# Evolution of State-Dependent Risk Preferences in Social-Modeling Games

Patrick Roos, Ryan Carr, and Dana Nau

Department of Computer Science,
Institute for Systems Research, and
Laboratory for Computational Cultural Dynamics

# Preface

- My research field is Artificial Intelligence
- Interdisciplinary research has interested me for a long time
  - I’ve worked with researchers in at least 8 different academic disciplines
  - Business, Computer Science, Electrical Engr., Industrial Engr., Mathematics, Mechanical Engr., Medicine, Political Science
- People in different fields can have *very* different notions of
  - what questions are important
  - what simplifying assumptions are appropriate
  - what answers are reasonable
  - how to describe what they’ve done
- This can make it hard to communicate intelligibly
  - If what I say doesn’t make sense to you, please stop me and I’ll try to clarify it

# Introduction

- Joint work with two talented PhD students:
  - Patrick Roos
  - Ryan Carr
- Analyses and simulations using several evolutionary-game models
- Objective
  - Explore some hypotheses about biological and cultural evolution of human risk preferences
  - Explore effects of risk-taking on social cooperation

LCCD

# Motivating Example

- Suppose you had to choose between two lotteries
  - Lottery A:
    - you're guaranteed to get \$4,900
  - Lottery B:
    - 50% chance that you'll get \$10,000
    - 50% chance that you'll get nothing
- Which lottery would you choose?

# Decision Making Under Risk

- A well-known decision-theoretic criterion
  - Maximize the expected value of the outcome
- From this point of view, Lottery B looks better
  - Its expected value is ½ ($10,000) + ½ ($0) = $5000
  - Lottery A's expected value is only $4900
- But Lottery B also has a higher risk, and people often are ***risk-averse***
  - Choose an option whose expected value is lower,
    if it avoids the possibility of an undesirable outcome

# Decision Making Under Risk

- There also are situations where people ***seek*** risk
  - Choose a risky option if it offers the possibility of escaping from a bad situation
- Example from American football
  - ***Hail Mary pass***: a very long forward pass with only a small chance of success, made in desperation when the clock is running out

# Human Risk Behavior

- Subject of much empirical and theoretical study
- Evidence that human risk preferences are ***state-dependent***
  - Like your current situation => risk-averse
  - Dislike your current situation enough => risk-seeking
- Several models of this
  - e.g., Prospect Theory, Security-Potential/Aspiration (SP/A) theory

# Objectives and Approach

**Questions we wanted to explore**

- How might state-dependent risk behavior have come about?
  - Several recent papers speculate about relation to evolutionary factors

    Houston, McNamara, & Steer. Do we expect natural selection to produce rational behaviour? *Philosophical Transactions of the Royal Society B* 362 (2007) 1531–1543

    J. R. Stevens. Rational decision making in primates: the bounded and the ecological. *In* Platt and Ghazanfar (eds.), *Primate Neuroethology*. Oxford University Press, 20110 (pp. 98-116)
- How might it relate to cultural evolution?
  - Boyd & Richerson. *Culture and the evolutionary process.* University of Chicago Press, 1988.

**Approach**

- Analyses and simulations using evolutionary-game models intended to reflect both biological and cultural evolution

# Evolutionary Simulations

- ***Evolutionary simulation***: a repeated stochastic game whose structure is intended to model certain aspects of evolutionary environments
  - Consists of a number of ***stages*** or ***generations***
- In each stage, there is a set of *k* agents (*k* is the ***total population size***)
  - The agents interact in some kind of game-theoretic scenario
  - Different agents have different ***strategies*** (ways of choosing actions)
  - Each agent gets a numeric ***payoff*** that's a stochastic function of the ***strategy profile*** (the strategies of all the agents)
- The payoffs are used in deciding what the set of agents will be at the next stage

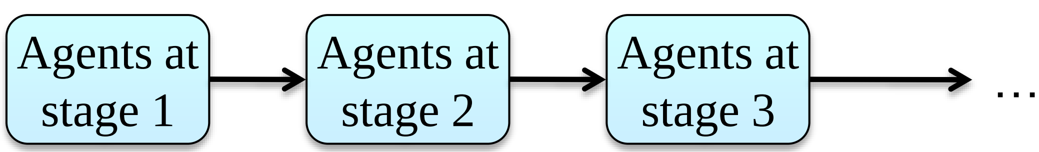

# Evolutionary Simulations

- Consider the set of all agents that use strategy $s$
  - In a biological setting, $s$ may represent a type of animal
  - In a cultural setting, $s$ may represent a learned behavior
- Over time, the number of agents using $s$ may grow or shrink depending on how well $s$ performs
  - How this happens depends on the ***reproduction dynamic*** (next slide)
- At the end of the simulation, $s$'s ***reproductive success***
  = proportion of agents that use $s$ = (number of agents that use $s$) / $k$,
  where $k$ = total population size

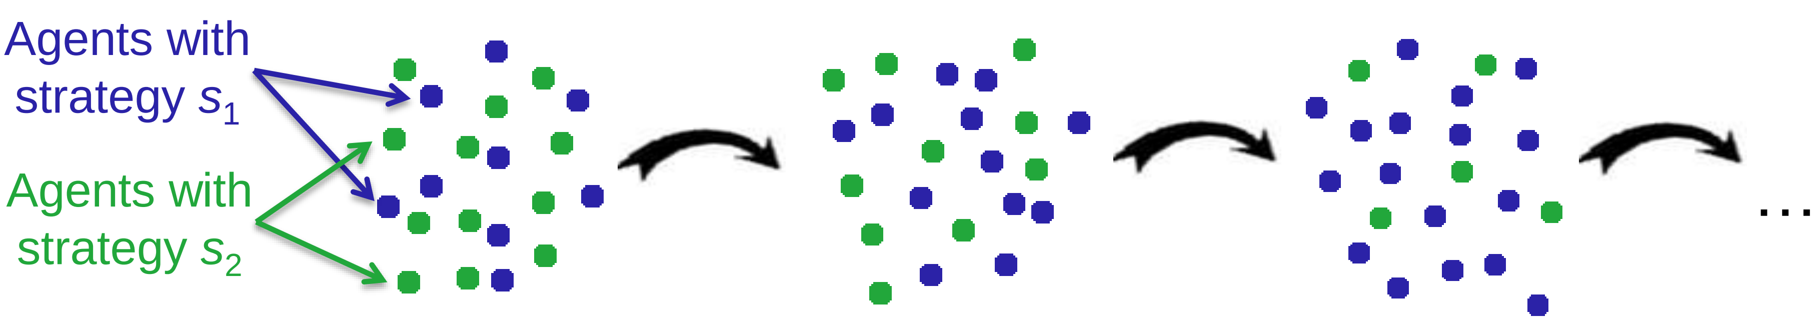


$A_1$ = {agents at stage 1}   $A_2$ = {agents at stage 2}   $A_3$ = {agents at stage 3}

# Reproduction Dynamics

- The ***reproduction dynamic*** is the mechanism for deciding
  - which strategies will disappear
  - which strategies will reproduce
  - how many progeny they'll have
- Many different possible reproduction dynamics
  - I'll briefly discuss two of them
- Later I'll generalize to others

