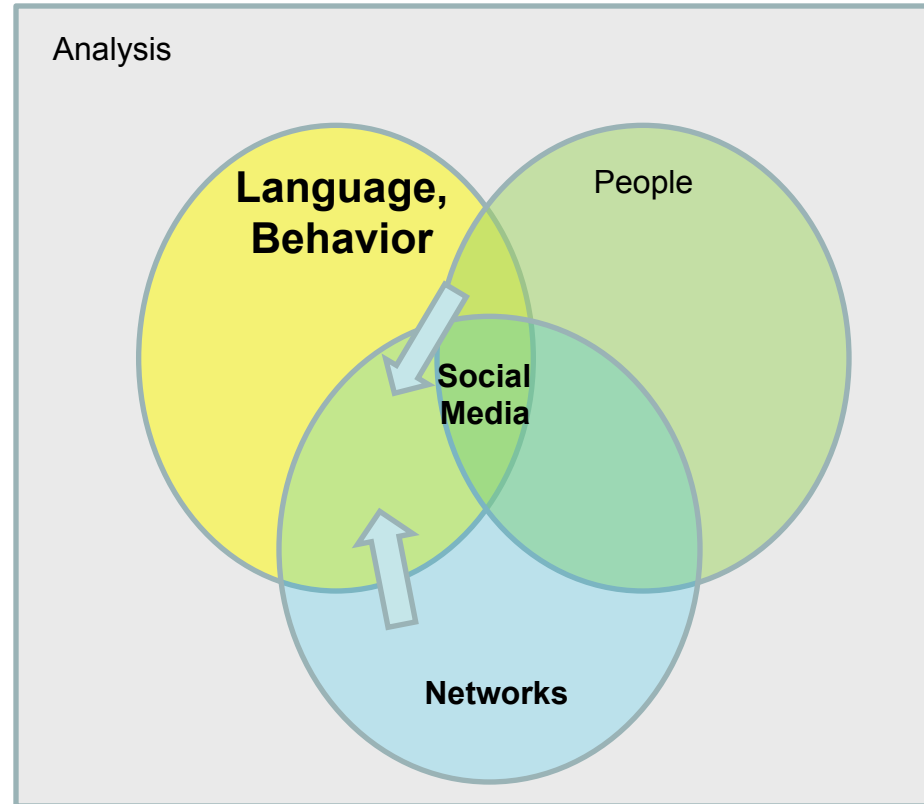


Diffusion in Networks

10-23-2012

Research questions & areas

- How do you model collections of documents in graphs? How do ideas spread through a network?
- ➔ Hybrid models of text and connections
- ➔ Models of diffusion and influence.
 - Viral marketing
 - Collaborative problem-solving.



Diffusion through social networks: *what* things spread

- Behavior
 - Smoking, product purchases, LOLCats, doing research in graphical models,
- Diseases
 - H1N1 flue, bubonic plague,
- Information
 - News, rumors, ...

Diffusion through social networks: *why* things spread


- Fun: i.e., why do things get popular?
 - Fashion, fads, internet memes, research ideas, ...
 - First-order approximation: *preferential attachment* in graphs
- Rational decisions
 - Decisions made publically with limited information
 - Specifically, decisions where *choice* is **public** but some *evidence* used in the choice is **private**
 - Decisions made about products (or behaviors, etc) that have “network effects” (aka “externalities”)
 - Specifically, the benefits and costs of the behavior are **not completely local** to the decision-maker

Review: Barabasi-Albert Networks


- *Science* **286** (1999)
- Start from a small number of node, add a new node with m links
- **Preferential Attachment**
 - Probability of these links to connect to existing nodes is proportional to the node's degree

$$P(k_i) = \frac{k_i}{\sum_j k_j}$$


- 'Rich gets richer'
- For citations: you're more likely to cite something you found from another citation.
- This creates 'hubs': few nodes with very large degrees.



