



Evolution of conditional and unconditional commitment

Tadeg Quillien

Center for Evolutionary Psychology, Department of Psychological and Brain Sciences, University of California, Santa Barbara, CA 93106-9660, USA

ARTICLE INFO

Article history:

Received 11 October 2019
Revised 12 February 2020
Accepted 14 February 2020
Available online 19 February 2020

Keywords:

Altruism
Cooperation
Signaling
Evolutionary game theory
Agent-based model

ABSTRACT

We show that altruism can evolve as a signaling device designed to solve commitment problems in interactions with outside options. In a simple evolutionary game-theoretic model, uncertainty about agents' incentives to stay in a relationship can cause the relationship to collapse, because of a vicious circle where being skeptical about one's partner's commitment makes one even more likely to leave the relationship. When agents have the possibility to send costly gifts to each other, analytical modeling and agent-based simulations show that gift-giving can evolve as a credible signal of commitment, which decreases the likelihood of relationship dissolution. Interestingly, different conventions can determine the meaning of the signal conveyed by the gift. Exactly two kinds of conventions are evolutionarily stable: according to the first convention, an agent who sends a gift signals that he intends to stay in the relationship if and only if he also receives a gift; according to the second convention, a gift signals unconditional commitment.

© 2020 Elsevier Ltd. All rights reserved.

1. Introduction

Altruistic behavior, whereby an organism delivers benefits to another organism at a cost to itself, can only evolve under certain conditions. Most mechanisms that have been proposed to account for the evolution of altruism leverage the fact that altruistic behavior directly generates benefits. Some of these mechanisms, such as reciprocity (Trivers, 1971; Axelrod and Hamilton, 1981; Nowak and Sigmund, 1998), kin selection (Hamilton, 1964), and partner choice (Noë and Hammerstein, 1994; McNamara et al., 2008) rely on positive assortment between altruists, whereby altruists are more likely to deliver benefits to other altruists than what would be expected by chance (Fletcher and Doebeli, 2008; Eshel and Cavalli-Sforza, 1982). Other theories emphasize the fact that altruism can evolve when the helper has a stake in the welfare of the recipient (Tooby and Cosmides, 1996; Aktipis et al., 2018; Dyble et al., 2018).

However, it is also possible to devise models where the costs, rather than the benefits, of altruism are responsible for its stability (Gintis et al., 2001). These models rely on costly signaling: the principle according to which individuals can use costly signals to honestly advertise their traits or intentions (Zahavi, 1975; Spence, 1973; Grafen, 1990).

Here we argue that a fundamental feature of social relations, the existence of outside options, constitutes an important selection pressure for the evolution of altruistic behavior as a form of costly

signaling. Individuals in an interaction often face a choice between continuing the interaction and doing something else, for instance finding a new partner. The expected payoff of the activity that an individual can pursue by forgoing the interaction is the individual's outside option. Researchers inspired by biological markets theory (Noë and Hammerstein, 1994) have argued that outside options can be instrumental to the evolution of moral standards, because they force individuals to treat their partners well enough that they will not want to leave the interaction (Debove et al., 2017; Geoffroy et al., 2019; McNamara et al., 2008; Takesue, 2017). Less attention has been paid, however, to another interesting consequence of outside options. In many social interactions, individuals are subject to the threat of unilateral desertion by their partner, often with substantial costs. A pregnant female abandoned by her mate has to pay the costs of childcare alone; fighters whose allies desert the battlefield lose the battle; a tenant may become suddenly homeless if the landlord decides to rent to someone else (Frank, 1988). In situations where the payoff structure features large rewards for successful interaction, but very tempting outside options, the risk of desertion may prevent individuals from engaging in profitable interactions, if they have no means of credibly convincing each other that they will not desert.

In some such cases, cultural institutions like marriages and leases solve the commitment problem: agents win each other's trust by voluntarily restricting their freedom, through the use of a legal document (Frank, 1988; Schelling, 1980; see also Martinez-Vaquero et al. (2015) and Pereira et al. (2017)). In the absence of enforceable contracts, how can individuals solve commitment

E-mail address: tadeguillien@ucsb.edu

problems? Because the interests of partners in a relationship are not necessarily identical, a simple verbal assurance from one's partner may not always be convincing (although see Robson, 1990; Santos et al., 2011; Skyrms, 2002). Signaling theory suggests that individuals can overcome this difficulty by sending a costly signal. Costs can make a signal reliable when the potential benefit an agent derives from sending the signal, minus the cost she incurs, is higher for honest signalers than for liars (Higham, 2013). Remarkably, there is a very natural reason why costly signals of commitment should be more beneficial to honest signalers: an individual who intends to walk away from an interaction will not be around to benefit from the effects of his signal on the behavior of his partner. By contrast, individuals who intend to stay in a relationship have a strong interest in convincing their partners to do the same. Therefore, costly signaling constitutes a very natural solution to commitment problems. When costly, acts of generosity can fulfill this signaling function.

Here we analyze a simple evolutionary game-theoretic model in order to formalize this argument. In our model, agents in a dyad can decide to work together or walk away from the relationship, but each agent is uncertain about how much the other agent would benefit from staying in the relationship. First, we show that such uncertainty can easily cause the collapse of an interaction, because of a recursive feedback dynamic, whereby the uncertainty of each agent makes the other agent ever more uncertain about the viability of the relationship. This suggests that commitment problems may arise easily in real-world interactions. Then we show that generosity can solve this problem: in an extended model where agents are allowed to send costly gifts to each other, gift-giving often evolves as a costly signal of commitment. We also explore which signaling conventions are likely to evolve, by considering a large strategy space where a very large number of signaling strategy profiles are possible.

Finally, we compare the evolutionary dynamics of costly and costless signaling in the model. Existing work in signaling theory has found that even 'cheap talk', i.e. costless signals, can foster coordination between individuals (Robson, 1990; Santos et al., 2011; Skyrms, 2002). In the present model, we find that cheap talk can indeed foster coordination, in the following way. Under certain conditions, the baseline (i.e. no-communication) version of the game we study has two stable equilibria: one where individuals always walk away from the interaction, and another, more efficient equilibrium, where individuals sometimes stay in the interaction. When costless signals are available, populations can escape the inefficient state, and reach the more efficient equilibrium which allows fruitful interactions. However, in situations where outside options are very tempting, and there are no equilibria of the game where agents are willing to stay in an interaction, costless signals cannot create such equilibria. By contrast, costly signals can create stable equilibria where agents successfully coordinate with each other, even when these equilibria do not exist in the no-communication version of the game.

2. Baseline model

Our model is an extension of the stag hunt, a game commonly used to study interactions with outside options (Skyrms, 2004). Each of two players in a dyad can either decide to Stay in the relationship (for instance, to hunt a stag together), or to Walk away from the joint interaction (for instance, forgoing the pursuit of the stag to catch a hare that passes by). If both players Stay, they each get a payoff S ; however if a player Stays while his partner Walks, he gets the null payoff (e.g., because hunting a stag on one's own is too difficult). A player who Walks gets a payoff of w regardless of the decision of his partner. In other words, pursuing a stag is potentially profitable, but also risky, because a

Table 1

Payoff structure of the baseline game. S_i and S_j are drawn from independent uniform distributions ranging from 0 to 1. The values of S_i and S_j are only known to player i and j , respectively.

	Stay	Walk
Stay	S_i, S_j	$0, w$
Walk	$w, 0$	w, w

hunter who decides to pursue a stag comes back empty-handed if his partner chooses to go after a hare.