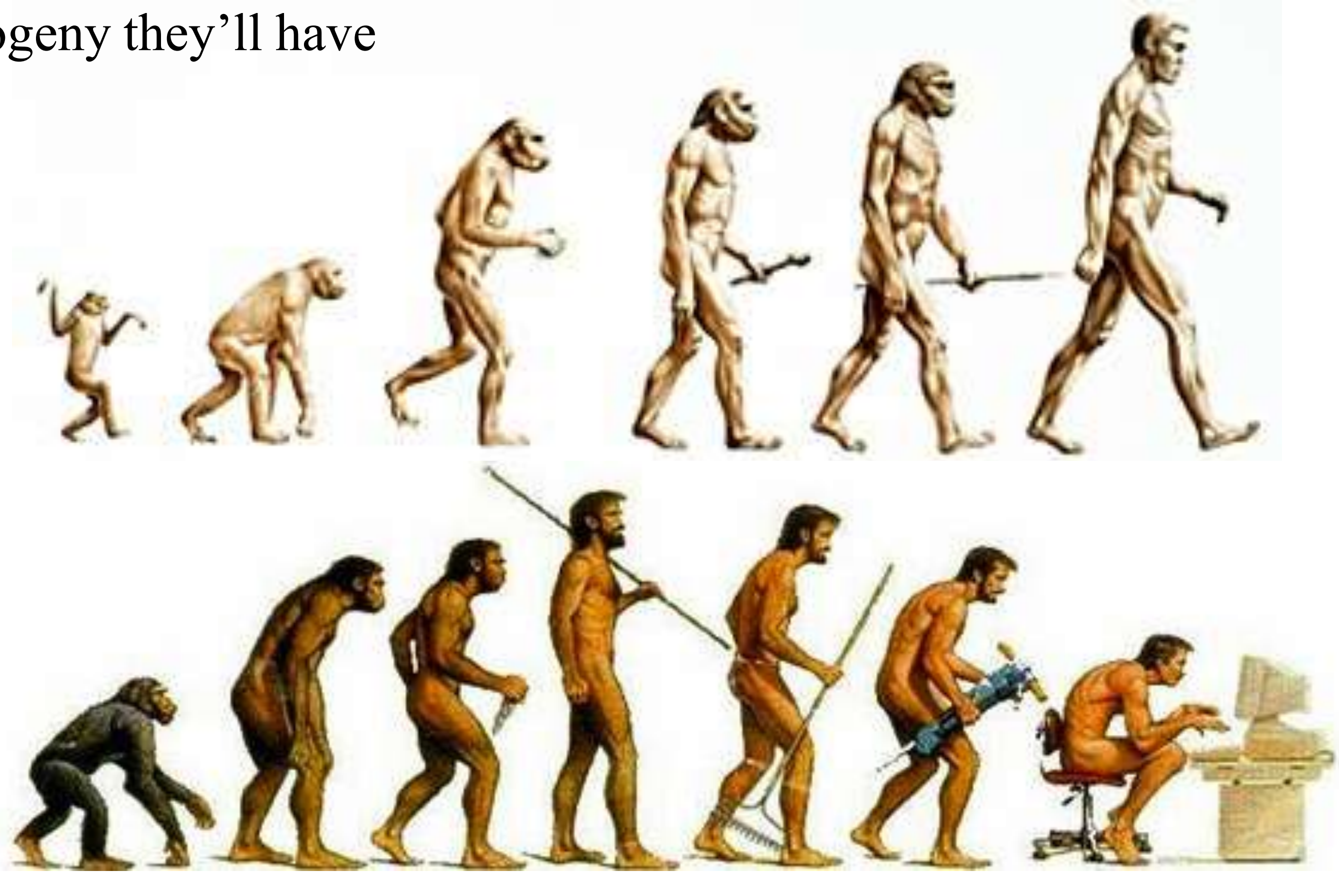

# Reproduction Dynamics

- The ***reproduction dynamic*** is the mechanism for deciding
  - which strategies will disappear
  - which strategies will reproduce
  - how many progeny they'll have
- Many different possible reproduction dynamics
  - I'll briefly discuss two of them
    - **No, not these two!**
- Later I'll generalize to others

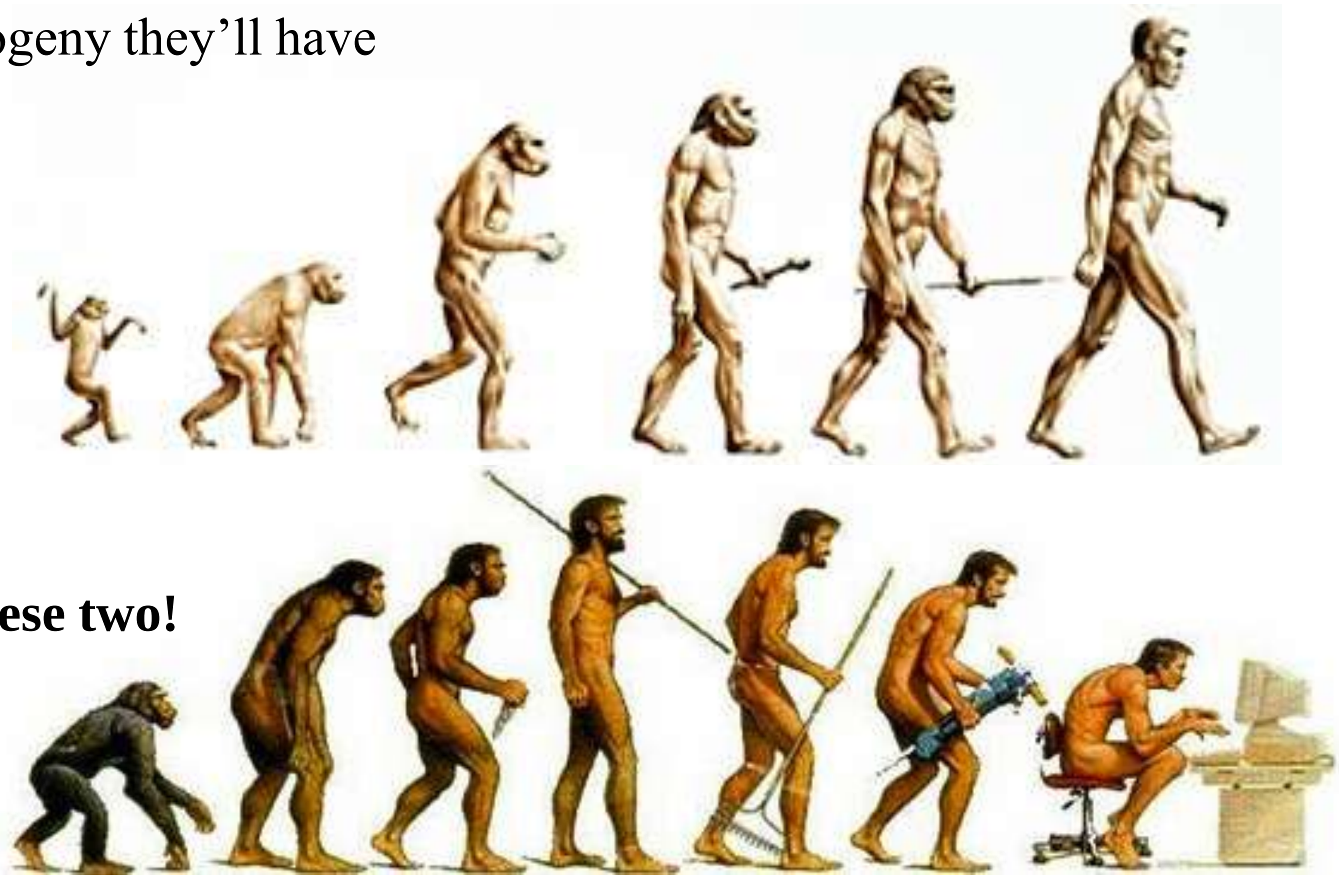

# Reproduction Dynamics

**Replicator dynamic**:

- A strategy's numbers grow or shrink proportionately to how much better or worse it does than the average
  - At stage $i$, let
    - $p$ = proportion of agents that use strategy $s$
    - $r$ = average payoff for those agents
    - $R$ = average payoff for all agents
  - At stage $i$+1, the proportion of agents that use $s$ will be $p\,(r\,/\,R)$
- Does well at reflecting growth of animal populations (where strategy ⇔ type of animal)

- Less clear whether or not it's the best model of economic or cultural behavior
  - Thomas Riechmann. Genetic algorithm learning and evolutionary games. *Journal of Economic Dynamics & Control* **25** (2001), 1019–1037

# Reproduction Dynamics

**Imitate-the-better dynamic**:

- At stage *i*, let $A_i$ = {all agents at stage *i*}
- To build $A_{i+1}$ , do the following steps *k* times:
  - Randomly choose 2 agents in $A_i$
  - Let *a* be the one that got the higher payoff
    (or choose *a* at random if both got the same payoff)
  - Add to $A_{i+1}$ an agent that uses *a*'s strategy
- A strategy's numbers grow if it does better than average
  - But the growth rate is different than with the replicator dynamic

- Evidence that this does well at modeling how behaviors spread when people copy the behavior of others
  - Offerman & Schotter. Imitation and luck: An experimental study on social sampling. *Games and Economic Behavior* **65**:2 (2009), 461–502

LCCD

# A Simple Lottery Game

- A repeated lottery game
- At each stage, agents make choices between two lotteries
  - The ***safe*** lottery: guaranteed reward of 4
  - The ***risky*** lottery: $P(0) = ½$ ; $P(8) = ½$
- Two pure (deterministic) strategies:
  - *S*: always choose the safe lottery
  - *R*: always choose the risky lottery

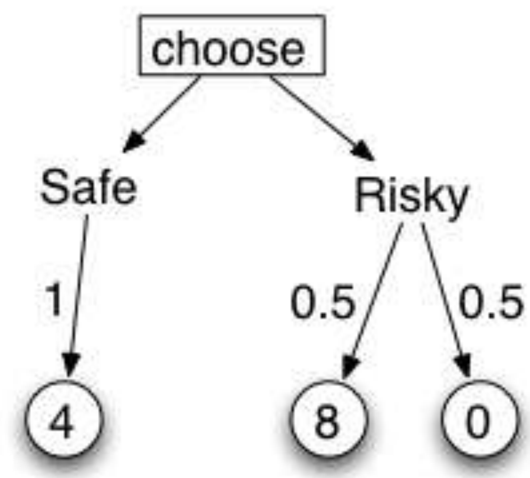

# Lottery Game, Replicator Dynamic

- At each stage, each strategy's average payoff is 4
  - Thus on average, each strategy's population size should stay roughly constant
- Verified by simulation for *S* and *R*
- We would have gotten similar results for any strategy that's a mixture of *S* and *R*