Random graph
(Erdos Renyi)



Preferential attachment
(Barabasi-Albert)



NEW YORK TIMES BESTSELLER

CHRIS ANDERSON


WHY THE FUTURE OF BUSINESS
IS SELLING LESS OF MORE

The LONGER Long Tail

LONGER

INCLUDES A NEW CHAPTER: THE LONG TAIL OF MARKETING

Degree distribution



$$p_k \propto k^{-\alpha}$$

Diffusion through social networks: *why* things spread

- Fun: i.e., why do things get popular?
 - Fashion, fads, internet memes, research ideas, ...
 - First-order approximation: *preferential attachment* in graphs
- Rational decisions:
 - **Decisions made publically with limited information**
 - Specifically, decisions where *choice* is **public** but some *evidence* used in the choice is **private**
 - Decisions made about products (or behaviors, etc) that have “network effects” (aka “externalities”)
 - Specifically, the benefits and costs of the behavior are **not completely local** to the decision-maker

Start with some simple cases in a non-networked world

Making decisions with other people's choices as an input

- Examples:
 - Picking a crowded restaurant in a new city, instead of an empty one
 - You have little information about the food quality
 - But all those people can't be wrong...?
 - Picking a popular course instead of a smaller one
 - ...

Making decisions with other people's choices as an input

- Simple example [Kleinberg 16.2]:
 - An urn with a 50/50 chance of having:
 - $\frac{2}{3}$ red balls and $\frac{1}{3}$ green balls; or
 - $\frac{1}{3}$ red balls and $\frac{2}{3}$ green balls
 - An experiment: in each trial $t=1,2,\dots$
 - Researcher t will
 - private – sample **one** ball (and then replaces it)
 - public – guess what's the majority class (red or green)
 - Everyone that guesses right gets a payoff
 - Nobody can change their guesses

Making decisions with other people's choices as an input

- Possible outcome:
 1. A red ball: the guess is red
 2. A red ball: the guess is red
 3. A green ball: the guess is ____ ?
 4. A green ball: the guess is ____ ?

Rational sequential decision-making leads to a cascade!

Thought experiment: suppose each researcher made a **bet** instead of voting -?

Diffusion through social networks: *why* things spread

- Fun: i.e., why do things get popular?
 - Fashion, fads, internet memes, research ideas, ...
 - First-order approximation: *preferential attachment* in graphs
- Rational decisions:
 - **Decisions made publically with limited information**
 - Specifically, decisions where *choice* is **public** but some *evidence* used in the choice is **private**
 - Decisions made about products (or behaviors, etc) that have “network effects” (aka “externalities”)
 - Specifically, the benefits and costs of the behavior are **not completely local** to the decision-maker

Next, an experiment in simple networked world

Making decisions with other people's choices as an input

[Salganic, Dodds, Watts – Science 2006]

- Artificial music market
 - 48 unknown songs, 14k participants
 - Participants *listen, rate, and (maybe) download*
 - Experiment 1:
 - “Independent world”: choose songs from a grid
 - “Social world” (x8): also see #prior downloads (**weak signal**)
 - Experiment 2:
 - “Independent world”: choose songs from a *list*
 - “Social world” (x8): list sorted by #prior downloads (**stronger signal**)
- Results: based on *average differences in market share*
 - Inequality: compare *different songs in same “world”*
 - Unpredictability: compare *same songs in different “worlds”*

Making decisions with other people's choices as an input

[Salganic, Dodds, Watts – Science 2006]




Fig. 1. Inequality of success for social influence (dark bars) and independent (light bars) worlds for (A) experiment 1 and (B) experiment 2. The success of a song is defined by m_i , its market share of downloads ($m_i = d_i / \sum_{k=1}^S d_k$, where d_i is song i 's download count and S is the number of songs). Success inequality is defined by the Gini coefficient $G = \frac{1}{2S} \sum_{i=1}^S \sum_{j=1}^S |m_i - m_j| / \sum_{k=1}^S m_k$, which represents the average difference in market share for two songs normalized to fall between 0 (complete equality) and 1 (maximum inequality). Differences between independent and social influence conditions are significant ($P < 0.01$) (18).