In a standard stag hunt, the payoff for successful coordination (the value of the stag) is common knowledge between the players. However, an important feature of real-life interactions is that one never fully knows how much one's partner values the interaction. Here we incorporate the realistic assumption that players are uncertain about each others' incentives. The value of S is not necessarily the same for player i and player j : S_i and S_j are drawn from independent uniform distributions ranging from 0 to 1. Furthermore, they are private information: player i knows the value of S_i , but not that of S_j , and vice-versa. The value of w is similar for both players and is common knowledge (see Table 1).

Successful coordination in this game is difficult to achieve. Since your partner might Walk, it is a sensible strategy for you to Walk, especially if the payoff for successful joint interaction is barely above your outside option. The fact that your partner can anticipate that you will think this way makes him even more likely to Walk - this in turn makes the option to Walk all the more tempting for you, and so forth. It is as if players had to read each others' minds in an infinitely recursive manner in order to determine their move. In this process, players are increasingly skeptical of each other's commitment.

Thinking of players as trying to 'read each other's minds' makes it easier to think about the problem they are faced with. But the recursive dynamic we describe does not require that players possess sophisticated cognitive abilities such as a Theory of Mind. It also applies to an evolving population of agents who use simple, genetically encoded decision rules. Consider a very large, well-mixed population from which individuals are randomly selected in pairs, play the game, and then reproduce as a function of their payoff. We look for evolutionarily stable strategies (ESS): strategies which, if played by most individuals in the population, cannot be invaded by another, initially rare strategy (Maynard Smith, 1982).

If there exists an ESS where Staying occurs with positive probability, then this ESS must be a threshold strategy of the form "Stay if $S_i > \beta$, Walk otherwise", where $\beta \in [0, 1]$. To see why, let $p > 0$ be the probability that the other player will Stay. Regardless of whether the other player plays a threshold strategy, the payoff-maximizing strategy is to Stay if $pS_i > w$, and Walk away if $pS_i < w$. This is the threshold strategy just described above.¹ (Note that agents with high values of β are less likely to Stay. For instance, at the extremes, an agent with $\beta = 1$ always Walks, while an agent with $\beta = 0$ always Stays.)

In a population where the average value of β is $\bar{\beta}$, the optimal strategy for an agent i is to stay if:

$$(1 - \bar{\beta})S_i > w$$

because the payoff of Staying is equal to S_i discounted by the probability that the other agent is staying. S is drawn from a uniform distribution, so there is probability $1 - \bar{\beta}$ that the other agent will Stay.

¹ We thank an anonymous reviewer for their input on how to make this step of the argument explicit.

This is equivalent to:

$$S_i > \frac{w}{1 - \beta}$$

β is precisely the threshold for S_i above which an agent Stays, so this equation can be read as saying that the best reply to β is given by the recursion function:

$$R(\beta) = \frac{w}{1 - \beta}$$

At equilibrium, β must be a best reply to itself:

$$\beta^* = \frac{w}{1 - \beta^*}$$

Solving the quadratic equation yields:

$$\beta^* = .5 \pm \frac{\sqrt{1 - 4w}}{2} \tag{1}$$

This equation describes two ESS (depending on whether \pm takes the value + or -). Because the strategy space of the game is continuous, strategies that are ESS will not necessarily be maintained by natural selection. Random fluctuations brought about by mutation will cause the average value of β in the population playing the ESS to actually be $\beta^* + \epsilon$, where ϵ is very small. In such a population, do selection pressures pull the population toward a state where $|\beta^* - \beta| < |\epsilon|$, i.e. toward a state closer to the ESS? If this is the case, then the population will converge back to the ESS even after a small deviation from equilibrium. Such an ESS is said to be a Continuously Stable Strategy (CSS, (Eshel, 1983; McNamara et al., 2003; Bulmer and Parker, 2002)). On the other hand, if selection pushes the population toward a state where $|\beta^* - \beta| \geq |\epsilon|$, then the population will be taken ever further away from the ESS after a small deviation from equilibrium, and such an ESS is not continuously stable.

Of the two ESS described by Eq. (1), only the equilibrium defined by the lower root (i.e., $\beta^- = .5 - \frac{\sqrt{1-4w}}{2}$) is CSS (see SI for proof). The equilibrium only exists if $1 - 4w > 0$, i.e. $w < 1/4$. For $w > 1/4$, the best reply to β is always greater than β , and the population settles to $\beta = 1$, where no player ever stays. That is, when outside options are too large, a runaway escalation of distrust makes successful coordination impossible. It is also easy to verify that ‘Always Walk’ (i.e. $\beta = 1$) is always CSS for $w > 0$: in a population where no player ever Stays, it is always suboptimal to Stay. Finally, one can show that in the interval $[0, 1/4]$, higher values of w are associated with a lower expected payoff for the players: paradoxically, better outside options make players worse off (see SI for proof).²

Whereas in a standard Stag Hunt, Staying is an equilibrium as long as $S > w$, in the present game Staying ceases to be a possible equilibrium whenever the temptation to Walk exceeds one-fourth of the highest possible payoff for successful coordination. In other words, taking into account the uncertainty that exists among partners about their respective incentives underscores the difficulty that outside options create for social interactions.

3. Signaling model

Can communication between the players make successful coordination more likely? We study an extended model, in which the game described above is preceded by a gift-giving phase. In this phase, each player has the possibility to pay a cost c (with

$c > 0$) to send a gift to the other player, yielding a benefit b to that player. Players must make this decision (to send or not send a gift) at the same time. Note that the term ‘gift’ is used here as a generic term meant to denote any behavior that can yield benefits to someone else: it could represent sharing one’s food with someone, holding the door to them, protecting them from a predator, etc. That is, the present argument could in principle apply to a wide range of altruistic behavior.

We assume that a player’s strategy in this extended game is a triplet (α, β, γ) , where each variable can take any value between 0 and 1. α defines the value of S_i above which the player sends a gift. β defines the value of S_i above which the player Stays in the absence of a gift from the other player; γ represents the same value in the presence of a gift. We look for Continuously Stable Strategies (CSS) and restrict our analysis to pure strategies.

When one intends to Walk away, there are only negative consequences to sending a gift. An agent intending to Walk does not get any benefits from sending a gift, because his payoff does not depend on what his partner does. By sending a gift, such an agent would only pay c without receiving any compensating benefits. This ensures that natural selection is stacked against deceptive signalers. Only agents who intend to Stay can possibly get benefits from sending a gift, which makes the signal conveyed by the gift honest. This raises the possibility that natural selection might favor strategies that use gift-giving as a reliable signal of one’s intention to Stay, enabling higher levels of coordination than in the no-communication version of the game.

Our 3-dimensional continuous strategy space in principle allows for an infinity of possible signaling strategies; among all of them, which kinds of strategies can actually be stable? Consider that an agent who has not sent a gift, but has received a gift from his partner, faces a dilemma. The fact that his partner sent a gift indicates that she leans toward Staying. However, her policy might be to Stay only if she has herself received a gift. Therefore it is unclear whether she will Stay or Walk.

One may suspect that a successful signaling strategy is one where a convention evolves to resolve this ambiguity. There are two possible such conventions: under the first convention, a gift is a signal of conditional commitment, which means “I will Stay if and only if you sent a gift”; under the second convention, a gift is a signal of unconditional commitment meaning “I will Stay regardless of whether you sent a gift”. When a whole population converges on one of these conventions, gifts are an unambiguous signal. In the following two sections we give an outline of the proofs showing the existence of a CSS for each convention (see SI for full analysis).

3.1. Conditional commitment

Formally, a conditional commitment strategy is one where $\alpha = \gamma$. At equilibrium, the strategy must also have $\beta^* = 1$, for the following reason. Clearly, we have $\beta^* > \gamma^*$ (i.e. agents who receive a signal are more likely to Stay), otherwise signaling would have no benefit. This implies that $\beta^* > \alpha^*$. Jointly, the facts that $\beta^* > \alpha^*$ and $\alpha^* = \gamma^*$ imply that agents who do not signal do not Stay. Therefore one should always Walk away from a non-signaling partner.

