

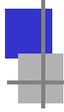


Title	Adaptive Bases of Human Rationality
Author(s)	Takahashi, Nobuyuki
Citation	SOCREAL 2007: Proceedings of the International Workshop on Philosophy and Ethics of Social Reality, Sapporo, Japan, 2007 / editor Tomoyuki Yamada, 1(71)-34(104)
Issue Date	2007
Doc URL	http://hdl.handle.net/2115/29940
Type	proceedings
Note	SOCREAL 2007: International Workshop on Philosophy and Ethics of Social Reality. Sapporo, Japan, 2007-03-09/10. Special Guest Lecture
File Information	Takahashi.pdf



[Instructions for use](#)

Adaptive Bases of Human Rationality



Nobuyuki Takahashi
Department of Behavioral Science
The Center for the Study of Cultural and
Ecological Foundations of the Mind
Hokkaido University

Part 1

General framework of CEFOM/21

21st Century COE "Cultural and Ecological Foundations of the Mind"



Rationality and Adaptationist Perspective

In studying human behavior, rationality has been one of the most central concepts in behavioral and social sciences, but ...

Normative camp (economics, mathematical modeling, etc.)

People should behave in a rational manner.



Empirical camp (psychology, cognitive science, etc.)

In reality people often behave in an irrational manner.



No predictable power at the macro level



Impractical desktop theory

Rationality and Adaptationist Perspective

This situation has begun to change since 1990s.

- The rise of experimental/behavioral economics
- The wide acceptance of evolutionary perspectives in psychology
- Interdisciplinary collaborations using evolutionary game theory as a common language

The currently shared idea

The myth of "tabula rasa" has been almost completely refuted by empirical evidence.

The human mind is not, in fact, infinitely malleable but instead works under certain constraints.

Rationality and Adaptationist Perspective

Why do certain constraints exist?

- We are endowed with psychological mechanisms that enable us to behave in certain ways **as adaptive tools for the social environment**.
- These psychological mechanisms have allowed humans to build and maintain societies.

✓The few research centers (e.g., George Mason University, Zurich University, the Max Planck Institute, the Santa Fe Institute, UCLA and UCSB) share the basic perspective.

✓CEFOM/21 (Center for the Study of Cultural and Ecological Foundations of the Mind, a 21st Century Center of Excellence) has been one of the few research centers around the world to study **the adaptive bases of human rationality both theoretically and empirically** since its establishment in 2002.

Part 2

From empirical findings to theory

Why do people cooperate in a one-shot PDG?

Puzzle

- Prisoner's dilemma game (PDG) is the most well-known game in research on human cooperation.

The most significant finding in empirical studies is that people actually cooperate very often even in a one-shot PDG where cooperation is irrational.

Why do people cooperate in a one-shot PDG?

Prisoner's Dilemma Game (PDG)

- There are two players, A and B.
- Each player has two behavioral choices: **Cooperation** or **Defection**.
- Each player's payoff is determined by a joint decision.
- Player A's payoff is indicated above and to the right of the diagonal, and player B's payoff is represented below and to the left.

Dominant choice: Defection

Defection produces an individually better outcome no matter what the choice of the partner is.

Player B's Choice	Player A's Choice	
	Cooperation	Defection
C	2 / 2	0 / 3
D	3 / 0	1 / 1

Why do people cooperate in a one-shot PDG?

So many people actually choose cooperation.

- A rational player should choose defection.
- However, even in a one-shot PD played between unrelated players under complete anonymity, **many people actually choose cooperation.**
- This is one of the most basic findings in experimental research during the past few decades.



Are ordinary people irrational?

Player B's Choice	Player A's Choice	
	Cooperation	Defection
C	2, 2	0, 3
D	3, 0	1, 1

Why do people cooperate in a one-shot PDG?

Choosing cooperation is not a simple logical error.

Traditional answer: Cooperation is a result of **confusion** (c.f., Andreoni, 1995).

People fail to understand the incentive structure of the game that makes defection a dominant choice.



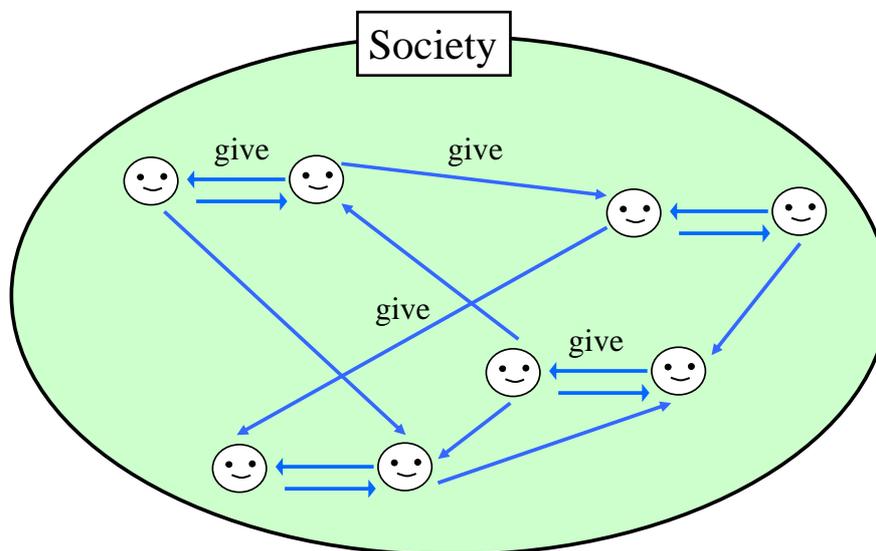
Our answer: **Social Exchange Heuristic (SEH)**

(Kiyonari, Tanida, & Yamagishi, 2000)

It is triggered by the construal of the situation as one involving **social exchange**, and once triggered, **subjectively transforms the nature of the exchange**, resulting in the perception of the situation as an **assurance game**. With the transformed perception of the situation, **people come to intuitively believe in the desirability of mutual cooperation.**

Why do people cooperate in a one-shot PDG?