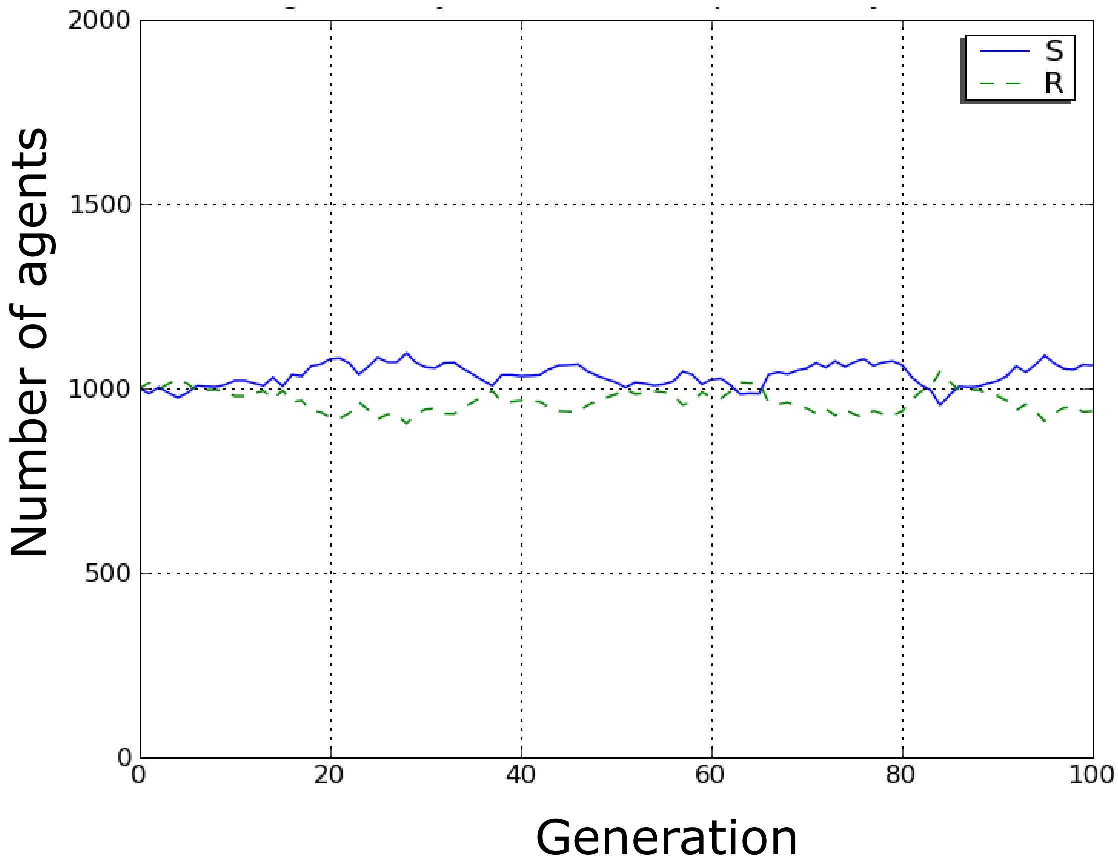

# Lottery Game, Imitate-the-Better Dynamic

- Pick any two agents, and let *s* and *t* be their strategies
- Regardless of what *s* and *t* are, each agent has equal probability of getting a higher payoff than the other
  - Again, each strategy's population size should stay roughly constant
- Verified by simulation for *S* and *R*
- Again, we would have gotten similar results for any mixture of *S* and *R*

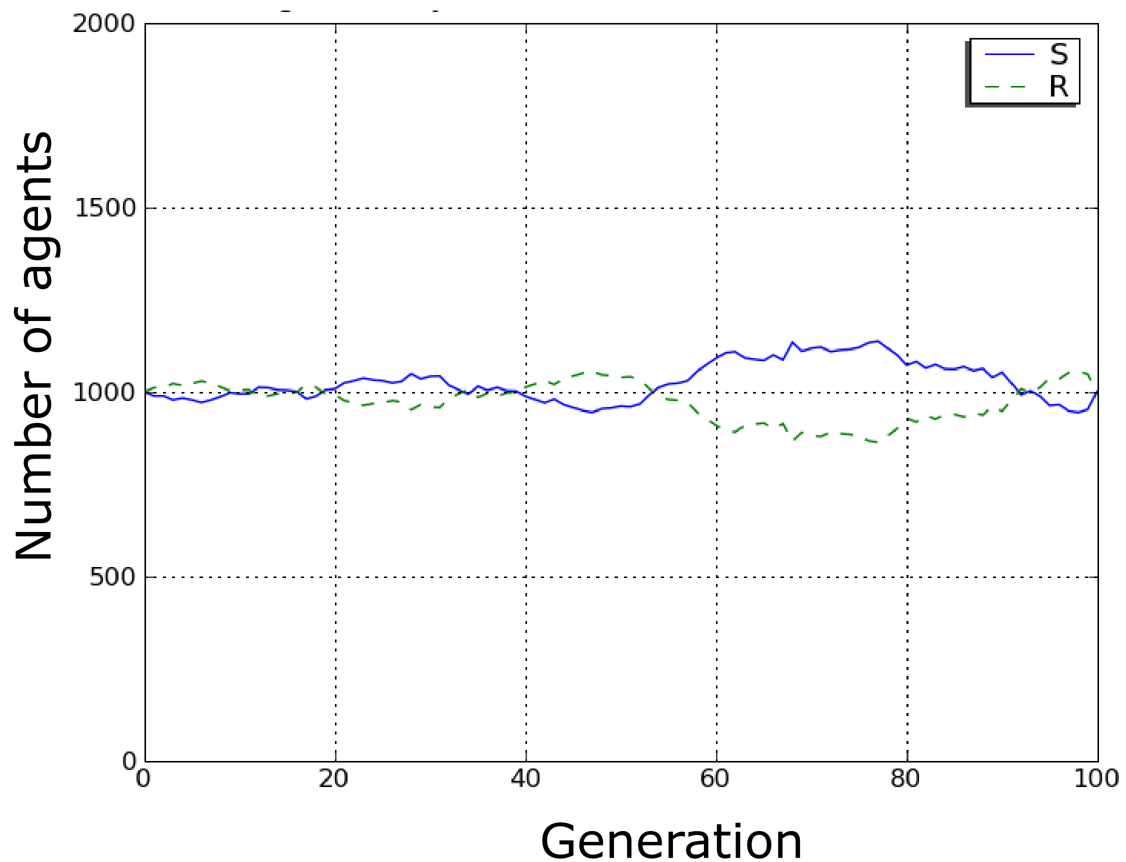

# Double Lottery Game

- At each stage, agents make ***two*** rounds of lottery choices
  1. Choose between the safe lottery and the risky lottery, get a payoff
  2. Choose between the safe lottery and the risky lottery again, and get an additional payoff