3.1.1. Evolutionary stability

In a population playing $\bar{\alpha} = \bar{\gamma}$, the best response strategy must also have $\alpha = \gamma$ (see SI). To compute the equilibrium value α^* (which is identical to γ^*), we compare the expected payoff of an agent who Signals and Stays conditionally (denoted $E(S)$) to that of an agent who does not Signal and does not Stay (denoted $E(-S)$). We have:

$$E(S) = \alpha^*(w - c) + (1 - \alpha^*)(S_i - c)$$

² As noted by an anonymous reviewer, a somewhat similar paradox exists in transportation networks: sometimes adding a new road to a network increases the average time users spend traveling from point A to point B, because the new road creates a less efficient Nash equilibrium (Frank, 1981).

(with probability α^* , the partner does not Signal, so the agent Walks away and gets w ; with probability $1 - \alpha^*$, the partner Signals and Stays, so coordination is successful and the agent gets S_i . In both cases, the agent pays the cost c of Signaling) and:

$$E(-S) = w$$

An agent should Signal and Stay conditionally when:

$$E(S) > E(-S)$$

→

$$\alpha^*(w - c) + (1 - \alpha^*)(S_i - c) > w$$

→

$$S_i > \frac{c + w - \alpha^*w}{1 - \alpha^*}$$

Therefore, α^* must satisfy:

$$\alpha^* = \frac{c + w - \alpha^*w}{1 - \alpha^*}$$

Solving the quadratic equation yields:

$$\alpha^* = \frac{1 + w}{2} \pm \frac{\sqrt{(1 - w)^2 - 4c}}{2}$$

This equation describes two ESS (depending on whether \pm takes the value + or -). Note that these equilibria only exist if $(1 - w)^2 - 4c > 0$, that is if $w < 1 - 2\sqrt{c}$. When $w > 1 - 2\sqrt{c}$, runaway escalation of distrust makes coordination impossible, and selection takes the population to $\bar{\alpha} = 1$. We denote the lower and the upper equilibria as $\alpha^- = \frac{1+w}{2} - \frac{\sqrt{(1-w)^2 - 4c}}{2}$ and $\alpha^+ = \frac{1+w}{2} + \frac{\sqrt{(1-w)^2 - 4c}}{2}$. We now show that only α^- is CSS.

3.1.2. Continuous stability

$R_\alpha(\alpha)$ is the function giving the best reply value of α (which by necessity is also the best reply value of γ , see SI) in a population where $\bar{\alpha} = \bar{\gamma}$.

As shown in Section 3.1.1, the best reply function for α is:

$$R_\alpha(\alpha) = \frac{c + w - \alpha w}{1 - \alpha}$$

Its derivative with respect to α is:

$$R'_\alpha(\alpha) = \frac{c}{(1 - \alpha)^2}$$

For an equilibrium α^* , if $0 < R'_\alpha(\alpha^*) < 1$ then α^* is CSS, and if $R'_\alpha(\alpha^*) > 1$ it is not a CSS (see e.g. Bulmer and Parker, 2002). Note that $R'_\alpha(\alpha^*) < 1 \iff 1 - \alpha^* > \sqrt{c}$. We first show that α^- is CSS. We have:

$$1 - \alpha^- = \frac{1 - w}{2} + \frac{\sqrt{(1 - w)^2 - 4c}}{2}$$

We know that $w < 1 - 2\sqrt{c}$, otherwise the equilibrium does not exist. This is equivalent to:

$$\frac{1 - w}{2} > \sqrt{c}$$

Together with $\frac{\sqrt{(1-w)^2 - 4c}}{2} > 0$, this implies that:

$$1 - \alpha^- > \sqrt{c}$$

$$\iff R'_\alpha(\alpha^-) < 1$$

meaning α^- is stable.

Since $\lim_{\alpha \rightarrow 1} R(\alpha) = \infty$, a simple graphical argument shows that $R'_\alpha(\alpha^+) > 1$, and the upper equilibrium is unstable (Fig. 1).

Fig. 2 shows the values of w and c for which the conditional commitment strategy is stable, together with the corresponding values of α^* .




Fig. 1. Why the α^+ equilibrium is unstable. The equilibria occur when $R(\alpha) = \alpha$, so they must lie at the intersections between the best-response function (red line) and the identity line (dashed line). Since $R'(\alpha^-) < 1$, at $\alpha = \alpha^-$ the red line crosses the dashed line from above, therefore at $\alpha = \alpha^+$ it must be crossing the dashed line from below. This means that $R'_\alpha(\alpha^+) > 1$, so that α^+ is unstable. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

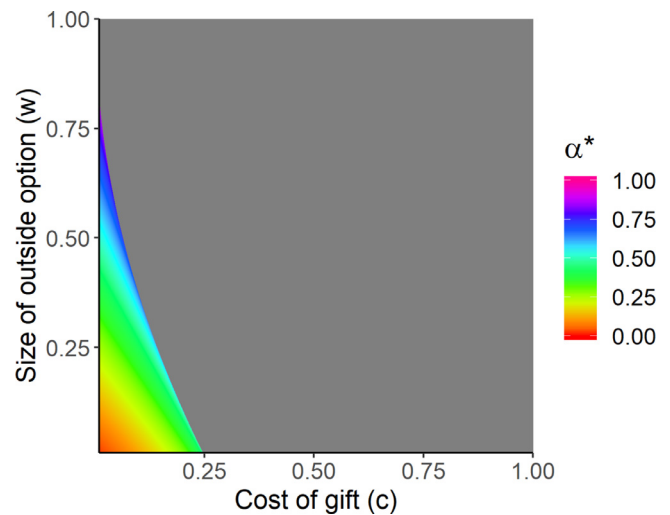


Fig. 2. Conditions for the stability of the Conditional Commitment equilibrium. Conditional commitment is stable for values of c and w in the colored area. At equilibrium, agents send a costly signal whenever they have $S_i > \alpha^*$; when they do so, they Stay if and only if they also receive a signal from their partner.

3.2. Unconditional commitment

Formally, an unconditional commitment strategy is one where $\alpha = \beta$. At equilibrium, the strategy must also have $\gamma^* = w$. The fact that $\alpha^* = \beta^*$ implies that a signaling agent always Stays, regardless of whether he receives a signal. Therefore, an agent who receives a signal is certain that his partner will Stay, and should Stay if and only if $S_i > w$.

3.2.1. Evolutionary stability

In a population playing $\bar{\alpha} = \bar{\beta}$, the best response strategy must also have $\alpha = \beta$ (see SI). To compute the equilibrium value α^* (which is identical to β^*), we compare the expected payoff of an agent who Signals and Stays unconditionally (denoted $E(S)$) to that of an agent who does not Signal and Stays only in the presence of a signal (denoted $E(-S)$). We have:

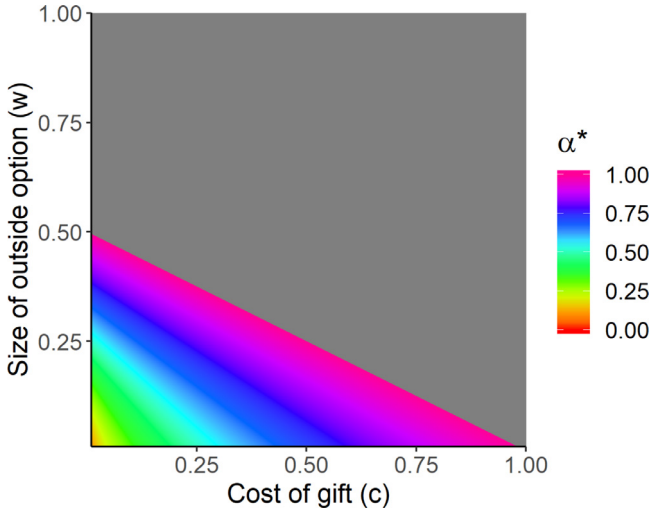


Fig. 3. Conditions for the stability of the Unconditional Commitment equilibrium. Unconditional commitment is stable for values of c and w in the colored area. At equilibrium, agents send a costly signal and Stay whenever they have $S_i > \alpha^*$.

$$E(S) = (1 - w)S_i - c$$

$$E(-S) = (1 - \alpha)S_i + \alpha w$$

(In the first case, the partner Stays with probability $1 - w$, and you pay the cost c of Signaling. In the second case, the partner Stays with probability $1 - \beta (= 1 - \alpha)$; and Walks with probability α ; it is possible to detect that she will Walk from the absence of a signal, and pocket the outside option in this case).