Social exchange is the foundation of society.



Why do people cooperate in a one-shot PDG?

Social exchange is the foundation of society.

Society

- Everyday, everybody is engaged in social exchange.
Social exchange – Giving / receiving resources
Resources can be anything from economic resources (e.g., money) to psychological/sociological resources (e.g., love, respect).
- Social exchange has been the most important activity throughout the history of human beings.
- Social exchange has a characteristic of PD. Although there is a possibility of being exploited, successful exchange brings mutual benefit.

Why do people cooperate in a one-shot PDG?

PDG as social exchange

In a social exchange relation, each party pays or does not pay a cost, c , to provide the other partner a benefit, b , where $b > c$.

Social exchange	PDG
Paying the cost (c)	Cooperation
Not paying the cost	Defection

A PDG matrix representing four outcomes of a social exchange

Player B's Choice	Player A's Choice	
	Cooperation	Defection
C	$b-c$ / $b-c$	$-c$ / b
D	b / $-c$	0 / 0

$b > b-c > 0 > -c$
 Temptation payoff Reward payoff Punishment payoff Sucker payoff

Why do people cooperate in a one-shot PDG?

Assurance Game (AG) and Prisoner's Dilemma Game (PDG)

Defection is the dominant choice in PDG.



There is no dominant choice in AG.

- When the partner defects, defection produces an individually better outcome.
- However, **when the partner cooperates, cooperation produces an individually better outcome.**

Assurance Game

Player B's Choice	Player A's Choice	
	Cooperation	Defection
C	2 / 2	0 / 1
D	1 / 0	1 / 1

Prisoner's Dilemma Game

Player B's Choice	Player A's Choice	
	Cooperation	Defection
C	2 / 2	0 / 3
D	3 / 0	1 / 1

Why do people cooperate in a one-shot PDG?

Heuristic as a rule of thumb

Heuristics are simple, efficient decision-making rules that work well under most circumstances, but in certain cases lead to systematic cognitive biases.

Why do people cooperate in a one-shot PDG?

Confusion vs. Social Exchange Heuristic

Confusion

Cooperation is irrational. People choose cooperation **because they are confused** (e.g., they are not smart enough, the instructions during the experiment are too vague). Therefore, **when they are not confused, they would choose defection.**

Social Exchange Heuristic (SEH)

It is an important design feature of human cognitive functioning that makes mutual cooperation possible in social exchange.

People choose cooperation when SEH is triggered. When it is triggered, **it subjectively transforms PD into AG**, and motivates people to seek mutual cooperation.

In order to distinguish which explanation is better, a series of experiments was conducted.

Why do people cooperate in a one-shot PDG?

Experimental Design

Reality (high, middle, low) × Partner's cooperation (known, unknown)

A) Realistic sense of exchange

Confusion – Realistic sense of exchange decreases confusion.
→ Realistic sense of exchange **decreases cooperation**.

SEH – Realistic sense of exchange triggers SEH.
→ Realistic sense of exchange **increases cooperation**.

B) Partner's cooperation

If players know that their partners decided to cooperate...

Confusion – Players face a less complex decision task than those who make decisions simultaneously.
→ Players **cooperate less**.

SEH – Sequential nature of the game promotes the sense of exchange, which triggers SEH.
→ Players **cooperate more**.

Why do people cooperate in a one-shot PDG?

Manipulation of Reality

Full experiment – The participants played a PDG once and were paid exactly the amount specified in the payoff matrix.

Reality – High

Vignette experiment – The participants were told to **imagine** that they had been participating in the experiment described in the vignette, and to decide whether they would have cooperated or defected if they had been in the experiment.

□ Vignette with money

Participants were shown the payoff matrix expressed **in terms of monetary value** (“yen”) as in the original, full experiment.

Reality – Middle

□ Vignette with score

Participants were shown the payoff matrix expressed **in terms of scores** (“points”). There was no mention of money.

Reality – Low

Why do people cooperate in a one-shot PDG?

Manipulation of partner's cooperation

What if players know that their partners decided to cooperate?

➤ Simultaneous-game condition

– Participants and their partners make their decisions simultaneously.

Players don't know what their partners would do.

➤ 2nd-player condition

– Participants were told that their partners would make the decision first. **After informed that their partners had chosen C, they made their own decisions.**

Players know that their partners have already decided to cooperate.

Why do people cooperate in a one-shot PDG?

Experimental Design

Reality (high, middle, low) × Partner's cooperation (known, unknown)

Payoff matrix used in the experiment

Participant's choice	Partner's choice	
	C	D
C	1200 / 1200	0 / 1800
D	1800 / 0	600 / 600

Why do people cooperate in a one-shot PDG?

Hypothesis regarding reality

Confusion – Realistic sense of exchange decreases confusion.
→ Realistic sense of exchange **decreases cooperation.**

→ Cooperation rate ...

Full < Vignette with money < Vignette with score
Reality (High) (Middle) (Low)

SEH – Realistic sense of exchange triggers SEH.
→ Realistic sense of exchange **increases cooperation.**

→ Cooperation rate ...

Full > Vignette with money > Vignette with score
Reality (High) (Middle) (Low)

Why do people cooperate in a one-shot PDG?

Hypothesis regarding partner's cooperation

If players know that their partners decided to cooperate...

Confusion – Players face a less complex decision task than those who make decisions simultaneously.
→ Players **cooperate less.**

→ Cooperation rate ...

