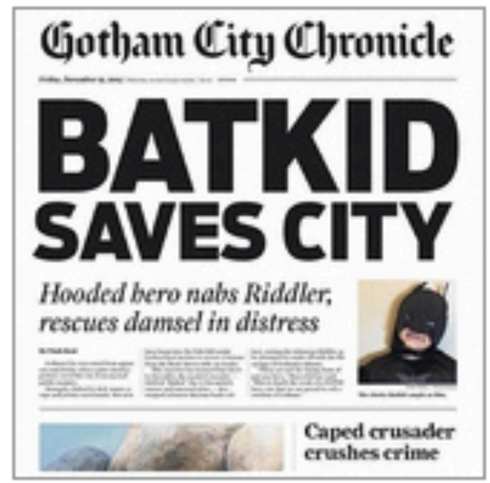

Photo by Shawn on Flicker.

Comments   Add a comment »

"Responsible Third-Party Punisher" sounds like a great idea for a caption for t-shirts.

by Frank IBC on Dec 17, 2013 9:54 am • link • report