

Cooperation investments: Building capacity and signaling intent

JULIEN LIE-PANIS

University of Guelph & MIT
jliep@protonmail.com

CHRISTIAN HILBE

Interdisciplinary Transformation
University (IT:U), 4020 Linz,
Austria
christian.hilbe@it-u.at

Abstract

People do not just cooperate – they regularly invest in their cooperative abilities. Partners learn to resolve conflicts, colleagues train for their roles within large teams, and people adopt the shared practices of their community. Often, these activities do not benefit others per se. Rather they facilitate future cooperative interactions. Yet such cooperation investments are rarely captured in evolutionary game theory, which focuses on cooperative decisions instead of the preparatory efforts that enable them. Here, we develop a model that allows individuals to incur upfront costs to lower their future cooperation costs. Through mathematical analysis and evolutionary simulations, we show that cooperation investments serve a dual function. First, they allow individuals to transform into more effective partners, enabling cooperation where it would otherwise be too costly. Second, they can serve as honest signals of cooperative intent. When their costs deter would-be cheaters, investments allow observers to identify trustworthy partners, further expanding the cooperative domain. Our model helps to explain why people are attentive to others’ cooperation investments. When someone takes the time to learn more about a partner’s hobbies or a community’s norms, they do more than build capacity – they signal an intent to cooperate with us.

Keywords: cooperation, cooperative capacity, signaling, reputation, evolutionary game theory

Cooperation takes work. Consider our closest relationships: across cultures, people build friendships (1) and romantic partnerships (2, 3) gradually, through sustained effort. Over time, partners learn to resolve conflicts (4), become more accurate at reading each other’s thoughts and feelings (5), and divide mental tasks, with each partner specializing in different kinds of knowledge (6). The demands of cooperation extend beyond intimate relationships, from professional roles to community life. Healers—from shamans to physicians—spend

years in training, developing the expertise that makes their care valuable to others (7, 8). Hunters and engineers similarly spend years learning specialized skills, and also learning to function in the large, coordinated teams on which success depends (9, 10). Community membership, too, requires investment: we absorb local practices, internalize group norms, and acquire the culturally specific knowledge that makes us effective members (11–15).

These behaviors are cooperation investments: costly efforts that build one’s cooperative capacity. Unlike cooperative acts, which directly benefit others, investments shape the individual, lowering their future cooperation costs. Most models in evolutionary game theory do not include such investments. They treat cooperation costs as fixed and outside individual control (16–36). As a result, such models cannot address how we become reliable partners, and how others draw inferences from our preparatory efforts.

To address this gap, we develop a formal model that allows for up-front investments into one’s cooperative abilities. The model is based on a repeated trust game between choosers (trust-givers) and actors (reciprocators). Actors differ in how much they value future payoffs – while some are more future-oriented, others care more about the present. This variation in actors’ evaluations of the future induces variation in their willingness to reciprocate a chooser’s trust. We study this repeated trust game with and without the possibility of cooperative investments: the possibility that actors pay an extra cost before any interaction occurs, to permanently lower the cost of reciprocating. With a comprehensive mathematical analysis, complemented by evolutionary simulations, we ask: Who would be willing to make such investments, and how do these investments facilitate cooperation?

Our results reveal that cooperation investments serve two distinct functions. First, they allow individuals to transform themselves into more effective partners. In this way, they enable cooperation where it would be too costly otherwise. Second, they can serve as honest signals of cooperative intent. When their costs deter impatient indi-