Simultaneous > 2nd-player
(unknown) (known)

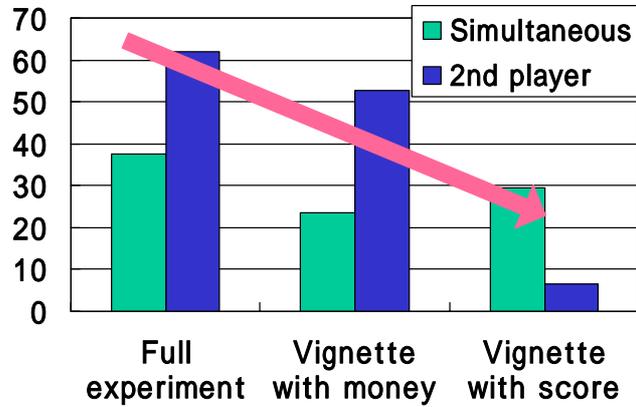
SEH – Sequential nature of the game promotes the sense of exchange, which triggers SEH.
→ Players **cooperate more.**

→ Cooperation rate ...

Simultaneous < 2nd-player
(unknown) (known)

Why do people cooperate in a one-shot PDG?

Result (1)



Cooperation rate ...

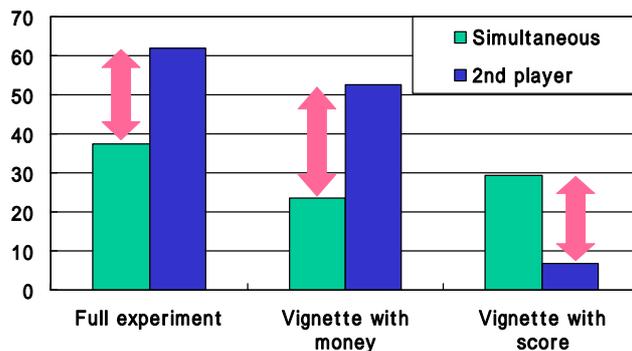
Full > Vignette with money > Vignette with score

Realistic sense of exchange promoted cooperation.

→ Confusion was refuted, and SEH was supported.

Why do people cooperate in a one-shot PDG?

Result (2)



In the full experiment and the vignette with money

Simultaneous < 2nd-player

→ Confusion was refuted, and SEH was supported.

In the vignette experiment with score,

Simultaneous > 2nd-player

→ Only when the outcome of the game was truly trivial, participants acted rationally.

Why do people cooperate in a one-shot PDG?

Is SEH Adaptive?

It is clear that people have SEH. However, by definition, SEH produces a logical error.



How can it be adaptive?

Error management as an adaptationist explanation for heuristics

- Decision-making under **uncertainty** often results in erroneous inferences, but some errors are more costly than others.
- SEH helps people reduce the likelihood of one type of error while increasing the likelihood of another type of error in social exchanges.

Why do people cooperate in a one-shot PDG?

Uncertain nature of social exchange

If this exchange opportunity is truly a one-shot game with no possibility of reputation, **D** is the dominant choice.

Player B's Choice	Player A's Choice	
	Cooperation	Defection
C	b-c / b-c	-c / b
D	b / -c	0 / 0

However two possibilities of receiving sanctions

- Possibility a) This can be **the beginning of a long-term relationship between you and your partner.**
 ⇒ Choosing **C** is adaptive as TFT dictates (e.g., Axelrod, 1984).
- Possibility b) **Information about your behavior (reputation) could potentially spread to other members of the community.**
 ⇒ Choosing **C** is adaptive in order to avoid ostracism from a community (Davis and Holt, 1993).

Managing our behaviors under such uncertainty is one of the most important adaptive tasks humans face.

Why do people cooperate in a one-shot PDG?

Our decisions can produce two types of errors

The consequences of 4 possible states of the social exchange

Inferences	True nature of exchanges	
	Defection is sanctioned	Defection is not sanctioned
Defection is sanctioned (punishment/ostracism)	<i>Correct inference</i> Choosing C Outcome: Gains from mutual cooperation	<i>Type I Error</i> Choosing C Outcome: Failure in saving the cost of cooperation
Defection is not sanctioned	<i>Type II Error</i> Choosing D Outcome: punishment and/or ostracism for defection	<i>Correct inference</i> Choosing D Outcome: Savings in the cost of cooperation

The cost of Type II error is huge in human societies.

Having SEH (i.e., perceiving an exchange situation under uncertainty as one in which defection is sanctioned) is adaptive.

SEH is a systematic bias in favor of committing **Type II errors** less often despite the cost of the increased probability of committing Type I errors.

Why do people cooperate in a one-shot PDG?

Summary of Part 2

Cooperation in a one-shot PD is a logical error.

However, many people actually cooperate.

Experimental results showed that ...

- Realistic sense of exchange promoted cooperation.
- People cooperated more when the partner has already cooperated in the sequential game than in the simultaneous game.

Humans have Social Exchange Heuristic.

Making logical errors is more adaptive than making the logically correct decisions in the “real” environment.

Part 3

From theory to empirical findings

What kind of altruism can be adaptive?

Altruism as unilateral resource giving

Altruistic behavior × Altruistic motivation



What kind of altruism can be adaptive?

Altruism as unilateral resource giving

Altruistic behavior × Altruistic motivation



What kind of altruism can be adaptive?

Altruism as unilateral resource giving

Altruistic behavior × Altruistic motivation



B loses his resource.

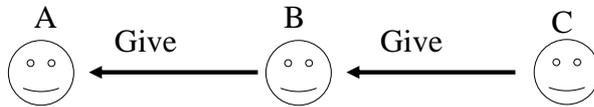
Giving is costly.

How can B's unilateral resource giving be rational?

What kind of altruism can be adaptive?

Indirect reciprocity

Reciprocation not by the recipient but by the third party



If B's giving is reciprocated not by A but by another person (C), B's giving can be rational.

← Indirect reciprocity

However, it has been shown that not indiscriminate giving but **discriminate giving** is necessary for the emergence of indirect reciprocity.

What kind of discriminate altruism can maintain indirect reciprocity?

What kind of altruism can be adaptive?

1. Theory – Mathematical analysis, computer simulation
2. Empirical findings

What kind of altruism can be adaptive?