More social influence → More skewed *distributions* of success

Making decisions with other people's choices as an input


More social influence → Less predictable *which* songs are successful

[Salganic, Dodds, Watts – Science 2006]

Fig. 2. Unpredictability of success for (A) experiment 1 and (B) experiment 2. In both experiments, success in the social influence condition was more unpredictable than in the independent condition. Moreover, the stronger social signal in experiment 2 leads to increased unpredictability. The measure of unpredictability u_i for a single song i is defined as the average difference in market share for that song between all pairs of realizations; i.e.,



$$u_i = \frac{1}{\binom{W}{2}} \sum_{j=1}^W \sum_{k=j+1}^W |m_{i,j} - m_{i,k}|, \text{ where } m_{i,j} \text{ is song } i\text{'s market share in world } j \text{ and } W \text{ is the number of worlds.}$$

The overall unpredictability measure $U = \frac{1}{S} \sum_{i=1}^S u_i$ is then the average of this measure over all S songs. For the independent condition, we randomly split the single world into two subpopulations to obtain differences in market shares, and we then averaged



Making decisions with other people's choices as an input

[Salganic, Dodds, Watts – Science 2006]



Diffusion through social networks: *why* things spread

- Fun: i.e., why do things get popular?
 - Fashion, fads, internet memes, research ideas, ...
 - First-order approximation: *preferential attachment* in graphs
- Rational decisions:
 - Decisions made publically with limited information
 - Specifically, decisions where *choice* is **public** but some *evidence* used in the choice is **private**
 - **Decisions made about products (or behaviors, etc) that have “network effects” (aka “externalities”)**
 - Specifically, the benefits and costs of the behavior are **not completely local** to the decision-maker

Start with some simple cases in a non-networked world


The dying seminar

[Schelling, “Micromotives and Macrobehavior”, p 102]

- Sad story:
 - A reading group starts up on a hot new topic
 - 20 people say they’ d be interested in attending!
 - Eight people show up for the first meeting
 - Six for the second meeting
 - Two for the third meeting
 - » After that nobody bothers
 - Diagnosis: no “critical mass”

The dying seminar


[Schelling, “Micromotives and Macrobehavior”, p 102]



The *value* of attending the seminar for each person depends on the (expected) number of *other people* attending.

The successful seminar


[Schelling, “Micromotives and Macrobehavior”, p 102]



Basic issue:

The *value* of attending the seminar for each person depends on the (expected) number of *other people* attending.

Review: economics 101






Figure 17.1: When there are no network efforts, the demand for a product at a fixed market price p can be found by locating the point where the curve $y = r(x)$ intersects the horizontal line $y = p$.




Figure 17.1: When there are no network efforts, the demand for a product at a fixed market price p can be found by locating the point where the curve $y = r(x)$ intersects the horizontal line $y = p$.




Figure 17.3: Suppose there are network effects and $f(0) = 0$, so that the good has no value to people when no one is using it. In this case, there can be multiple self-fulfilling expectations equilibria: at $z = 0$, and also at the points where the curve $r(z)f(z)$ crosses the horizontal line at height p^* .

Expect $f(0)=0$ and $r(1)=0$

Diffusion through social networks: *why* things spread

- Fun: i.e., why do things get popular?
 - Fashion, fads, internet memes, research ideas, ...
 - First-order approximation: *preferential attachment* in graphs
- Rational decisions:
 - Decisions made publically with limited information
 - Specifically, decisions where *choice* is **public** but some *evidence* used in the choice is **private**
 - Decisions made about products (or behaviors, etc) that have “network effects” (aka “externalities”)
 - Specifically, the benefits and costs of the behavior are **not completely local** to the decision-maker

Start with some simple cases in a non-networked world

Prisoner's dilemma

	Don't confess	Confess
Don't Confess	-1, -1	-10, -1
Confess	-1, -10	-4, -4

“If you confess, and your partner doesn't confess, then you will be released and your partner will be charged with the crime. Your confession will be sufficient to convict him of the robbery and he will be sent to prison for 10 years. If you both confess, then we don't need either of you to testify against the other, and you will both be convicted of the robbery. (Although in this case your sentence will be less — 4 years only — because of your guilty plea.) Finally, if neither of you confesses, then we can't convict either of you of the robbery, so we will charge each of you with resisting arrest. Your partner is being offered the same deal. Do you want to confess?”