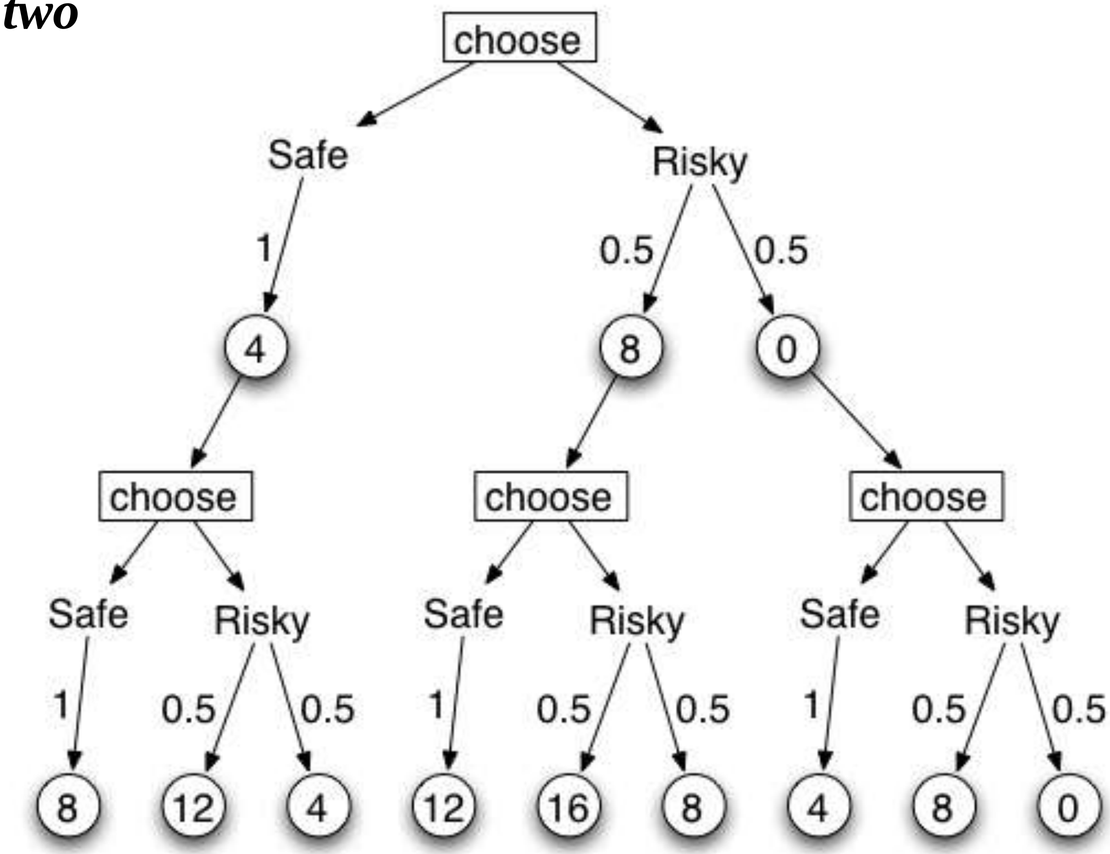

# Double Lottery Game

There are 6 pure strategies:

- *S*: choose Safe both times
- *SR*: 1st time choose Safe
  2nd time choose Risky
- *RS*: 1st time Risky
  2nd time Safe
- *R*: Risky both times
- *RwS*: 1st time Risky
  - 2nd time: if 1st time was a win (payoff 8), then Safe, otherwise Risky
- *RwR*: 1st time Risky
  - 2nd time: if 1st time was a win (payoff 8), then Risky, otherwise Safe

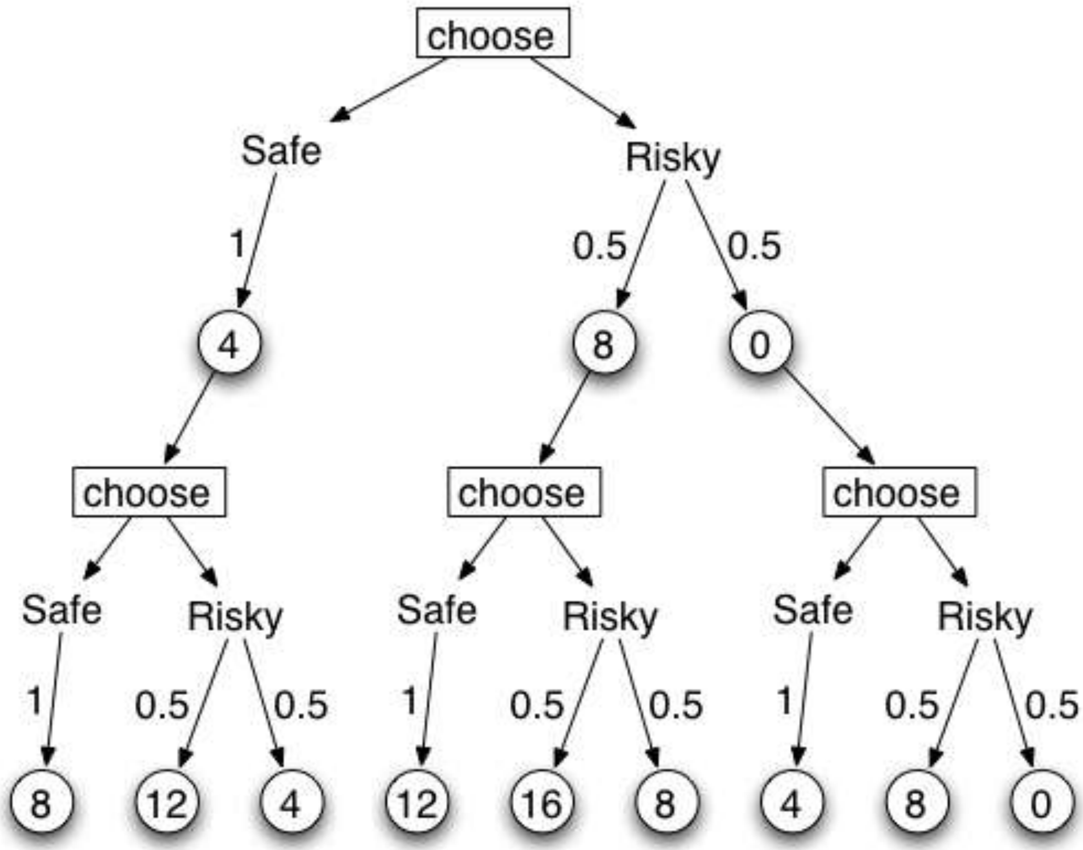

# Distribution of Payoffs for Each Strategy

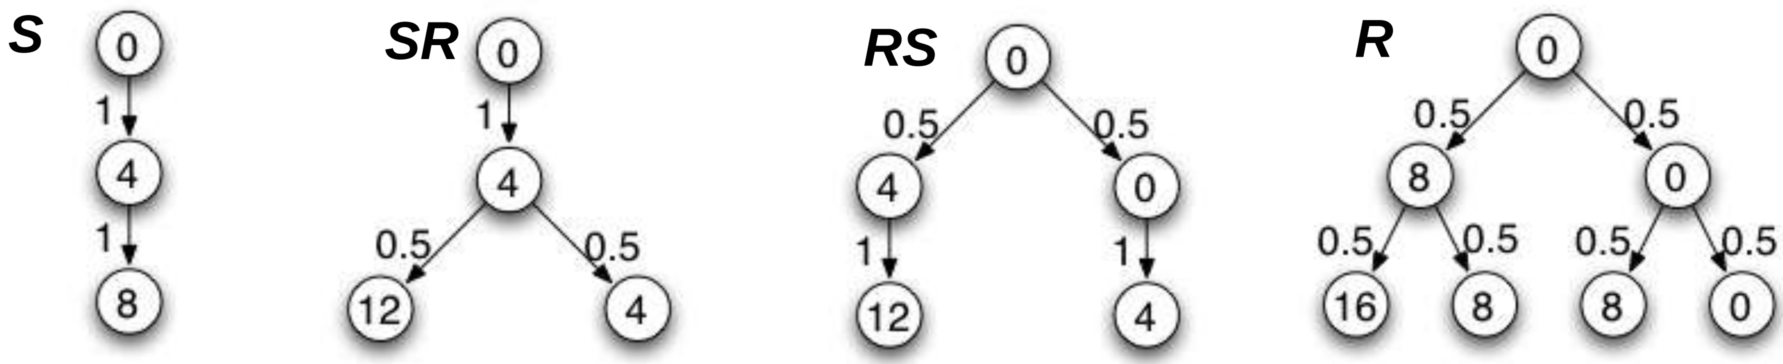


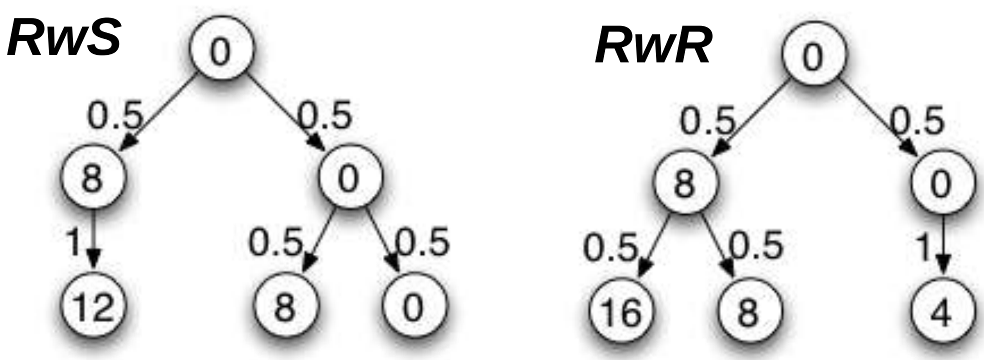


- For every strategy, the expected value is 8
- But the distribution of payoffs differs

| | *S* | *SR* | | *RS* | | *R* | | | *RwS* | | | *RwR* | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Payoff | **8** | **12** | **4** | **12** | **4** | **16** | **8** | **0** | **12** | **8** | **0** | **16** | **8** | **4** |
| Probability | 1 | ½ | ½ | ½ | ½ | ¼ | ½ | ¼ | ½ | ¼ | ¼ | ¼ | ¼ | ½ |

LCCD

# Double Lottery Game, Replicator Dynamic

- At each stage, each strategy's expected payoff is 8
  - Thus on average, each strategy's population size should stay roughly constant
- Verified by simulation for all 6 strategies