Mathematical analysis and computer simulation using evolutionary game theory

Framework of giving game

- A pair consisting of a **donor** and **recipient** is chosen randomly from a population.
- The donor decides whether to give his resource to his recipient with a cost of c . When a donor gives, the recipient receives the benefit b ($b > c$).
- Each player has a reputation score S which has two values: “Good” or “Bad.”
 - The donor gives if he thinks the recipient’s score is Good.
 - The donor doesn’t give if he thinks the recipient’s score is Bad.

What kind of altruism can be adaptive?

Mathematical analysis and computer simulation using evolutionary game theory

Framework of giving game

- The donor decides recipient’s score by using two types of information based on recipient’s previous behavior as a donor.
 - (a) Their previous behavior (**gave** or **didn’t give**)
 - (b) Previous recipient’s reputation score (**Good** or **Bad**)

Four types of recipient

- Each player has 4 rules that determine whether each type of recipients is considered good or bad.
- Strategies: sets of 4 rules – 16 strategies are possible.

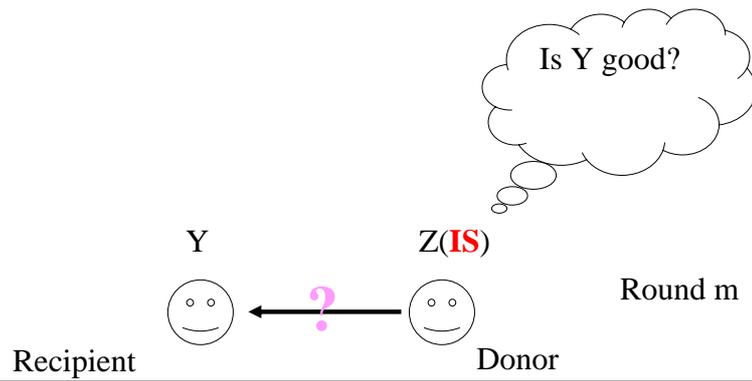
Strategies dictate what kind of recipients are regarded as good.

(a) Current recipient’s previous behavior	(b) Previous recipient’s S	
	Good	Bad
Gave	Good or Bad	Good or Bad
Didn’t give	Good or Bad	Good or Bad

What kind of altruism can be adaptive?

Previously proposed solutions (1)

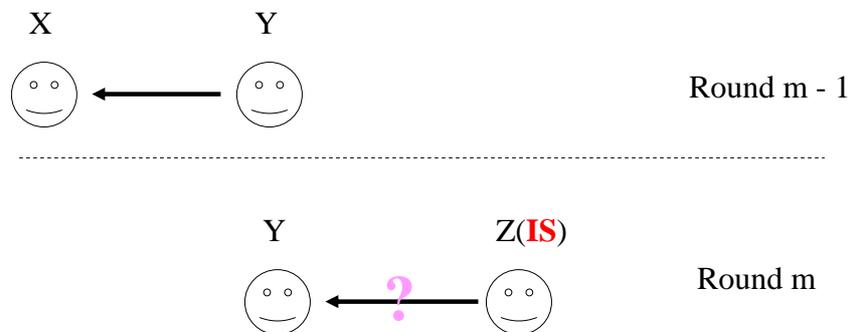
Image Scoring Strategy (Nowak and Sigmund, 1998a, b)



What kind of altruism can be adaptive?

Previously proposed solutions (1)

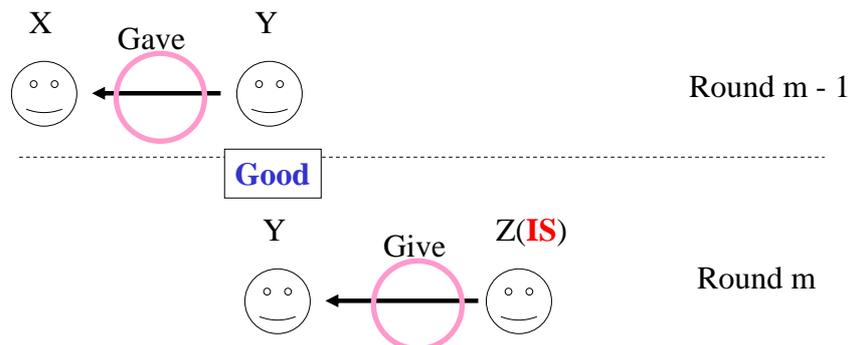
Image Scoring Strategy (Nowak and Sigmund, 1998a, b)



What kind of altruism can be adaptive?

Previously proposed solutions (1)

Image Scoring Strategy (Nowak and Sigmund, 1998a, b)

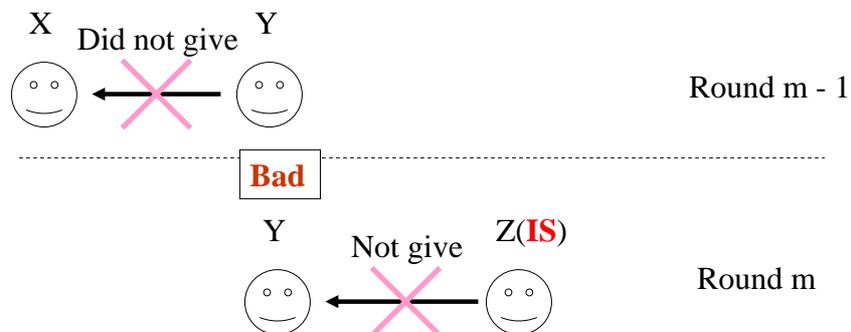


What kind of altruism can be adaptive?

Previously proposed solutions (1)

Image Scoring Strategy (Nowak and Sigmund, 1998a, b)

- IS** gives the recipient who gave to his/her recipient, and does not give to the recipient who did not give in the previous round.
- IS** is a variant of Tit-For-Tat.



What kind of altruism can be adaptive?