Prisoner's dilemma

	Cooperate	Defect
Cooperate	-1, -1	-10, -1
Defect	-1, -10	-4, -4

Good strategy in *repeated* game with another player is “Tit for tat”:

- Round 1: cooperate
- Round n , for $n > 1$: do whatever your partner did on round $n-1$, ie
 - Defect if he defected
 - Cooperate if he cooperated

Why Copy Others? Insights from the Social Learning Strategies Tournament

L. Rendell,^{1*} R. Boyd,² D. Cownden,³ M. Enquist,^{4,5} K. Eriksson,^{5,6} M. W. Feldman,⁷ L. Fogarty,¹ S. Ghirlanda,^{5,8} T. Lillicrap,⁹ K. N. Laland^{1*}

Social learning (learning through observation or interaction with other individuals) is widespread in nature and is central to the remarkable success of humanity, yet it remains unclear why copying is profitable and how to copy most effectively. To address these questions, we organized a computer tournament in which entrants submitted strategies specifying how to use social learning and its asocial alternative (for example, trial-and-error learning) to acquire adaptive behavior in a complex environment. Most current theory predicts the emergence of mixed strategies that rely on some combination of the two types of learning. In the tournament, however, strategies that relied heavily on social learning were found to be remarkably successful, even when asocial information was no more costly than social information. Social learning proved advantageous because individuals frequently demonstrated the highest-payoff behavior in their repertoire, inadvertently filtering information for copiers. The winning strategy (discountmachine) relied nearly exclusively on social learning and weighted information according to the time since acquisition.

Social learning tournament

- Environment: 100-armed bandit
 - 100 actions with very different payoffs (exponentially distributed)
 - Payoffs vary over time (re-drawn with prob p_c)
- Strategies:
 - Exploit: Play an “old” action (previously observed)
 - Innovate: test payoff of a “new” action
 - Observe: observe other players “Exploit” payoff (plus noise)
- Lifecycle:
 - Agents “die” and “reproduce” (based on fitness)
- 104 entries
- Pairwise round-robin, then melee among top 10

Social learning tournament

- Environment: 100-armed bandit
 - 100 actions with very different payoffs (exponentially distributed)
 - Payoffs vary over time (re-drawn with prob p_c)
- Strategies:
 - Exploit: Play an “old” action (previously observed)
 - Innovate: test payoff of a “new” action
 - Observe: observe other players “Exploit” payoff (plus noise)
 - Might return nothing
 - Might sample multiple players in one turn (between 1-12)
 - Lots of observations \sim “parasitic” strategies

Some intuition

- Successful strategies are hard to find
 - if there are 100 “arms” with widely varying payoffs
- Successful agents can’t hide their successful strategies
 - if you Exploit, you can be Observed

What does theory tell us?

- (Asocial) k-armed bandit learning in a changing environment is difficult
- Noise in observations is an expected weakness of social learning, but copying high-payoff behavior is an expected strength
- Social learning might be better for individuals, but worse for a population

What works?