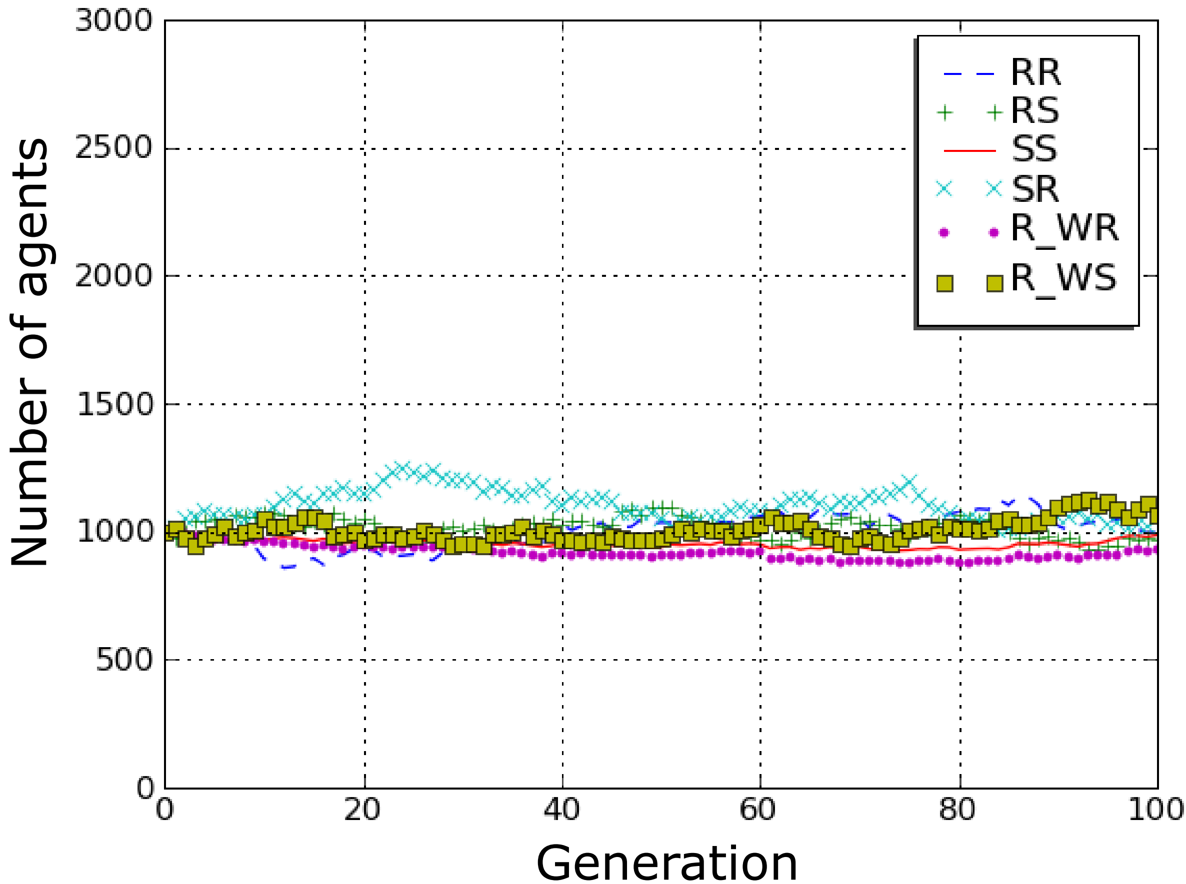


LCCD

# Double Lottery Game, Imitate-the-Better Dynamic

- For imitate-the-better, do the following *k* times:
  - Choose two agents *a* and *b*, and compare their payofs
    - Reproduce the one that got a higher payoff
    - If they got the same payoff, choose either of them at random

- Suppose *a* uses *S* and *b* uses *SR*
  - *P*(*b* gets 4) = ½ => *a* reproduces
  - *P*(*b* gets 12) = ½ => *b* reproduces
- Thus *a* and *b* are equally likely to reproduce

| | *a* | *b* | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | *S* | *SR* | | *RS* | | *R* | | | *RwS* | | | *RwR* | | |
| Payoff | **8** | **12** | **4** | **12** | **4** | **16** | **8** | **0** | **12** | **8** | **0** | **16** | **8** | **4** |
| Probability | 1 | ½ | ½ | ½ | ½ | ¼ | ½ | ¼ | ½ | ¼ | ¼ | ¼ | ¼ | ½ |

# Double Lottery Game, Imitate-the-Better Dynamic

- Suppose *a* uses *S* and *b* uses *RwS*
  - *P*(*b* gets 0) = ¼ => *a* reproduces
  - *P*(*b* gets 8) = ¼ => *a* and *b* equally likely to reproduce
  - *P*(*b* gets 12) = ½ => *b* reproduces
- Thus
  - *P*(*a* reproduces) = ¼ + ½ (¼) = 0.375
  - *P*(*b* reproduces) = ½ + ½ (¼) = 0.625
- *RwS* dominates *S*

| | a | | | | | | | | | b | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | *S* | *SR* | | *RS* | | *R* | | | *RwS* | | | *RwR* | | |
| Payoff | **8** | **12** | **4** | **12** | **4** | **16** | **8** | **0** | **12** | **8** | **0** | **16** | **8** | **4** |
| Probability | 1 | ½ | ½ | ½ | ½ | ¼ | ½ | ¼ | ½ | ¼ | ¼ | ¼ | ¼ | ½ |

# Double Lottery Game, Imitate-the-Better Dynamic

- In general:
  - *RwS* dominates *S*, *R*, and *RwR*
    - In a pair where one of the agents uses one of those strategies and the other uses *RwS*, the *RwS* agent is more likely to reproduce
  - For all other pairs of strategies, neither dominates the other
    - Both are equally likely to reproduce

Dominated by *RwS*

| | *S* | *SR* | | *RS* | | *R* | | | *RwS* | | | *RwR* | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Payoff | **8** | **12** | **4** | **12** | **4** | **16** | **8** | **0** | **12** | **8** | **0** | **16** | **8** | **4** |
| Probability | 1 | ½ | ½ | ½ | ½ | ¼ | ½ | ¼ | ½ | ¼ | ¼ | ¼ | ¼ | ½ |

LCCD

# Double Lottery Game, Imitate-the-Better Dynamic

- Start with equal numbers of all 6 strategies
- *RwS* has an advantage whenever it's paired with *S*, *R*, or *RwR*
    - *RwS* should increase until *S*, *R*, and *RwR* become extinct
- For all other pairs of strategies, neither has an advantage
    - Once *S*, *R*, and *RwR* are extinct, the population should stabilize
- Verified by simulation ➔