Previously proposed solutions (1)

However...

Image scoring strategy (GGBB) was rejected by later studies (Leimar and Hammerstein, 2001; Panchanathan and Boyd, 2003).

(a) Current recipient's previous behavior	(b) Previous recipient's S	
	Good	Bad
Gave	Good	Good
Didn't give	Bad	Bad

What kind of altruism can be adaptive?

Previously proposed solutions (2)

Standing Strategy

(Leimar & Hammerstein, 2001; Panchanathan & Boyd, 2003)

- Like IS, **standing** gives to the recipient who gave to his/her recipient in the previous round.
- Unlike IS, **standing** uses **2nd order information**.
Standing does not always consider not-giving “**bad**.”
Standing distinguishes between justifiable not-giving and unjustifiable not-giving.

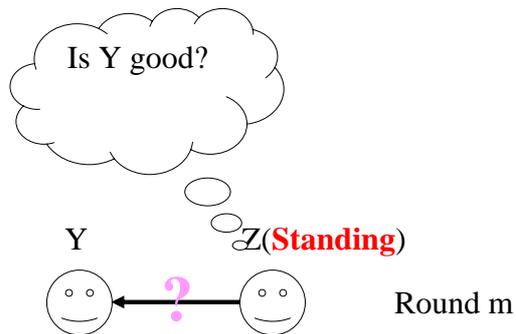
(a) Current recipient's previous behavior	(b) Previous recipient's S	
	Good	Bad
Gave	Good	Good
Didn't give	Bad	Good

What kind of altruism can be adaptive?

Previously proposed solutions (2)

Justifiable not-giving

If the recipient did not give to a **bad** person, **standing** gives to the recipient.

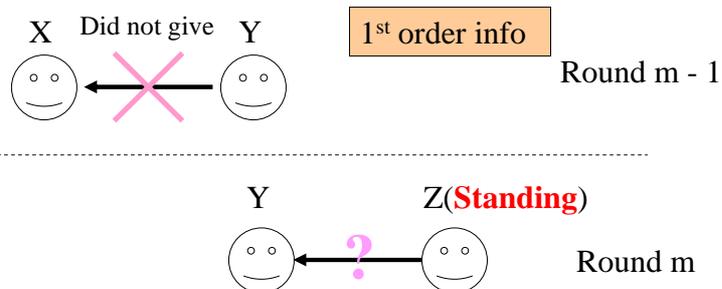


What kind of altruism can be adaptive?

Previously proposed solutions (2)

Justifiable not-giving

If the recipient did not give to a **bad** person, **standing** gives to the recipient.

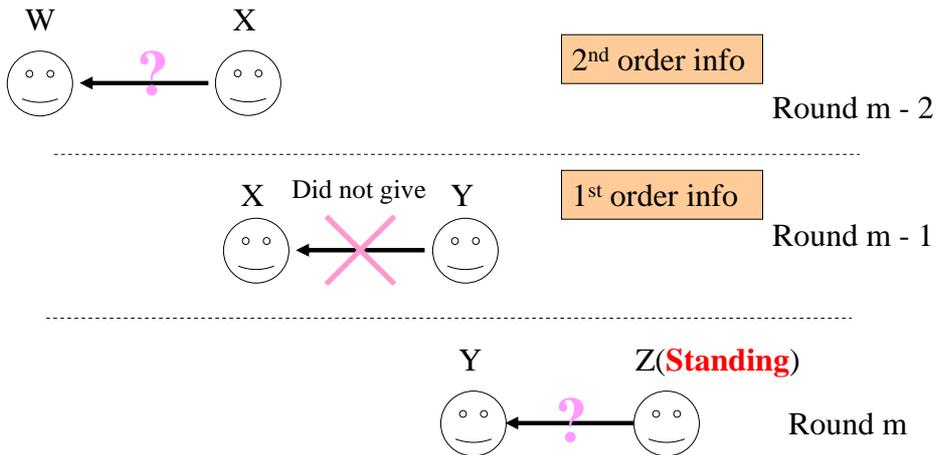


What kind of altruism can be adaptive?

Previously proposed solutions (2)

Justifiable not-giving

If the recipient did not give to a **bad** person, **standing** gives to the recipient.

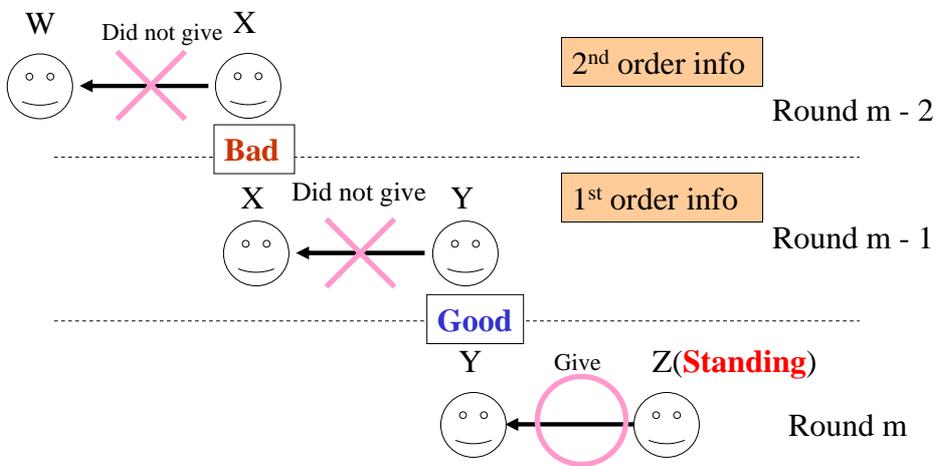


What kind of altruism can be adaptive?

Previously proposed solutions (2)

Justifiable not-giving

If the recipient did not give to a **bad** person, **standing** gives to the recipient.