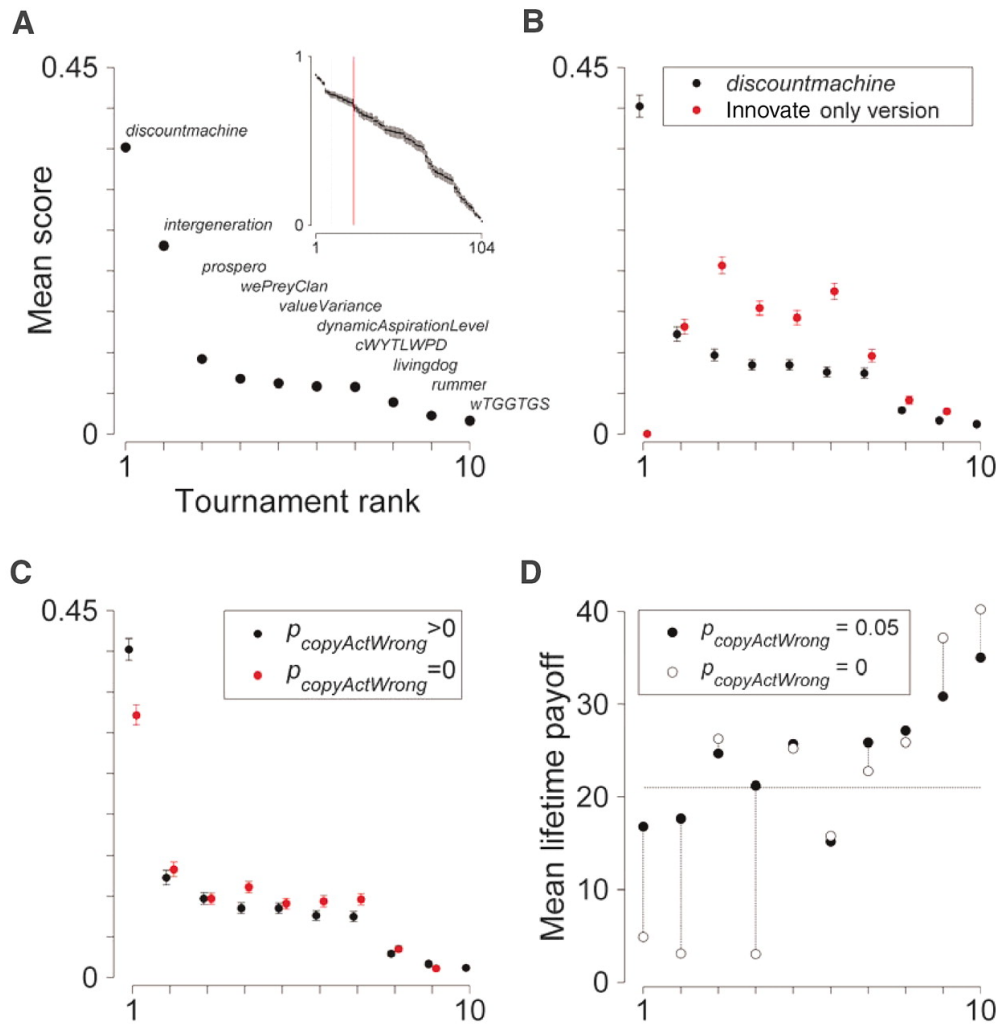
Table 1. Parameters of the AIC best-fit model predicting strategy scores in the first, pairwise, tournament stage. Adjusted $R^2 = 0.76$. Dash entry indicates not applicable.

Predictor	Effect size (β weight)	β	SE	t	$p(> t)$
(Intercept)	—	0.32	0.06	5.43	<0.0001
1) Proportion of learning that is Observe	0.42	0.43	0.06	7.15	<0.0001
2) Variance in rounds to first Exploit*	−0.42	−0.06	0.01	−6.62	<0.0001
Proportion of learning moves	−0.17	−0.34	0.12	−2.79	0.0063
Average rounds between learning moves	0.16	0.01	<0.01	3.09	0.0026
Estimate p_c ? (yes = 1, no = 0)	−0.07	−0.04	0.03	−1.47	0.1452

*We used the natural log of this predictor to give a better linear relationship.

- 1) Observe, don't Innovate
- 2) Exploit what's been learned regularly


Fig. 1 Performance of entered strategies.