An agent should Signal and Stay unconditionally when:

$$E(S) > E(-S)$$

→

$$(1 - w)S_i - c > (1 - \alpha)S_i + \alpha w$$

→

$$S_i > \frac{c + \alpha w}{\alpha - w}$$

Therefore α^* must satisfy:

$$\alpha^* = \frac{c + \alpha^* w}{\alpha^* - w}$$

Solving the quadratic equation yields:

$$\alpha^* = w + \sqrt{w^2 + c}$$

Therefore, the strategy is a signaling ESS whenever $w + \sqrt{w^2 + c} < 1$, i.e., whenever $w < \frac{1-c}{2}$. Fig. 3 shows the values of w and c for which the unconditional commitment strategy is ESS, together with the corresponding values of α^* . We now show that the strategy is also a CSS.

3.2.2. Continuous stability

The method employed Section 3.1.2 to determine continuous stability does not work as well here because we have $R'(\alpha^*) = -1$, which is not enough information to determine whether the equilibrium is stable or unstable. This method consisted in looking at the best reply to a given strategy – instead, in the following, for a given strategy we look at strategies directly in this strategy’s neighborhood and ask whether they have higher fitness (see McNamara et al., 2003).

Let $E(\alpha, \bar{\alpha})$ be the expected payoff of an agent playing α in a population playing $\bar{\alpha}$. Then, let $f(\bar{\alpha})$ be the partial derivative of

$E(\alpha, \bar{\alpha})$ with respect to α , evaluated at $\alpha = \bar{\alpha}$. Intuitively, $f(\bar{\alpha})$ gives the direction of selection on α in a population playing $\bar{\alpha}$: e.g. when $f(\bar{\alpha})$ is negative, an agent with a value of α slightly below the population average enjoys a higher expected payoff than an agent playing the resident strategy. We go yet one level deeper and consider the partial derivative of this partial derivative. An equilibrium α^* is continuously stable if:

$$\frac{\partial f}{\partial \bar{\alpha}} \Big|_{\bar{\alpha}=\alpha^*} < 0$$

In other words, if, at equilibrium, a change in the population value $\bar{\alpha}$ in a given direction causes a change in the selection pressure on α in the opposite direction, the equilibrium is stable (see McNamara et al., 2003). The intuition is that such negative feedback keeps pushing the population back to equilibrium.

$f(\bar{\alpha})$ is simply $E(-S) - E(S)$, i.e. the difference in expected payoff between an agent who does not signal and Stays conditionally, and an agent who signals and Stays unconditionally (see SI for proof). From 3.2.1 we know that $E(-S) = (1 - \bar{\alpha})S_i + \bar{\alpha}w$ and $E(S) = (1 - w)S_i - c$. Then:

$$f = (1 - \bar{\alpha})S_i + \bar{\alpha}w - (1 - w)S_i + c$$

$$= (w - \bar{\alpha})\bar{\alpha} + \bar{\alpha}w + c$$

$$= -\bar{\alpha}^2 + 2w\bar{\alpha} + c$$

Therefore we have:

$$\frac{\partial f}{\partial \bar{\alpha}} \Big|_{\bar{\alpha}=\alpha^*} = -2\alpha^* + 2w$$

$$\frac{\partial f}{\partial \bar{\alpha}} \Big|_{\bar{\alpha}=\alpha^*} < 0$$

and the ESS is continuously stable.

3.3. Other unstable signaling equilibria

We also find that no other signaling ESS is a CSS. It is easy to see how a signaling ESS must have $\beta^* > \gamma^*$ (i.e., agents are more likely to Stay when they receive a signal), otherwise Signaling has no benefits. Note also that α^* cannot be smaller than both β^* and γ^* : an agent playing this strategy would sometimes Signal when he strictly does not intend to Stay, which is suboptimal.

There are four different types of strategies that obey these constraints: strategies with $\alpha > \beta > \gamma$; $\beta > \alpha > \gamma$; $\alpha = \gamma$; or $\alpha = \beta$. Strategies with $\alpha = \gamma$ and $\alpha = \beta$ correspond to the ‘Conditional commitment’ and ‘Unconditional commitment’ strategies analyzed in the previous sections. In the SI we find that there exist ESS with $\alpha > \beta > \gamma$ and with $\beta > \alpha > \gamma$, but that they are all continuously unstable. In other words, the only signaling strategies that are CSS are those where signals are unambiguous (in the sense that if one knows the signal an agent has sent and the signal it has received, one can know with certainty whether the agent will Stay).

3.4. Can costly signaling strategies invade non-signaling, ‘Always Walk’ strategies?

Consider a population of silent agents who Always Walk, i.e. a population with $\bar{\alpha} = \bar{\beta} = 1$. In this population, there is no selection pressure on γ , since no signals are sent. Therefore, random drift can take $\bar{\gamma}$ to a low value, such that if an agent signals, its partner is very likely to Stay. This might create incentives for signaling.

Indeed, unconditional commitment strategies can invade a population with $\bar{\alpha} = \bar{\beta} = 1$. They can do so under the same conditions that make the unconditional commitment strategy a CSS. Intuitively, this is suggested by the fact that for a given combination

of c and w , there is only one single CSS with $\alpha^* = \beta^*$ (as shown in Section 3.2.2). When the conditions are such that this CSS has $\alpha^* = \beta^*$ below 1, a population with $\bar{\alpha} = \bar{\beta} = 1$ will be invaded by the CSS, provided that drift has taken γ to a low enough value. (The analysis in Section 3.2.2 only applies to signaling strategies, but we show in the SI that the same logic applies to the case with $\bar{\alpha} = \bar{\beta} = 1$).

It is easy to see that, in a population with $\bar{\alpha} = \bar{\beta} = 1$, conditional commitment strategies cannot invade. When initially rare, a mutant who signals, and then stays only if it receives a signal, will always signal vacuously (the fact that the resident is silent means that the mutant always Walks, thus wasting its costly signal). However, conditional commitment strategies have a large basin of attraction. The upper boundary of this basin of attraction is $\alpha^+ = \frac{1+w}{2} + \frac{\sqrt{(1-w)^2 - 4c}}{2}$, i.e. the continuously unstable ESS as described in Section 3.1.1 (To give an example, for $w = .3$ and $c = .05$, the upper boundary value of the basin of attraction is $\alpha^+ = .92$; see also Fig. S2 in SI). This suggests that, in an analysis of the evolutionary dynamics of a small population, the conditional commitment signaling strategy may be able to invade the silent population, if stochastic events take the value of $\bar{\alpha}$ and $\bar{\gamma}$ below the upper boundary of the basin of attraction. Indeed, we find that in our agent-based simulations, conditional commitment strategies can be evolutionarily successful even in simulations where the initial population plays a non-signaling, 'Always Walk' strategy (see Section 4).