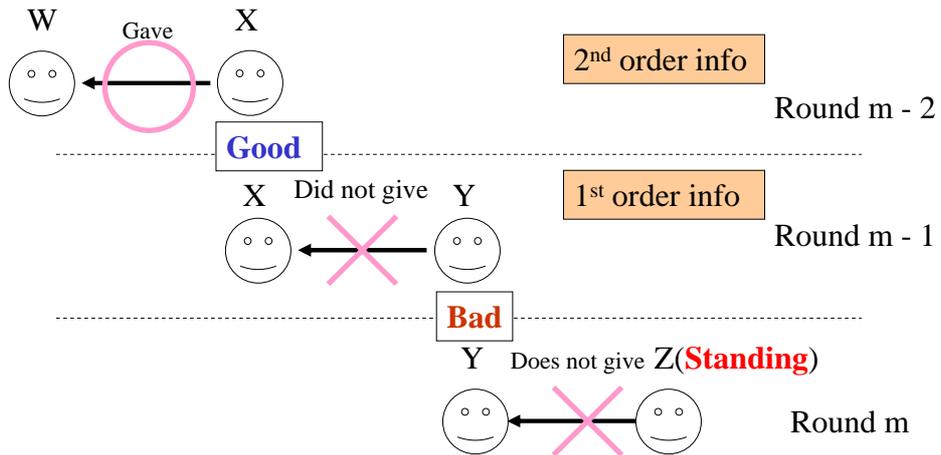


What kind of altruism can be adaptive?

Previously proposed solutions (2)

Unjustifiable not-giving

If the recipient did not give to a **Good** person, **standing** does not give to the recipient.



What kind of altruism can be adaptive?

Previously proposed solutions (2)

However...

Standing strategy (GGBG) was rejected by later studies (Takahashi and Mashima, 2006).

(a) Current recipient's previous behavior	(b) Previous recipient's S	
	Good	Bad
Gave	Good	Good
Didn't give	Bad	Good

What kind of altruism can be adaptive?

Common characteristics of rejected strategies

It does not matter whom one gives to.

Giving is always regarded as **Good**.

Image scoring strategy (GGBB)

(a) Current recipient's previous behavior	(b) Previous recipient's S	
	Good	Bad
Gave	Good	Good
Didn't give	Bad	Bad

Standing strategy (GGBG)

(a) Current recipient's previous behavior	(b) Previous recipient's S	
	Good	Bad
Gave	Good	Good
Didn't give	Bad	Good

What kind of altruism can be adaptive?

New solutions (Takahashi and Mashima, 2006)

Strict discriminator strategy (SDISC) and **Extra standing** strategy (ES) (Takahashi and Mashima, 2003)

(a) Current recipient's previous behavior	(b) Previous recipient's S	
	Good	Bad
Gave	Good	Bad
Didn't give	Bad	Bad

(a) Current recipient's previous behavior	(b) Previous recipient's S	
	Good	Bad
Gave	Good	Bad
Didn't give	Bad	Good

•Unlike IS and Standing, **SDISC** and **ES** does not always consider giving “good.”



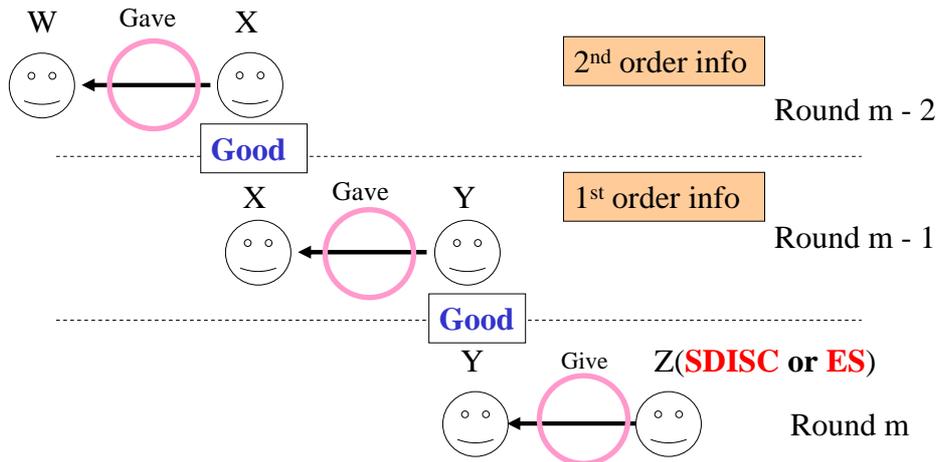
They distinguishes between “good” givers and “bad” givers.

What kind of altruism can be adaptive?

Strict discriminator strategy and extra standing strategy

Helping a “good” person is “good”!

If the recipient gave to a **good** recipient, **SDISC** or **ES** gives to the recipient.

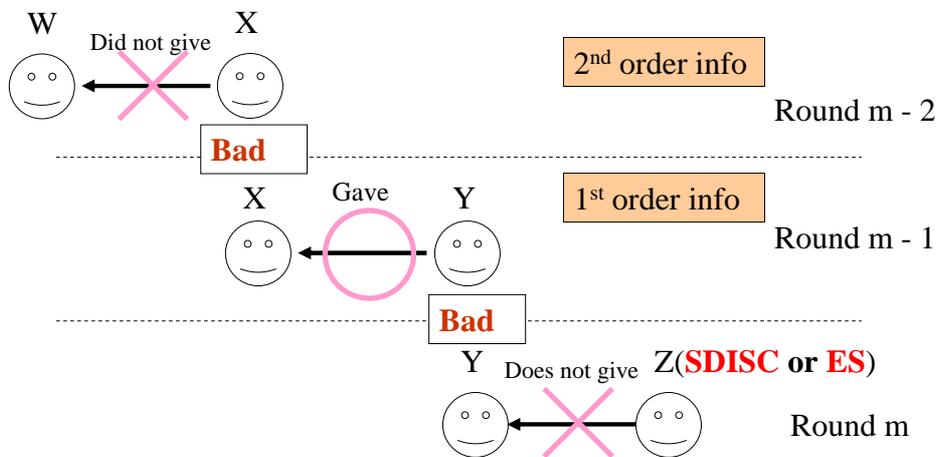


What kind of altruism can be adaptive?