Performance of entered strategies. (A) Ranked overall strategy scores in the final stage of the tournament ((Inset) Scores for all 104 entered strategies. B) Ranked scores from those final-stage simulations in which conditions were chosen at random (33), and under the same conditions but with the tournament winner, discountmachine, recoded to learn only with Innovate and never Observe (red). (C) As in (B) but comparing original results with pcopyActWrong fixed at 0 (red). (D) Average individual fitness, measured as mean lifetime payoff, in populations containing only single strategies for each of the final-stage contestants, ranked by tournament placing. Data are average values from the last quarter of single simulations, which were run under the same conditions as the first stage of the tournament and also under the same conditions except with pcopyActWrong = 0. The horizontal dashed line represents the mean lifetime payoff of individuals when all strategies are played together under the same conditions. Strategies relying exclusively on social learning are those ranked 1, 2, and 4.

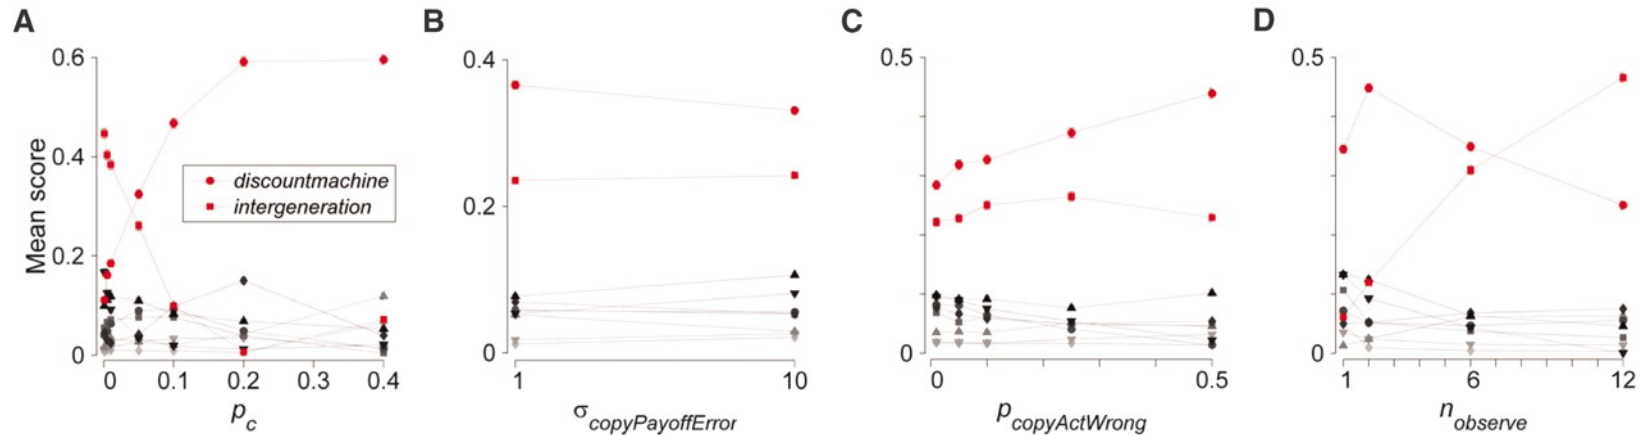
L Rendell et al. Science 2010;328:208-213

Fig. 2 Key variables affecting strategy performance.



L Rendell et al. Science 2010;328:208-213

Fig. 4 Social learning dominates irrespective of cost across a broad range of conditions.