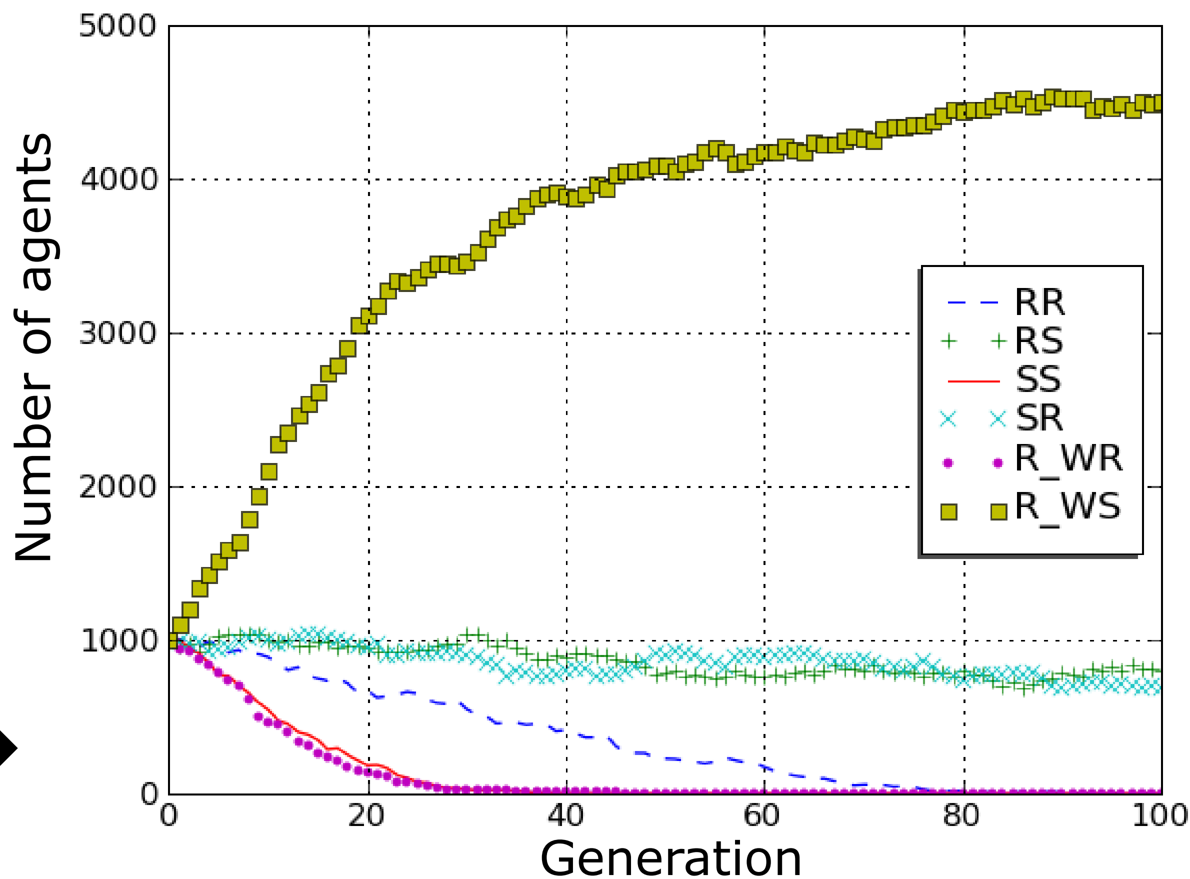

# Discussion

- Lots of different possible reproduction dynamics
- The replicator dynamic and the imitate-the-better dynamic are thought to be good models of biological and cultural evolution, respectively
  - But we're not sure that either of them is a 100% accurate model, so let's look at other reproduction dynamics
- Hypothesis:
  - For *any* reproduction dynamic other than the replicator dynamic, a strategy other than utility maximization is likely to do best
- To test this hypothesis, we need to examine
  - Other reproduction dynamics
  - Games in which the safe and risky lotteries have different expected payoffs
- That's what I'll discuss next …

# 1. Other Reproduction Dynamics

- ***Imitation dynamics*** are a parameterized class of reproduction dynamics with a parameter $0 \leq \alpha \leq 1$

  [Hofbauer & Sigmund. Evolutionary game dynamics. *Bulletin of the American Mathematical Society* **40** (2003), 479–519]

  - Case $\alpha = 0$: imitate-the-better
  - Case $\alpha = 1$: replicator dynamic
  - Case $0 < \alpha < 1$: in between

- **Theorem:** For $0 < \alpha < 1$, *RwS* is evolutionarily stable.
- In a population that includes any mixture of *RwS* and the other strategies, *RwS* will go to 100% and the others will go extinct

LCCD

# 2. Other Expected Payoffs

***S***


***SR***

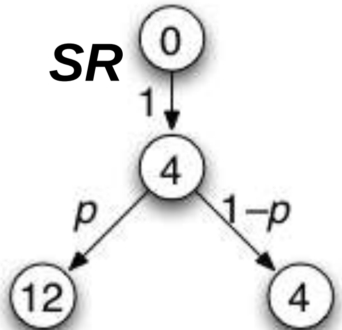


***RS***

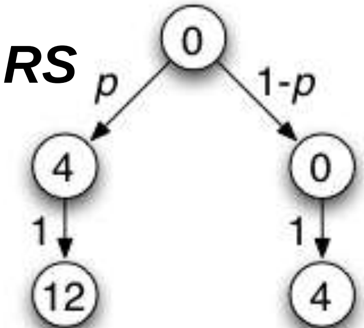


***R***

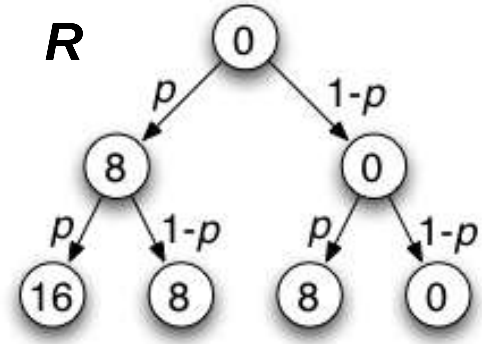


***RwS***

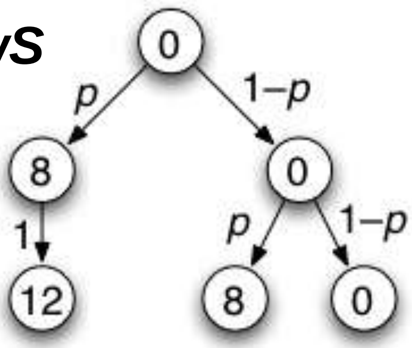


***RwR***

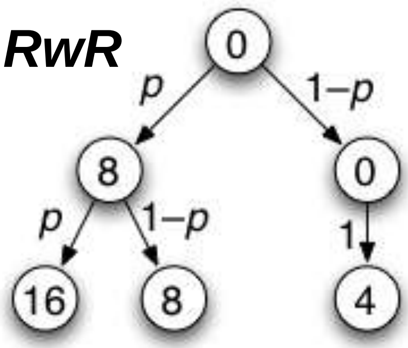


- For the risky lottery, let $P(8) = p$ and $P(0) = 1\text{–}p$
  - Expected value is $8p$
- Safe lottery's payoff is still 4

| | *S* | *SR* | | *RS* | | *R* | | | *RwS* | | | *RwR* | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Payoff | **8** | **12** | **4** | **12** | **4** | **16** | **8** | **0** | **12** | **8** | **0** | **16** | **8** | **4** |
| Prob. | 1 | $p$ | $1\text{–}p$ | $p$ | $1\text{–}p$ | $p^2$ | $2p(1\text{–}p)$ | $(1\text{–}p)^2$ | $p$ | $p(1\text{–}p)$ | $(1\text{–}p)^2$ | $p^2$ | $p(1\text{–}p)$ | $1\text{–}p$ |

# Double Lottery Game

- For all values of *p* and *α*, compare *RwS* to *S* and *R*

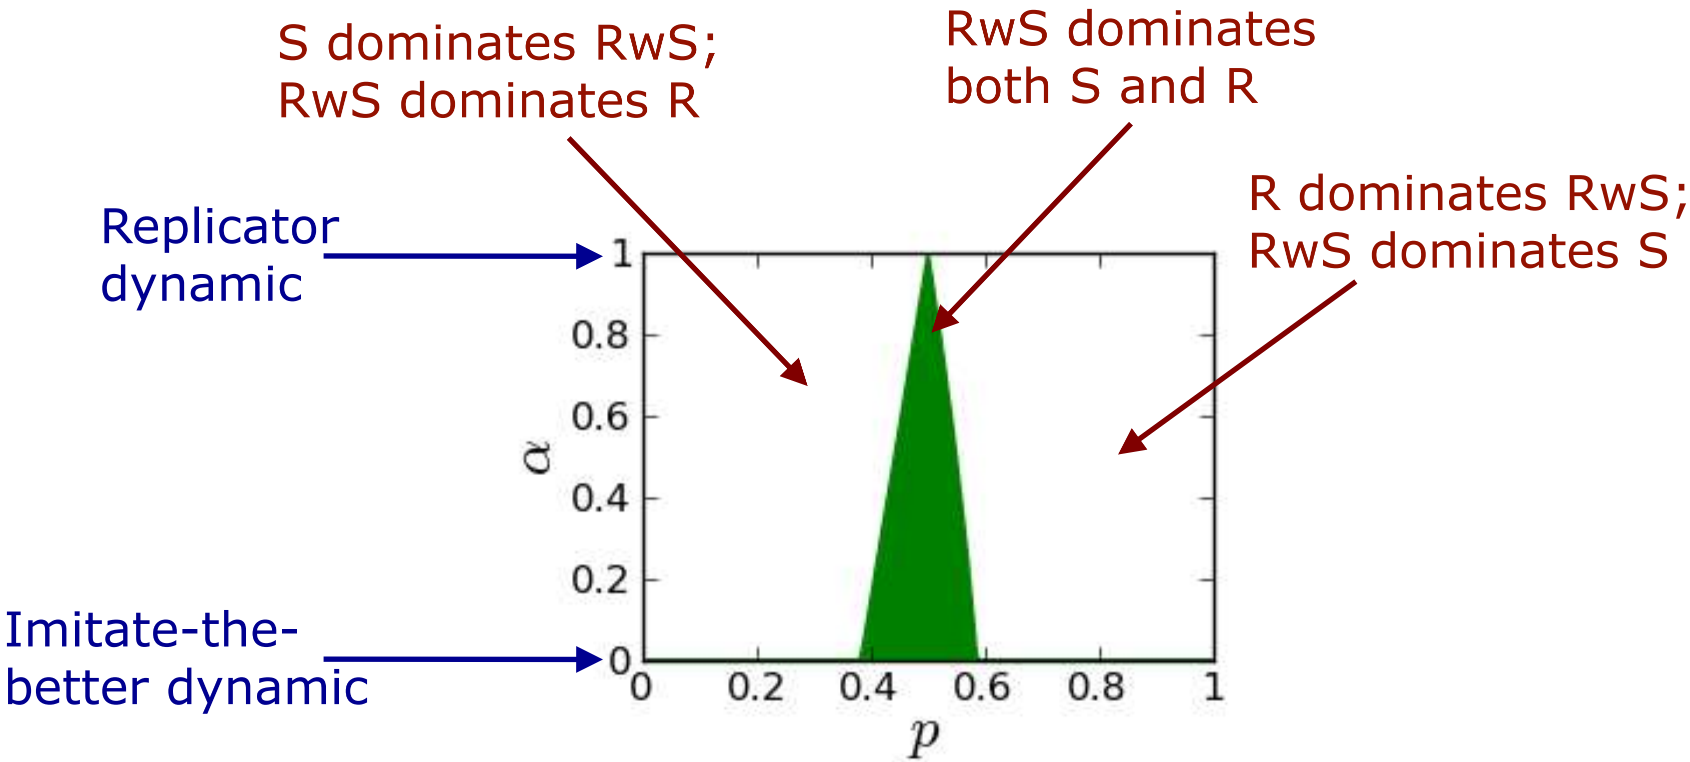

# More Complex Interactions

- In the lottery games, each agent's payoff depended only on its own choices
  - ➢ What about situations in which the agents interact?
  - ➢ Instead of lotteries, use non-zero-sum games
- We used the Stag Hunt