Strict discriminator strategy and extra standing strategy

Helping a “good” person is “good”!

If the recipient gave to a **bad** recipient, **SDISC** or **ES** does not give to the recipient.



What kind of altruism can be adaptive?

Common characteristics of new solutions

- Unlike IS and Standing, **SDISC** and **ES** do not give to previous givers who gave to a “**bad**” recipient.



They regard giving to **Bad** as **Bad**.

SDISC and **ES** regards “saints” (indiscriminate altruists) as “traitors.”

Strict discriminator strategy (GBBB) Extra standing strategy (GBBG)

(a) Current recipient's previous behavior	(b) Previous recipient's S	
	Good	Bad
Gave	Good	Bad
Didn't give	Bad	Bad

(a) Current recipient's previous behavior	(b) Previous recipient's S	
	Good	Bad
Gave	Good	Bad
Didn't give	Bad	Good

Three keys to make indirect reciprocity possible

- (a) To regard a person who didn't give despite having had a chance to give to a **Good** person as “**Bad**” (i.e., to exclude a free rider)
- (b) To regard a person who gave to a **Good** person as “**Good**” (i.e., to help a conditional giver)
- (c) To regard a person who gave to a **Bad** person as “**Bad**” (i.e., to exclude an unconditional giver).

The most important characteristic for the emergence of indirect reciprocity

(a) Current recipient's previous behavior	(b) Previous recipient's S	
	Good	Bad
Gave	Good	Bad
Didn't give	Bad	Bad

(a) Current recipient's previous behavior	(b) Previous recipient's S	
	Good	Bad
Gave	Good	Bad
Didn't give	Bad	Good

What kind of altruism can be adaptive?

1. Theory – Mathematical analysis, computer simulation
2. Empirical findings

What kind of altruism can be adaptive?

Do people actually regard those who gave to bad as bad and exclude them?



Experiment

What kind of altruism can be adaptive?

Experimental setting

8-person Giving game

- ✓ There were 17 rounds.
- ✓ In each round, participants were endowed 50 yen and asked to decide whether to
 - ✓ give it to one of the other 7 participants (The endowment they gave were doubled and given to their recipients: recipients received 100 yen) OR
 - ✓ keep it for themselves (They received 50 yen).
- ✓ After each round, they received the feedback information (=how much they received in the round).

What kind of altruism can be adaptive?

Experimental setting

8-person Giving game

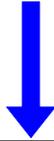
- ✓ There were 17 rounds.
- ✓ In each round, participants were endowed 50 yen and asked to decide whether to
 - ✓ give it to one of the other 7 participants (The endowment they gave were doubled and given to their recipients: recipients received 100 yen) OR
 - ✓ keep it for themselves (They received 50 yen).
- ✓ After each round, they received the feedback information (=how much they received in the round).
- ✓ When they made their decisions, they could see
 - 1st-order information
 - 2nd-order information about all members.

What kind of altruism can be adaptive?

Displayed information

This person...

Gave in the last
round



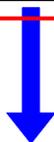
**to the person
who had not
given two
rounds before**

What kind of altruism can be adaptive?

Displayed information

This person...

Gave in the last
round



**to the person
who had not
given two
rounds before**

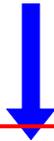
1) 1st-order information:
whether each target **gave** or
didn't give in the last round

What kind of altruism can be adaptive?

Displayed information

This person...

Gave in the last round



to the person who had not given two rounds before

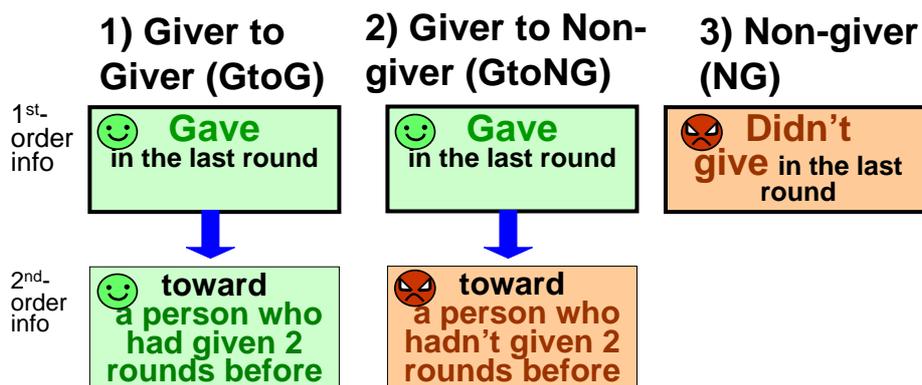
1) 1st-order information: whether each target **gave** or **didn't give** in the last round

2) 2nd-order information: (If 1st-order information was "gave",) whether the person had given to a **giver** or a **non-giver** in the last round

What kind of altruism can be adaptive?