3.5. Comparison of conditional and unconditional commitment strategies

Conditional and unconditional commitment strategies are stable under slightly different sets of conditions. A strategy of conditional commitment can be stable even when outside options are quite high, but not when gifts are too costly (formally, when $w < 1 - 2\sqrt{c}$, see Fig. 2). Conversely, a strategy of unconditional commitment can be stable even when gifts are quite costly, but not when outside options are too high (formally, when $w < \frac{1-c}{2}$, see Fig. 3). Notably, however, both strategies allow successful coordination to occur even in some regions of parameter space where $w > 1/4$: they enable agents to coordinate in a wider range of conditions than in the baseline version of the game.

4. Agent-based simulations

We conducted agent-based simulations in order to further explore our evolutionary model. The simulations had the following goals.

First, while the analytical model studied above focused on costly gift-giving ($c > 0$), we wanted to explore the potential effects of costless gifts ($c = 0$) on evolutionary dynamics. Notably, existing research has shown that costless signals can foster coordination in games like the Stag Hunt (Robson, 1990; Santos et al., 2011; Skyrms, 2002); we wanted to find to which extent this applied to the present model.

Second, we wanted to check the robustness of the main findings from the analytical model. While the analytical model relies on equilibrium analysis and assumes a very large population, agent-based simulations actually simulate the process of evolution over time in a finite population of agents, where chance plays a larger role. Would costly gift-giving evolve in the simulations, and would natural selection design both conditional commitment and unconditional commitment strategies?

We also wanted to rule out the following possibility: maybe evolution designs agents who send costly signals because they are signals, but not because they are costly. In other words, we

needed to rule out the possibility that costly gift-giving evolves simply because agents lack a costless signaling alternative. Therefore, we set out to find whether evolution would design costly gift-giving strategies even when costless gifts were also available. This scenario is complicated to study analytically, but tractable with simulations.

4.1. What can cheap talk achieve?

In order to see how much coordination can be fostered by costless and costly signals, we ran a first set of simulations. Across simulations, we independently manipulated the following three factors (see Appendix for details of the simulations).

First, we varied the game that was being played. In a 'baseline' condition, agents played the baseline game (described above in Section 2), and had no opportunities for communication. In the other simulations, agents played one of two modified versions of the signaling game (described above in Section 3). In a 'cheap talk' condition, agents played a variant of the game where an agent could send up to two different costless gifts. For each of the two possible gifts, the agent had an evolvable threshold rule such that it would send the gift if it drew a value of S_i above its threshold. For each type of gift, agents also had a threshold decision rule determining the value of S_i above which they would Stay, given that they had received that gift (agents who received both types of gifts Stayed if they had S_i larger than at least one of their thresholds). In a 'costly signal' condition, agents played the same game as in the 'cheap talk' condition, except that one of the gifts was costly, with $c = .05$.

Second, we varied the strategy played by the initial population. In half of the simulations, the initial population had strategies with threshold values set to 1, such that agents in the first generation effectively played the strategy "Never send any signal and Always Walk regardless of the signal you receive" (in the baseline version, this was simply "Always Walk"). In the other simulations, the initial population had threshold values set to .5, such that agents Stayed if they had $S_i > .5$ and (in the 'cheap talk' and 'costly signal' conditions) sent both types of gift if they had $S_i > .5$.

Finally, we varied the size of the outside option (w) across simulations.

We find that cheap talk fosters coordination compared to the baseline model, but does so mostly when the initial population plays 'Always Walk' (see Figs. 4 and 5). In the baseline model, 'Always Walk' is theoretically always a stable equilibrium. Populations that played the baseline game and started at 'Always Walk' sometimes reached the more efficient cooperative equilibrium, but they remained stuck at 'Always Walk' when outside options were too large ($w \geq .15$). By contrast, populations that had access to costless signals were able to reach the efficient equilibrium at a much greater frequency. Costless signals allow mutants to invade the 'Always Walk' strategy by using a 'secret handshake' (see Robson, 1990): if they draw a high enough value of S_i , they send a signal, then Stay only if they receive a signal. This allows them to enjoy successful interactions when meeting other mutants, but safely pocket the outside option when meeting an agent playing the resident strategy.

While the costless and costly signaling models support the same amount of successful coordination for small values of w , for $w > \frac{1}{4}$ the amount of coordination supported by costless signaling starts declining at a much faster rate than for costly signaling. We can understand why by looking at the evolutionary dynamics of a given simulation run with $w > \frac{1}{4}$ (see figures S5 and S6): costless signaling allows 'secret handshake' invasions to occur, but these invasions do not lead to a stable equilibrium, as 'Always Walk' systematically re-invades. By contrast, costly signaling strategies are more stable. This is consistent with the result from the analytical

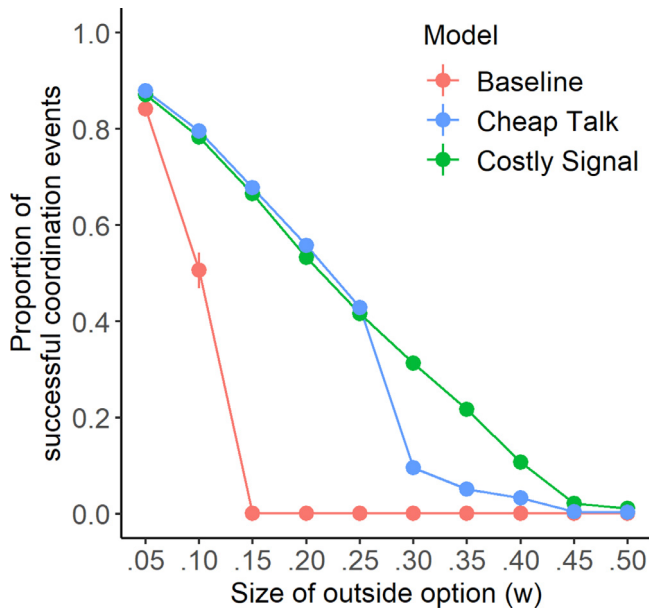


Fig. 4. Proportion of successful interactions as a function of w , for three different models, when the first generation plays 'Always Walk'. The y-axis represents the average proportion of encounters where both agents Stayed in the interaction, for the last 1000 generations of a given simulation. Each data point represents an average of 100 simulations. Error bars (most of them too small to be visible) represent the standard error of the mean (where variance is computed across simulations). Agents in the 'Cheap Talk' condition could send up to two different costless gifts. The 'Costly Signal' condition was similar, except that one of the two possible gifts had cost $c = .05$.




Fig. 5. Proportion of successful interactions as a function of w , for three different models, when the first generation does not play 'Always Walk'. The y-axis represents the average proportion of encounters where both agents Stayed in the interaction, for the last 1000 generations of a given simulation. Each data point represents an average of 100 simulations. Error bars (most of them too small to be visible) represent the standard error of the mean (where variance is computed across simulations). Agents in the 'Cheap Talk' condition could send up to two different costless gifts. The 'Costly Signal' condition was similar, except that one of the two possible gifts had cost $c = .05$.