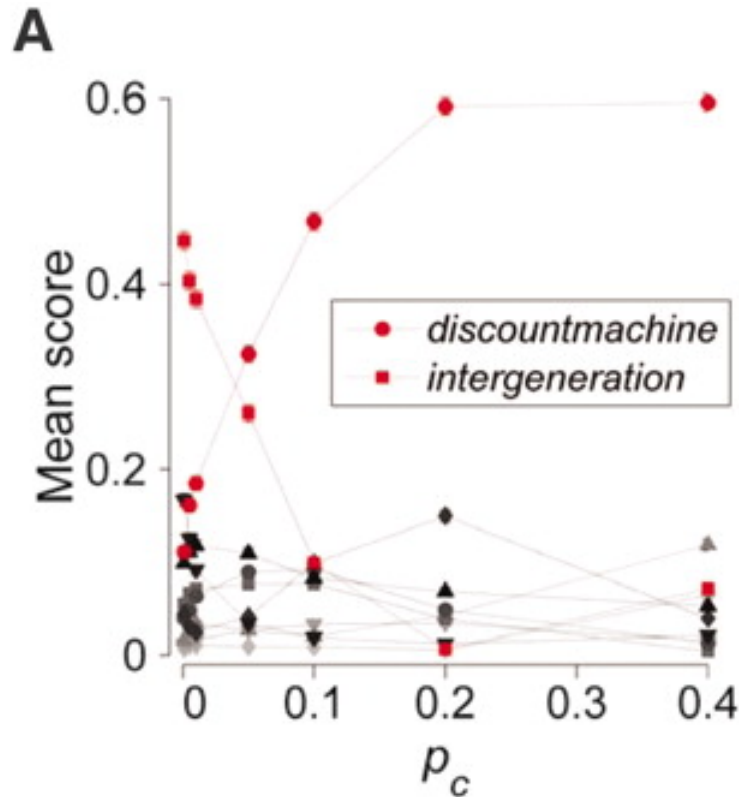


Social learning dominates irrespective of cost across a broad range of conditions. Plots show mean strategy scores (±variance) across systematic melee conditions with respect to (A) variation in the rate of environmental change (p_c), (B) $\sigma_{\text{copyPayoffError}}$, the standard deviation of a normally distributed error applied to payoffs returned by Observe, (C) $p_{\text{copyActWrong}}$, the probability that Observe returned a behavior selected, at random from those not actually observed, and (D) the number of other agents sampled when playing Observe (n_{observe}). First and second place strategies are highlighted; the rank of the other strategies is indicated by shading with darker shading indicating higher rank. Error bars are ± SEM but mostly not visible because all SEMs < 0.01.

L Rendell et al. Science 2010;328:208-213



Fig. 4 Social learning dominates irrespective of cost across a broad range of conditions.




Social learning dominates irrespective of cost across a broad range of conditions. Plots show mean strategy scores (\pm variance) across systematic melee conditions with respect to (A) variation in the rate of environmental change (p_c), (B) $\sigma_{\text{copyPayoffError}}$, the standard deviation of a normally distributed error applied to payoffs returned by Observe, (C) $p_{\text{copyActWrong}}$, the probability that Observe returned a behavior selected, at random from those not actually observed, and (D) the number of other agents sampled when playing Observe (nobserve). First and second place strategies are highlighted; the rank of the other strategies is indicated by shading with darker shading indicating higher rank. Error bars are \pm SEM but mostly not visible because all SEMs < 0.01 .

L Rendell et al. Science 2010;328:208-213



Fig. 4 Social learning dominates irrespective of cost across a broad range of conditions.