Stag Hunt

Chicken Game

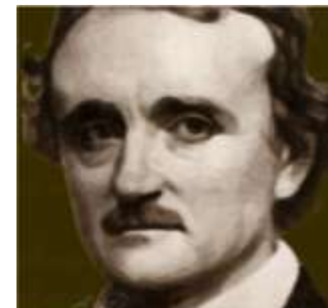

Matching Pennies

Prisoner's Dilemma

Roshambo

Battle of the Sexes

Ultimatum Game

# Stag Hunt

- Simple model of a situation where one must decide whether to work alone or cooperate with others
- Two hunters, each hunting for food
- Hunting for hare: solitary activity
  - Small payoff (4), but *safe*:
    - Same payoff, regardless of what the other hunter does
- Hunting for stag: cooperative activity
  - Possibility of a much bigger payoff (8), but *risky*:
    - Payoff = 8 only if the other hunter cooperates
  - In an evolutionary game setting, *P*(payoff = 8) depends on the relative proportions of stag hunters and hare hunters at stage *i*

Stag Hunt

| *Hunter 1* \ *Hunter 2* | Stag (risky) | Hare (safe) |
|---|---|---|
| Stag (risky) | 8, 8 | 0, 4 |
| Hare (safe) | 4, 0 | 4, 4 |

Nash equilibria

# Evolutionary Double Stag Hunt

- Instead of two lotteries at each stage, have two Stag Hunt games
  - Randomly divide the agents into pairs,
    - Each pair plays Stag Hunt
  - Randomly divide the agents into pairs again
    - Each pair plays another Stag Hunt
- 6 pure strategies (by analogy with the double lottery game)
- But initially we'll just be interested in two of them
  - *Stag*: hunt stag both times (like the *R* strategy in the double lottery game)
  - *Hare*: hunt hare both times (like the *S* strategy)
- Consider the case where every agent uses either *Stag* or *Hare*

Stag Hunt

| *Hunter 1* \ *Hunter 2* | Stag (risky) | Hare (safe) |
|---|---|---|
| Stag (risky) | 8, 8 | 0, 4 |
| Hare (safe) | 4, 0 | 4, 4 |

LCCD

# Evolutionary Double Stag Hunt

- Let $p_i$ = proportion of *Stag* agents at stage $i$
- Payoff for *Hare* is 4 + 4 = 8, regardless of the other players' strategies
- Payoff distribution for *Stag*:
  - $P$(play against *Stag* twice) = $p_i^2$
    => payoff = 8 + 8 = 16
  - $P$(play against *Hare* twice) = $(1–p_i)^2$
    => payoff = 0
  - $P$(play once against each) = $2p_i(1–p_i)$
    => payoff = 0 + 8 = 8

Stag Hunt

| *Hunter 2* / *Hunter 1* | Stag (risky) | Hare (safe) |
|---|---|---|
| Stag (risky) | 8, 8 | 0, 4 |
| Hare (safe) | 4, 0 | 4, 4 |

Double Stag Hunt

| | *Hare* | *Stag* | | |
|---|---|---|---|---|
| Payoff | **8** | **16** | **8** | **0** |
| Prob. | 1 | $p_i^2$ | $2p_i(1–p_i)$ | $(1–p_i)^2$ |

- Same formulas as for the double lottery, but with $p_i$ instead of $p$
  - Amount of risk depends on how many agents of each type at stage $i$
- Examine what happens with replicator and imitate-the-better dynamics

# Replicator Dynamic

- Proportion of *Stag* agents at stage $i$+1 is
  - $p_{i+1} = p_i\, s_i \,/\, R_i$
  - where
    - $s_i$ = *Stag*'s average payoff $= 16\,p_i^2 + 16\,p_i\,(1-p_i) + 0\,(1-p_i)^2 = 16\,p_i$
    - $R_i$ = average payoff for all agents $= (p_i\,s_i + 8(1\text{-}p_i)) = 16\,p_i^2 - 8\,p_i + 8$
  - Thus $p_{i+1} = 16p_i^2 \,/\, (16p_i^2 - 8\,p_i + 8)$
    - If $p_1 = ½$, then $p_i = ½$ for all $i$ (more about this later)
    - If $p_1 < ½$, then $p_i \to 0$
    - If $p_1 > ½$, then $p_i \to 1$

| | *Hare* | *Stag* | | |
|---|---|---|---|---|
| Payoff | **8** | **16** | **8** | **0** |
| Prob. | 1 | $p_i^2$ | $2p_i\,(1-p_i)$ | $(1-p_i)^2$ |

- Larger group gets a bigger average payoff

=> group grows even larger → even bigger avg. payoff

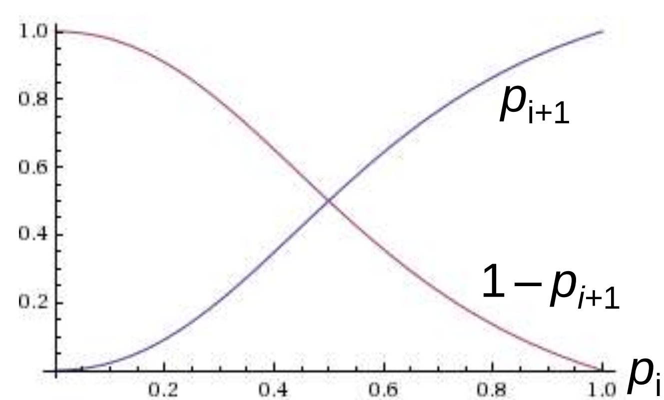

# Replicator Dynamic (continued)

- On the previous slide, I said
    - If $p_1 = ½$, then $p_i = ½$ for all $i$
  - ➢ That neglects the effects of random variation
- Random variation => eventually we'll get a stage $j$ for which $p_j \neq ½$
    - If $p_j < ½$, then $p_i \to 0$
    - If $p_j > ½$, then $p_i \to 1$
  - ➢ $p_i \to 0$ and $p_i \to 1$ are equally likely
- Confirmed by simulation:
  - ➢ 200 simulation runs, each starting with 3000 *Stag* and 3000 *Hare*
    - 101 runs converged to 100% *Stag*
    - 99 runs converged to 100% *Hare*

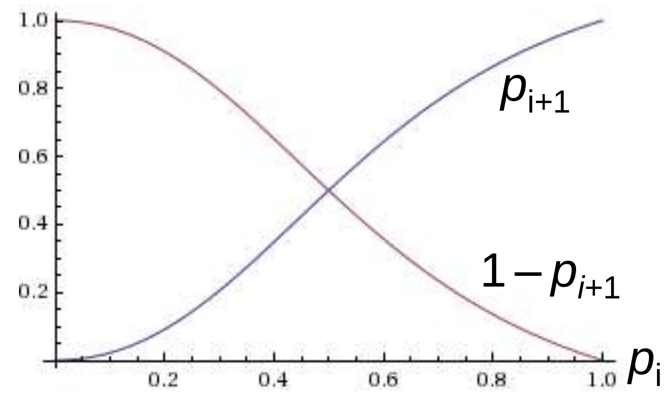

# Imitate-the-Better Dynamic

- Compare pairs of randomly chosen agents
  - Reproduce the one with the higher payoff
  - Same payoff => probability ½ for each

| | *Hare* | *Stag* | | |
|---|---|---|---|---|
| Payoff | **8** | **16** | **8** | **0** |
| Prob. | 1 | $p_i^2$ | $2p_i(1-p_i)$ | $(1-p_i)^2$ |