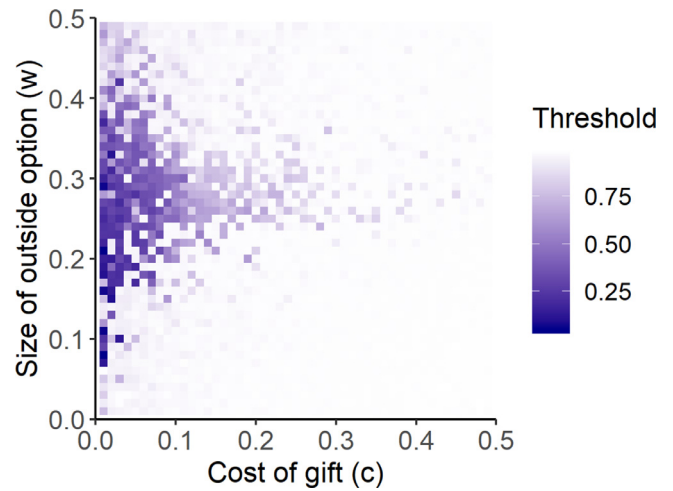


Fig. 6. Costly gift-giving across simulations. The darkness of a tile is proportional to the prevalence of gift-giving in the population in the last 1000 generations of one simulation. Formally, it represents the average value of α_1 , the value of S above which an agent sends a costly gift – lower values of α_1 represent more frequent gift-giving.

model that 'Always Walk' is the only possible equilibrium in the baseline version of the game for $w > \frac{1}{4}$.

Why are costless signaling strategies not stable beyond $w > \frac{1}{4}$? Intuitively, costless signals foster coordination by relying on *transient information* (Skyrms, 2002): if at a given point in the evolutionary process, the signals sent by agents in a population contain information about their intention to Stay, this makes it possible for the population to escape a state where agents always Walk. However, if the signal is costless, it will lose its information content over time, because there is no selection pressure against agents who send a signal when in fact they intend to Walk. As a result, drift can take populations to a state where agents always have the same signaling behavior regardless of whether they have incentives to Stay. Therefore, in conditions where Staying is irrational when agents do not have reliable information about their partners' incentives (their value of S), costless signals do not create new cooperative equilibria.

4.2. When does costly gift-giving evolve, and which form does it take?

In a second set of simulations, we systematically explored the level of costly gift-giving that evolves across different values of the outside option (w) and the magnitude of the costly gift ($c > 0$). In these simulations, agents played the 'costly signal' version of the game, as describe just above (in Section 4.1). Every simulation was initialized with a population of agents who never send gifts and always Walk regardless of the signal they receive.

We find that costly gift-giving reliably evolves, in a portion of parameter space where gift costs are small enough, and outside options are neither too high nor too low (see Fig. 6). Intuitively, outside options that are too tempting inhibit successful interactions even when a signaling system is available, and outside options that are too low do not distract agents from successful coordination, making costly gift-giving unnecessary. We also find that two qualitatively different conventions govern gift-giving. Consistent with the analytical model, when gift costs are low, populations tend to evolve conditional-commitment strategies, where agents who send a gift Stay if and only if they also receive a gift (Fig. 7a). When gift costs are high, populations tend to evolve unconditional-commitment strategies, where sending a gift

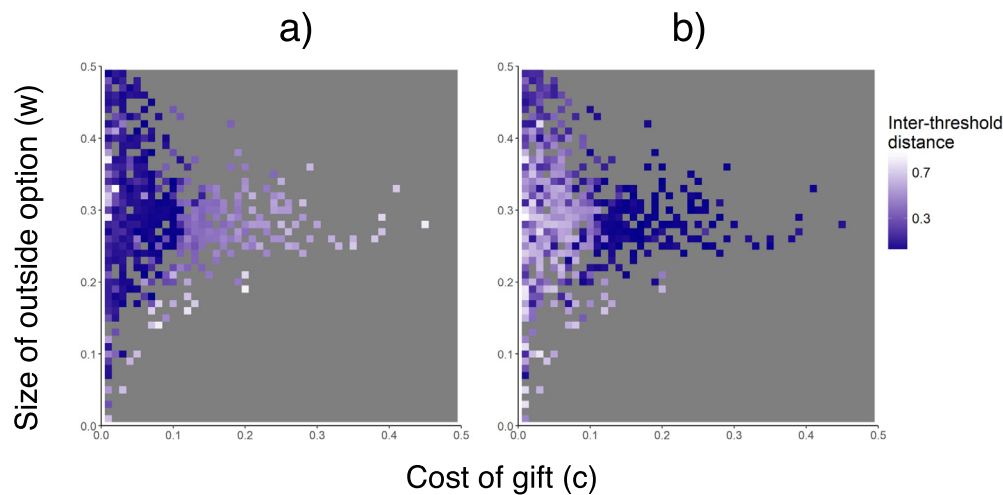


Fig. 7. Prevalence of conditional (a) and unconditional (b) commitment strategies. Colored tiles represent simulations where costly gift-giving evolved ($\alpha_1 < .9$). On panel **a**), dark tiles represent simulations where conditional commitment strategies have evolved; the brightness of a tile represents the absolute difference between the threshold value of S above which agents send a costly gift (α_1) and the threshold value of S above which agents Stay if they have received a costly gift (γ_1) – dark values therefore correspond to $\alpha_1 \approx \gamma_1$. On panel **b**), dark tiles represent simulations where unconditional commitment strategies prevailed; the brightness of a tile represents the absolute difference between the threshold value of S above which agents send a costly gift (α_1) and the threshold value of S above which agents Stay if they have received a costless gift (γ_0) – dark tiles therefore correspond to $\alpha_1 \approx \gamma_0$.

conveys one's intention to Stay in the relationship no matter what (Fig. 7b).

Costless gift-giving also reliably evolved. However, in simulations where $w > \frac{1}{4}$, agents very rarely Stayed when they received only a costless gift (see Figs. S7 and S8).

The present model makes the assumption that the payoff for successful coordination, S_i , is drawn from a uniform distribution. Because the choice of a probability distribution can sometimes qualitatively alter the dynamics of a model (Jagau and van Veelen, 2017), we also ran a replication of our second set of simulations, where S_i was drawn from a normal distribution, truncated at 0 and 1, with mean .5 and standard deviation .25. Compared to a uniform distribution, a normal distribution of S_i makes agents slightly less uncertain about their partner's incentives. As a result, we predictably found slightly lower levels of costly gift-giving in these simulations, albeit the qualitative pattern of results was identical (fig S9-10).

5. Discussion

We have shown that generosity can evolve as a costly-signaling mechanism to solve a commitment problem in interactions with outside options. However, in our model, the amount of benefits created by the gift is irrelevant to the evolutionary outcomes (note that the parameter b does not appear in any equation). In principle, many different types of costly actions could serve as a reliable commitment signal; why would one use generosity instead of any other costly behavior? There are several reasons why individuals may preferentially use benefit delivery to signal commitment. First, organisms have limited attention, and are expected to allocate more of it to fitness-relevant stimuli. Second, individuals have an intrinsic interest in their interaction partners staying alive and healthy (Tooby and Cosmides, 1996). Third, selection pressures such as reciprocity and kinship will also select for cognitive designs that deliver benefits to close others. Note that for these reasons, individuals might actually use the receipt of benefits as a cue that the helper values them and intends to cultivate the relationship. Our simulations have shown that signaling equilibria can be reached even from an initial state where agents consider that gifts do not convey commitment. However, it might be more realistic to assume that individuals draw positive inferences from

benefit delivery even before the dynamics of costly signaling are put into motion.

For simplicity, our analysis focused on a two-stage game. A more exhaustive analysis of the commitment problem would benefit from taking into account the fact that in many interactions, temptations to leave the interaction can occur at multiple points over time. Therefore there may be a selection pressure for inter-actants to send commitment signals to each other periodically. Future work could attempt to extend the present model along those lines, for instance by studying an indefinitely repeated game where agents' estimates of their incentives can change over time.

The idea that outside options can create a selection pressure for costly signaling of commitment has been explored in previous modeling work (Bergstrom et al., 2008; Bolle, 2001; Camerer, 1988; Carmichael and MacLeod, 1997; Sozou and Seymour, 2005) and has received support in ethnographic and experimental work in humans (Bliege-Bird et al., 2018; Yamaguchi et al., 2015; Komiya et al., 2019); our results contribute two novel insights.