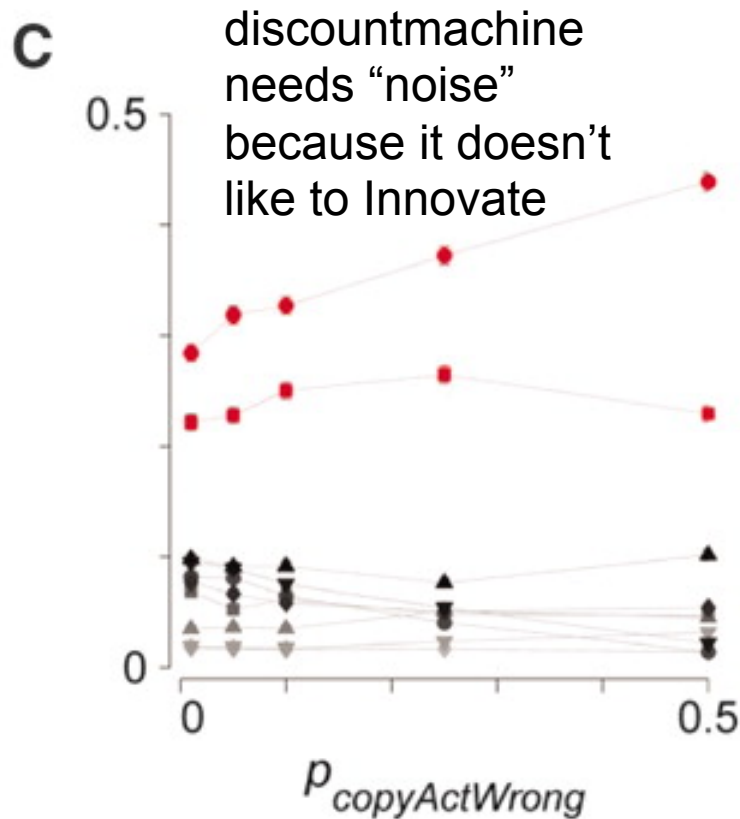


Social learning dominates irrespective of cost across a broad range of conditions. Plots show mean strategy scores (\pm variance) across systematic melee conditions with respect to (A) variation in the rate of environmental change (pc), (B) $\sigma(\text{copyPayoffError})$, the standard deviation of a normally distributed error applied to payoffs returned by Observe, (C) $p(\text{copyActWrong})$, the probability that Observe returned a behavior selected, at random from those not actually observed, and (D) the number of other agents sampled when playing Observe (nobserve). First and second place strategies are highlighted; the rank of the other strategies is indicated by shading with darker shading indicating higher rank. Error bars are \pm SEM but mostly not visible because all SEMs < 0.01 .

L Rendell et al. Science 2010;328:208-213




Fig. 4 Social learning dominates irrespective of cost across a broad range of conditions.



Social learning dominates irrespective of cost across a broad range of conditions. Plots show mean strategy scores (\pm variance) across systematic melee conditions with respect to (A) variation in the rate of environmental change (p_c), (B) $\sigma_{copyPayoffError}$, the standard deviation of a normally distributed error applied to payoffs returned by Observe, (C) $p(copyActWrong)$, the probability that Observe returned a behavior selected, at random from those not actually observed, and (D) the number of other agents sampled when playing Observe ($n_{observe}$). First and second place strategies are highlighted; the rank of the other strategies is indicated by shading with darker shading indicating higher rank. Error bars are \pm SEM but mostly not visible because all SEMs < 0.01 .

L Rendell et al. Science 2010;328:208-213

Fig. 4 Social learning dominates irrespective of cost across a broad range of conditions.



Social learning dominates irrespective of cost across a broad range of conditions. Plots show mean strategy scores (\pm variance) across systematic melee conditions with respect to (A) variation in the rate of environmental change (pc), (B) $\sigma_{copyPayoffError}$, the standard deviation of a normally distributed error applied to payoffs returned by Observe, (C) $p_{copyActWrong}$, the probability that Observe returned a behavior selected, at random from those not actually observed, and (D) **the number of other agents sampled when playing Observe (nobserve)**. First and second place strategies are highlighted; the rank of the other strategies is indicated by shading with darker shading indicating higher rank. Error bars are \pm SEM but mostly not visible because all SEMs < 0.01 .

L Rendell et al. Science 2010;328:208-213



What does the experiment say?

- Copying works!
- Noise in observations is an expected weakness of social learning, but copying high-payoff behavior is an expected strength?
 - No: actually, error does not affect success of social-learning (much)
- Social learning might be better for individuals, but worse for a population
 - No: mean individual fitness is higher for populations that do more social learning in mixed-strategy populations....
 - ...but it is lower with a population of poor social learners
 -and mean payoff with mixed-strategy populations is lower than in a ground of poor social learners
- Inverse correlation between mean individual fitness of population “pure” for population s and how well s did in the tournament