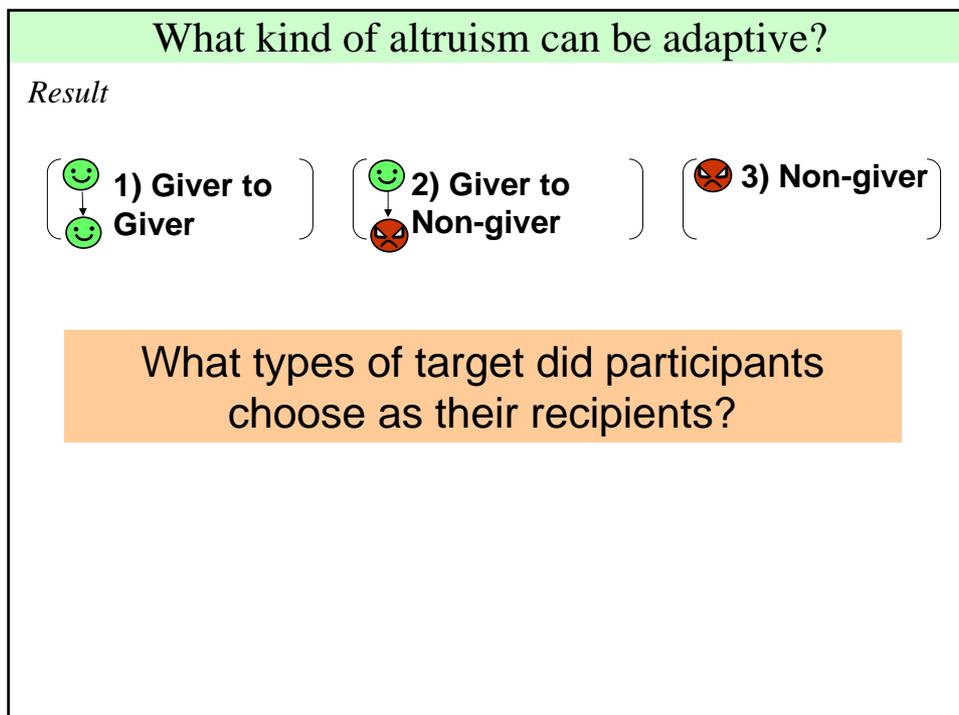
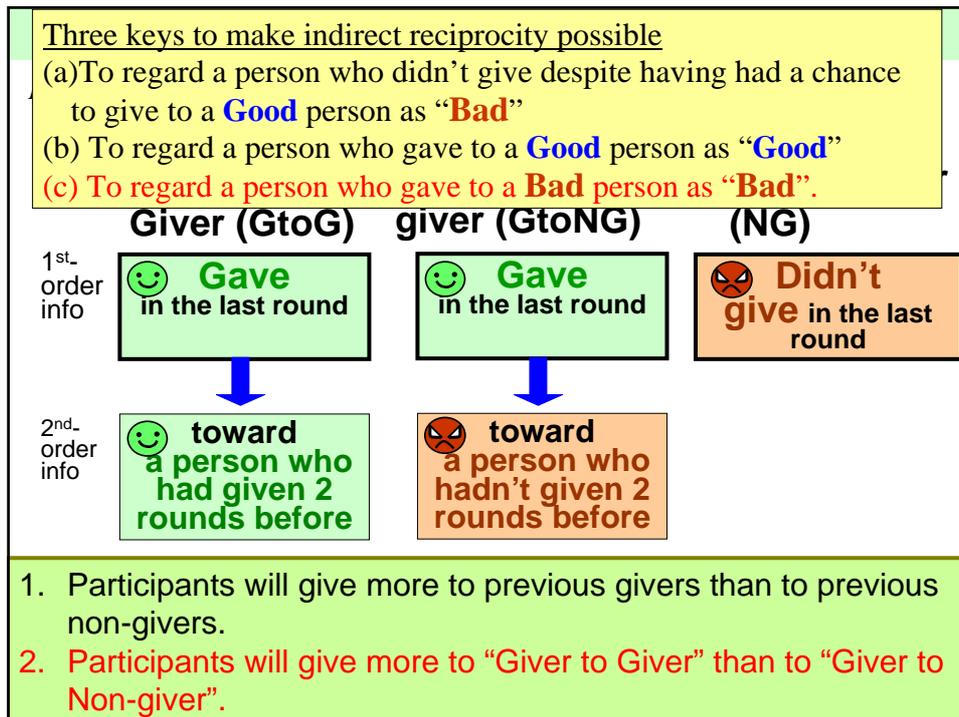
Dependent variable (DV)

There were always three types of targets.



DV: The preference score for each type of target

This DV indicates to what extent each participant gave to (or avoided) the target of each type selectively rather than randomly.



What kind of altruism can be adaptive?

Result

Average of preference of each type (SD)

 1) Giver to  Giver	 2) Giver to  Non-giver	 3) Non-giver
0.14 (0.23)	> 0.00 (0.18)	> -0.14 (0.13)

Results of a comparison between 3 types (ANOVA) showed that the effect of type of target is statistically significant ($F(2,62)=13.23, p<.0001$).

Giver to Giver > Giver to Non-giver > Non-giver

Participants gave more to previous givers than to previous non-givers.
Participants gave more to “Giver to Giver” than to “Giver to Non-giver”.



Two predictions were supported.

What kind of altruism can be adaptive?

Conclusion

Actual behavioral patterns are consistent with what theoretical studies have argued to be adaptive.

To **exclude** not only free riders but also **unconditional cooperators who benefit free riders** is needed.

Last remarks

When we study human rationality, we need to consider its adaptive bases.

Adaptationist approach argues that many of our psychological traits have been shaped by adaptive tasks we (or our ancestors) faced.

The mechanism that has shaped our psychological traits is still unspecified. It can be natural selection, sexual selection, individual learning, cultural evolution (i.e., imitation of successful strategies), and so on.

Although CEFOM/21 will end this month, we will continue our effort to be at the cutting-edge of research on the fundamental sociality of the human mind.

Thank you for your attention

For details, please read the following articles.

Kiyonari, T., Tanida, S. and Yamagishi, T. (2000). Social exchange and reciprocity: Confusion or a heuristic. *Evolution and Human Behavior*, 21(6), 411-427.

Yamagishi, T., Terai, S., Kiyonari, T., Mifune, N. and Kanazawa, S. (in press). The social exchange heuristic: Managing errors in social exchange. *Rationality and Society*.

Takahashi, N. and Mashima, R. (2006). The importance of subjectivity in perceptual errors on the emergence of indirect reciprocity. *Journal of Theoretical Biology*, 243:418-436.