First, in existing models, outside options typically threaten social interactions because they allow 'cheaters' to find new victims right after having exploited their current partner (Bergstrom et al., 2008; Bolle, 2001; Carmichael and MacLeod, 1997). This assumes that the payoff structure of interactions allows agents to 'cheat' on their partner, as in a Prisoner's Dilemma or a Trust Game. We show that outside options can threaten the success of social interactions because of a more general reason: whenever the success of an interaction requires investment on the part of both parties, uncertainty about whether one's partner will invest can make one reluctant to invest, which in turn will make one's partner even less likely to invest, and so on. This problem exists even in the absence of incentives for exploitative behavior. Our results do not require the existence of a dichotomy between cooperators and defectors in the population, but rely on differences in the social incentives of agents (see André, 2010).

Second, the continuous strategy space we explored allowed us to test the stability of a wide range of possible signaling strategies. We then found that two qualitatively different conventions for gift-giving can evolve; these two conventions correspond to the two possible signaling systems that make the intentions of a gift sender unambiguous. Camerer, using a discrete strategy space, also finds the existence of two signaling equilibria (Camerer, 1988). In

his model, unconditional commitment is a possible strategy only when the game is strongly asymmetric, such that it is common knowledge that one of the players is very likely to Stay (this removes the need for the other player to seek more information). By contrast, in the current model, unconditional-commitment can be a stable equilibrium despite the fact that signalers are completely uncertain about their partners' incentives before gifts are sent.

The existence of different possible gift-giving norms comes from the fact that agents face a second-order coordination problem: the use of costly signals allows them to solve the first-order coordination problem only if they agree about the meaning of the signal. This creates a selection pressure for agents to follow the convention that governs gift-giving in their population.

This result may have implications for understanding the variation that has been documented in cooperative behavior across relationship types (Fiske, 1992), across cultures (Henrich et al., 2010; 2001), across time (Pinker, 2012; Baumard et al., 2015), and across social ecologies (Nettle et al., 2011; Safra et al., 2016). Empirical and theoretical efforts to explain variability in cooperative behavior (Skyrms, 2004; Henrich et al., 2010; 2001; Nettle et al., 2011; Safra et al., 2016; Baumard et al., 2015; Lettinga et al., 2019) typically aim at accounting for variability in the level of cooperation, i.e. variation on a one-dimensional continuum ranging from selfishness to pure altruism. By contrast, modelling approaches such as the current one have the potential to explain differences in the qualitative properties of generosity norms.

Across the two kinds of gift-giving equilibria identified in the present model, generosity has a different *meaning*. That is, the inferences that agents can draw about the author of a generous act are not the same under the conditional commitment norm as under the unconditional commitment norm. Correspondingly, generosity is elicited under different conditions, and takes different forms, in the two different families of equilibria (i.e. gifts tend to be more costly and less frequent when they signal unconditional commitment). Of course, we do not suggest that actual social relationships exactly mimic our toy model: there are probably more than two possible norms governing the meaning of a commitment signal in the real world. Rather, our results show that a signaling approach can in principle explain some of the variability in fine-grained properties of cooperative behavior. Many possible conventions can govern the meaning of a generous act, and in different contexts, individuals coordinate on different such conventions, much as people can coordinate on different conventions for the meaning of a word.

Acknowledgement

For helpful suggestions and comments, I would like to thank Leda Cosmides, Dan Conroy-Beam, the members of the Center for Evolutionary Psychology research seminar at UCSB, and two anonymous reviewers.

Appendix A. agent-based simulations

General setup

Each simulation consisted of a population of 300 agents allowed to evolve during 25,000 generations. An agent's genotype is made of 5 genes (α_0 , α_1 , β , γ_0 , γ_1) which can take any value between 0 and 1. α_0 and α_1 define the values of S_i above which the agent sends a costless and a costly gift, respectively (agents are allowed to send both types of gifts at the same time). Genes β , γ_0 and γ_1 determine behavior in the following way. An agent Stays if either:

- It received no gift and has $S_i > \beta$;
- It received a costless gift and has $S_i > \gamma_0$;
- It received a costly gift and has $S_i > \gamma_1$;

-It received both gifts and has either $S_i > \gamma_0$ or $S_i > \gamma_1$; and Walks otherwise.

At the beginning of each generation, agents were paired randomly, then each pair played one game. Then all agents died and reproduced asexually. An agent's expected number of offspring was equal to its total payoff from the game, divided by the average population payoff (payoffs were first standardized such that the agent with the lowest payoff had payoff .01; this was achieved by subtracting, from the payoff of each agent, the payoff of the agent with the lowest payoff, then adding .01 to the payoff of each agent. This procedure ensured that no agent had a negative payoff before the selection phase).

Population size was kept fixed at 300 during the entire simulation. Each offspring inherited its parent genotype, subject to a .02 independent probability of mutation for each gene; mutation added random noise (drawn from a normal distribution with mean 0 and standard deviation .05) to the value of a gene, subject to the constraint that gene values could not leave the interval [0,1].

For each simulation, we recorded the average value of each gene in the population for the last 1000 generations of that simulation (total simulation time was 25,000 generations). Inter-threshold distance scores for Fig. 7 were computed by taking, for each of these 1000 generations, the absolute difference between the average values of the two relevant genes in the population, and averaging across generations. Simulation software is written in JavaScript and was run using Google Chrome (version 79.0.3945 for first set of simulations, version 62.0.3202 for second set) and can be executed using any modern internet browser; code and data from the simulations are available in Supplementary materials.

First set of simulations

In this set of simulations, across simulations agents played either the baseline game, a 'cheap talk' variant where the baseline game was preceded by a gift-giving phase where both gifts were costless, or a 'costly signal' variant where the baseline game was preceded by a gift-giving phase where one gift was costly and the other costless. These variants of the game were implemented in the following way. In the baseline variant, we forced the values of every gene other than β to 1, and prevented mutation from acting on these genes. In the 'cheap talk' variant, we set $c = 0$, such that the 'costly' gift was effectively costless. In the 'costly signal' version, we set $c = 0.05$.

Orthogonal to this manipulation, in half the simulations, the simulation was initialized with $\alpha_0 = \alpha_1 = \beta = \gamma_0 = \gamma_1 = 1$, for each agent in the first generation. In the other half, the simulation was initialized with $\alpha_0 = \alpha_1 = \beta = \gamma_0 = \gamma_1 = 0.5$, except for the baseline game, for which only β had an initial value of .5.

We also varied the value of w from .05 to .50, in increments of .05, across simulations. For every combination of the parameters we manipulated, we conducted 100 simulations, for a total of $3 * 2 * 10 * 100 = 6000$ simulations.

Second set of simulations

For this set of simulations, agents played the equivalent of the 'costly signal' variant described in the previous section, except that the value of c was varied across simulations. In each simulation, the population was initialized with a first generation playing $\alpha_0 = \alpha_1 = \beta = \gamma_0 = \gamma_1 = 1$, for each agent. We conducted 2500 simulations, one for each combination of w and c between (.01, .01) and (.50, .50), and 2500 additional simulations, spanning the same parameter space, for the model where S is normally distributed.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.jtbi.2020.110204.