- $p_{i+1} = P(\textit{Stag} \text{ vs } \textit{Stag}) \cdot 1 + P(\textit{Hare} \text{ vs } \textit{Hare}) \cdot 0$
  $+ P(\textit{Stag} \text{ vs } \textit{Hare})\,[P(\textit{Stag}\text{'s payoff is } 16) + \tfrac{1}{2} P(\textit{Stag}\text{'s payoff is } 8)]$

  $= p_i^2 + 2p_i(1-p_i)\,[p_i^2 + p_i(1-p_i)] = 3p_i^2 - 2p_i^3$

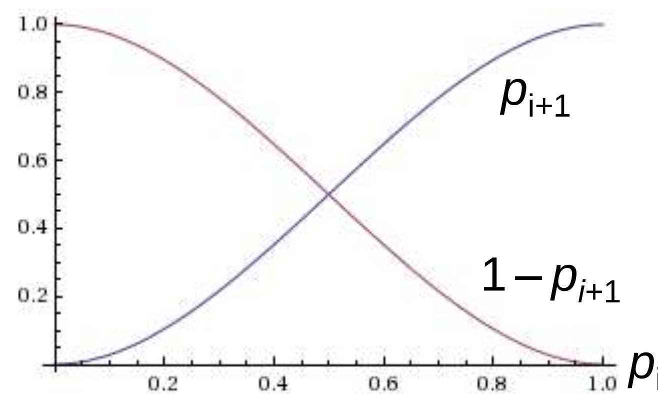


  - Outcome similar to before:
    - If $p_1 > ½$, then $p_j \to 1$
    - If $p_1 < ½$, then $p_j \to 0$
    - If $p_1 = ½$ then $p_i = ½$ for all $i$ (neglecting random variation)
      - Random variation => $p_i \to 0$ or $p_i \to 1$, each equally likely
  - Simulation results similar to before:
    - 101 runs converged to *Stag*, 99 converged to *Hare*

# Double Stag Hunt with *RwS*

- In the Double Stag Hunt, *RwS* does *conditional cooperation*
  - 1st time: hunt stag (risky choice)
  - 2nd time: If payoff was 8 (other hunter cooperated) the 1st time,
    - then hunt hare (safe)
    - otherwise hunt stag (risky)
- Suppose we start with equal numbers of *Stag* and *Hare* agents, and a very small number of *RwS* agents
- Would anyone care to guess what will happen?

# *Stag*, *Hare*, and *RwS*

- 200 simulation runs, starting with 3000 *Stag* agents, 3000 *Hare* agents, 30 *RwS* agents
  - Didn't converge to *RwS*
  - With the replicator dynamic, *RwS* made convergence to *Stag* slightly more likely
  - With the imitate-the-better dynamic, *RwS* made convergence to *Stag* much more likely

| | **Without *RwS*** | | **With 30 *RwS*** | |
|---|---|---|---|---|
| | Replicator | Imitate-the-better | Replicator | Imitate-the-better |
| **Converge to *Stag*** | 101 | 101 | 110 | 138 |
| **Converge to *Hare*** | 99 | 99 | 90 | 62 |
| **Converge to *RwS*** | —— | —— | 0 | 0 |

# *RwS* Catalyzes Growth of *Stag*

- The following effect occurs with both the replicator dynamic and the imitate-the-better dynamic:
  - In the 1st stag hunt, *RwS* plays *Stag*
    - Slightly increases the *Stag* strategy's payoff
  - In the 2nd stag hunt
    - Nearly equal probabilities that *RwS* won or lost the 1st stag hunt
    - => nearly equal probabilities that it will play *Stag* or *Hare*
    - => not much effect on the *Stag* strategy's payoff
  - Overall, a slight advantage for *Stag*
    - => slightly more likely to converge to *Stag*

# *RwS* Catalyzes Growth of *Stag*

- With the imitate-the-better dynamic, *RwS* has another, stronger effect
- Initially, equal numbers of *Stag* and *Hare*
  => *RwS* has an advantage over *Hare* (like *RwS* and *S* in the double lottery)
  => *RwS* agents increase, *Hare* agents decrease
- But fewer *Hare*
  => *Stag* gets higher payoffs
  => *Stag* agents increase
  => *Stag* gets even higher payoffs
- Eventually *Stag* has an advantage over *both RwS* and *Hare*

  => converge to all *Stag*; *RwS* and *Hare* both go extinct

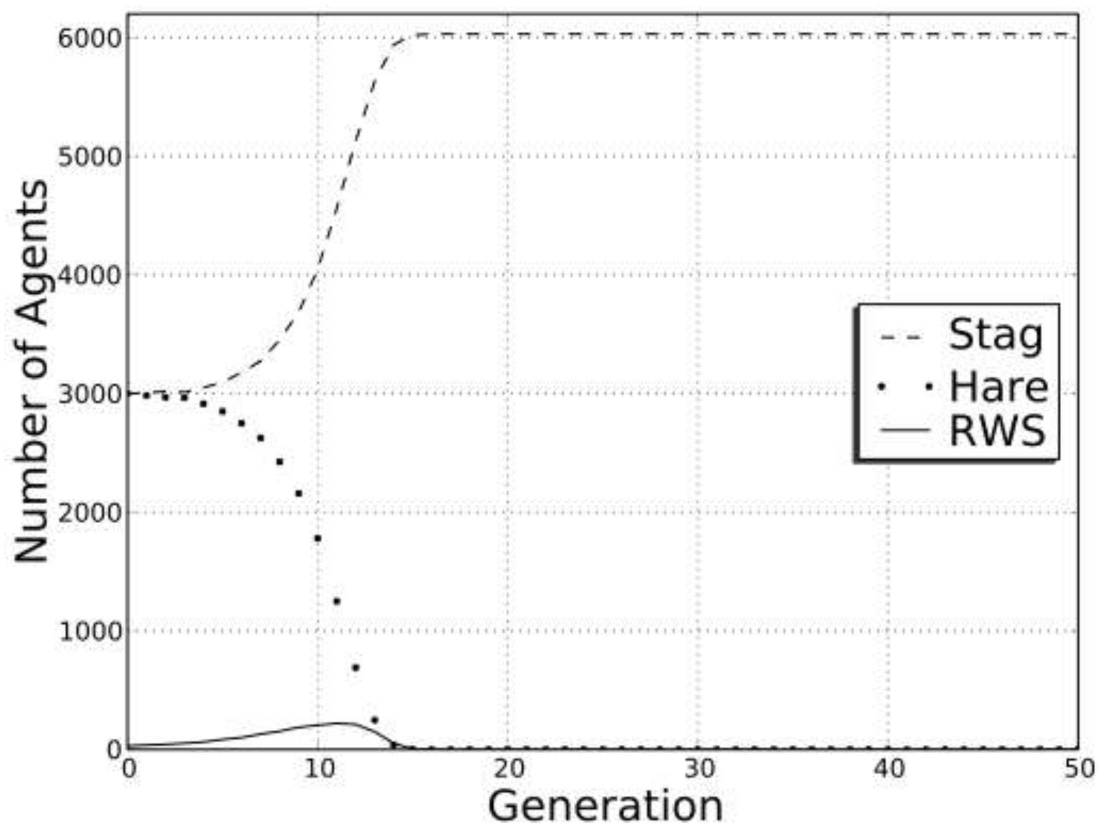


LCCD

# Conclusion

- Initial steps in exploring risk preferences through evolutionary games
- Double lottery game
  - Analogy between *RwS*'s behavior (conditional risk-taking) and human risk preferences
  - With all imitation dynamics except the replicator dynamic, *RwS* has an evolutionary advantage
  - This suggests a possible reason why state-dependent risk preferences might spread
    - But certainly not the only one, and we want to explore others
- Double stag hunt game
  - Example of how to extend our results to games of social cooperation
  - Conditional cooperation (*RwS*) promoted the evolution of cooperation (*Stag*) in a situation where cooperating required a risky decision
    - *RwS* did this more strongly with the imitate-the-better dynamic


- How to reach me
  - Dana Nau, nau@cs.umd.edu
  - http://www.cs.umd.edu/users/nau
- Publications based on this work:
  - P. Roos and D. Nau. Conditionally risky behavior vs expected value maximization in evolutionary games. In *Sixth Conference of the European Social Simulation Association (ESSA 2009)*, Sept. 2009.
  - P. Roos and D. S. Nau. State-dependent risk preferences in evolutionary games. In Chai, Salerno, and Mabry, editors, *Advances in Social Computing: Third International Conference on Social Computing, Behavioral Modeling, and Prediction, SBP 2010*, volume LNCS 6007, pp. 23–31. Springer, Mar. 2010.
  - P. Roos and D. Nau. Risk preference and sequential choice in evolutionary games. *Advances in Complex Systems*, 2010 (to appear).
  - P. Roos, R. Carr, and D. Nau. Evolution of state-dependent risk preferences. *Submitted for journal publication.*

LCCD