References

- Aktipis, A., Cronk, L., Alcock, J., Ayers, J.D., Baciú, C., Balliet, D., Boddy, A.M., Curry, O.S., Krems, J.A., Muñoz, A., et al., 2018. Understanding cooperation through fitness interdependence. *Nat. Hum. Behav.* 2 (7), 429.
- André, J.-B., 2010. The evolution of reciprocity: social types or social incentives? *Am. Nat.* 175 (2), 197–210.
- Axelrod, R., Hamilton, W.D., 1981. The evolution of cooperation. *Science* 211 (4489), 1390–1396.
- Baumard, N., Hyafil, A., Morris, I., Boyer, P., 2015. Increased affluence explains the emergence of ascetic wisdoms and moralizing religions. *Curr. Biol.* 25 (1), 10–15.
- Bergstrom, C.T., Kerr, B., Lachmann, M., 2008. Building Trust by Wasting Time. In: *Moral markets: The critical role of values in the economy*. Princeton University Press, pp. 142–156.
- Bliege-Bird, R., Ready, E., Power, E.A., 2018. The social significance of subtle signals. *Nat. Hum. Behav.* 1.
- Bolle, F., 2001. Why to buy your darling flowers: on cooperation and exploitation. *Theory Decis.* 50 (1), 1–28.
- Bulmer, M., Parker, G., 2002. The evolution of anisogamy: a game-theoretic approach. *Proc. R. Soc. Lond. Ser. B* 269 (1507), 2381–2388.
- Camerer, C., 1988. Gifts as economic signals and social symbols. *Am. J. Sociol.* 94, S180–S214.
- Carmichael, H.L., MacLeod, W.B., 1997. Gift giving and the evolution of cooperation. *Int. Econ. Rev.* 485–509.
- Debove, S., Baumard, N., André, J.-B., 2017. On the evolutionary origins of equity. *PLoS One* 12 (3), e0173636.
- Dyble, M., Gardner, A., Vinicius, L., Migliano, A., 2018. Inclusive fitness for in-laws. *Biol. Lett.* 14 (10), 20180515.
- Eshel, I., 1983. Evolutionary and continuous stability. *J. Theor. Biol.* 103 (1), 99–111.
- Eshel, I., Cavalli-Sforza, L.L., 1982. Assortment of encounters and evolution of cooperativeness. *Proceedings of the National Academy of Sciences* 79 (4), 1331–1335.
- Fiske, A.P., 1992. The four elementary forms of sociality: framework for a unified theory of social relations. *Psychol. Rev.* 99 (4), 689.
- Fletcher, J.A., Doebeli, M., 2008. A simple and general explanation for the evolution of altruism. *Proceedings of the Royal Society B: Biological Sciences* 276 (1654), 13–19.
- Frank, M., 1981. The Braess paradox. *Math. Program.* 20 (1), 283–302.
- Frank, R.H., 1988. *Passions Within Reason: The Strategic Role of the Emotions*. WW Norton & Co.
- Geoffroy, F., Baumard, N., André, J.-B., 2019. Why cooperation is not running away. *J. Evol. Biol.* 32 (10), 1069–1081.
- Gintis, H., Smith, E.A., Bowles, S., 2001. Costly signaling and cooperation. *J. Theor. Biol.* 213 (1), 103–119.
- Grafen, A., 1990. Biological signals as handicaps. *J. Theor. Biol.* 144 (4), 517–546.
- Hamilton, W.D., 1964. The genetical evolution of social behaviour. *J. Theor. Biol.* 7 (1), 17–52.
- Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., McElreath, R., 2001. In search of homo economicus: behavioral experiments in 15 small-scale societies. *Am. Econ. Rev.* 91 (2), 73–78.
- Henrich, J., Ensminger, J., McElreath, R., Barr, A., Barrett, C., Bolyanatz, A., Cardenas, J.C., Gurven, M., Gwako, E., Henrich, N., et al., 2010. Markets, religion, community size, and the evolution of fairness and punishment. *Science* 327 (5972), 1480–1484.
- Higham, J.P., 2013. How does honest costly signaling work? *Behav. Ecol.* 25 (1), 8–11.
- Jagau, S., van Veelen, M., 2017. A general evolutionary framework for the role of intuition and deliberation in cooperation. *Nat. Hum. Behav.* 1 (8), 152.
- Komiya, A., Ohtsubo, Y., Nakanishi, D., Oishi, S., 2019. Gift-giving in romantic couples serves as a commitment signal: relational mobility is associated with more frequent gift-giving. *Evol. Hum. Behav.* 40 (2), 160–166.
- Lettinga, N., Jacquet, P.O., André, J., Baumard, N., Chevallier, C., 2019. Environmental harshness is associated with lower investment in collective actions. *BioRxiv* 663518.
- Martinez-Vaquero, L.A., Han, T.A., Pereira, L.M., Lenaerts, T., 2015. Apology and forgiveness evolve to resolve failures in cooperative agreements. *Sci. Rep.* 5, 10639.
- Maynard Smith, J., 1982. *Evolution and the Theory of Games*. Cambridge University Press.
- McNamara, J.M., Barta, Z., Fromhage, L., Houston, A.I., 2008. The coevolution of choosiness and cooperation. *Nature* 451 (7175), 189.
- McNamara, J.M., Houston, A.I., Barta, Z., Osorno, J.-L., 2003. Should young ever be better off with one parent than with two? *Behav. Ecol.* 14 (3), 301–310.
- Nettle, D., Colléony, A., Cockerill, M., 2011. Variation in cooperative behaviour within a single city. *PLoS One* 6 (10), e26922.
- Noë, R., Hammerstein, P., 1994. Biological markets: supply and demand determine the effect of partner choice in cooperation, mutualism and mating. *Behav. Ecol. Sociobiol.* 35 (1), 1–11.
- Nowak, M.A., Sigmund, K., 1998. Evolution of indirect reciprocity by image scoring. *Nature* 393 (6685), 573.
- Pereira, L.M., Lenaerts, T., et al., 2017. Evolution of commitment and level of participation in public goods games. *Auton. Agent. Multi Agent Syst.* 31 (3), 561–583.
- Pinker, S., 2012. *The Better Angels of our Nature: Why Violence has Declined*. Penguin Group USA.
- Robson, A.J., 1990. Efficiency in evolutionary games: darwin, nash and the secret handshake. *J. Theor. Biol.* 144 (3), 379–396.
- Safra, L., Tecu, T., Lambert, S., Sheskin, M., Baumard, N., Chevallier, C., 2016. Neighborhood deprivation negatively impacts childrens prosocial behavior. *Front. Psychol.* 7, 1760.
- Santos, F.C., Pacheco, J.M., Skyrms, B., 2011. Co-evolution of pre-play signaling and cooperation. *J. Theor. Biol.* 274 (1), 30–35.
- Schelling, T.C., 1980. *The strategy of conflict*. Harvard university press.
- Skyrms, B., 2002. Signals, evolution, and the explanatory power of transient information. *Philos. Sci.* 69 (3), 407–428.
- Skyrms, B., 2004. *The Stag Hunt and the Evolution of Social Structure*. Cambridge University Press.
- Sozou, P.D., Seymour, R.M., 2005. Costly but worthless gifts facilitate courtship. *Proc. R. Soc. B* 272 (1575), 1877–1884.
- Spence, M., 1973. Job market signaling. *Q. J. Econ.* 87 (3), 355–374.
- Takesue, H., 2017. Partner selection and emergence of the merit-based equity norm. *J. Theor. Biol.* 416, 45–51.
- Tooby, J., Cosmides, L., 1996. Friendship and the banker's paradox: Other pathways to the evolution of adaptations for altruism. In: *Proceedings-British Academy*, 88. Oxford University Press, pp. 119–144.
- Trivers, R.L., 1971. The evolution of reciprocal altruism. *Q. Rev. Biol.* 46 (1), 35–57.
- Yamaguchi, M., Smith, A., Ohtsubo, Y., 2015. Commitment signals in friendship and romantic relationships. *Evol. Hum. Behav.* 36 (6), 467–474.
- Zahavi, A., 1975. Mate selection: a selection for a handicap. *J. Theor. Biol.* 53 (1), 205–214.