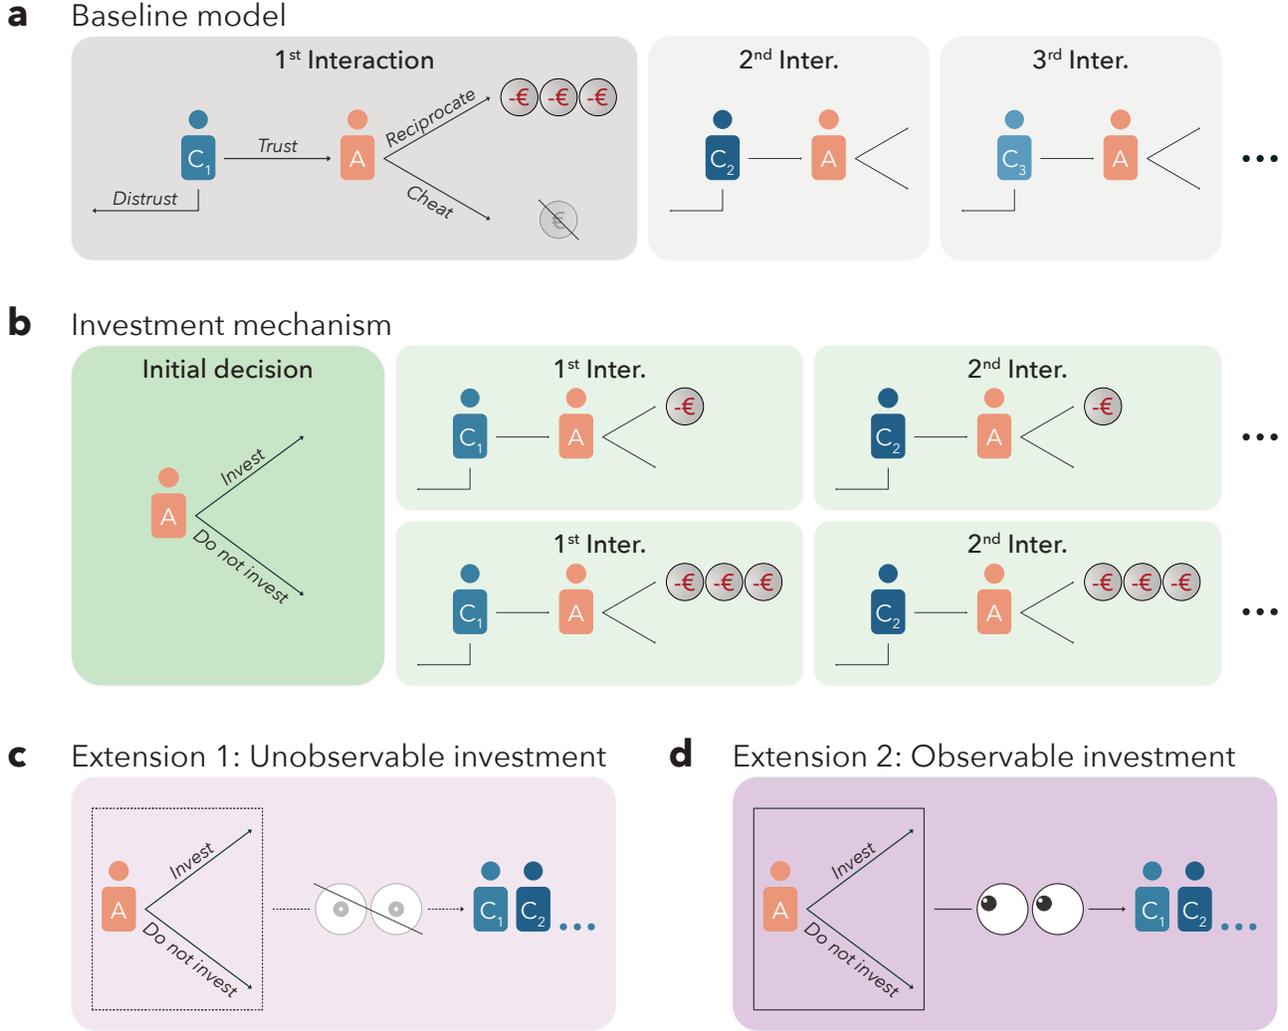


Fig. 1 | Model overview. We model cooperation investments in three steps. **a**, In the baseline model, a long-run actor (A) plays a repeated trust game with a series of short-run choosers (C_1, C_2, C_3, \dots). In each interaction, a chooser decides whether to trust the actor based on their reputation from past behavior. If trusted, the actor chooses whether to reciprocate or cheat. By default, the cost of reciprocation c_H is high, making cooperation between the actor and choosers unsustainable. **b**, In both our model extensions, we allow the actor to invest in cooperation before interacting with choosers, by incurring a cost k to permanently lower their cost of reciprocation to $c_L < c_H$. **c**, In the first extension, choosers do not observe the actor’s investment decision; in the second extension, they do, allowing them to condition trust on whether the actor initially invested.

viduals, investments allow observers to identify trustworthy partners, further expanding the cooperative domain.

This observation in turn can explain why people are often attentive to others’ investments. By taking the time to work through personal conflict (37) or to learn another community’s language (38), we do not only make us more effective cooperators in future. We also indicate that we intend to cooperate in the first place. By clarifying this dual function of cooperation investments, our model offers a framework for understanding how humans develop, evaluate and structure cooperation.

Results

We introduce our model in three steps. First, we establish a baseline model in which actors face fixed cooperation costs (Fig. 1a). Second, we extend the model by introducing an additional stage prior to the actual interaction. In this prior stage, actors can make an investment to reduce their future cooperation costs (Fig. 1b). Whether or not the investment is made remains an actor’s private knowledge (Fig. 1c). In a third step, we study what happens when investment decisions are observable to an ac-

tor’s potential partners (**Fig. 1d**). To make the underlying mechanisms most transparent, we initially work with a simplified version of our model. In particular, we only allow for two types of actors with two possible time preferences (more present-oriented or more future-oriented). Later on, we consider the case of a continuum of types. In the main text, we present a summary of the respective findings; further details can be found in the **Methods** section and the **Supplementary Information**.

Baseline model

The baseline scenario follows the setup of Lie-Panis & André (23): We consider a repeated trust game in which an actor repeatedly encounters different choosers (**Fig. 1a**). The actor makes cooperative decisions across rounds and acquires a reputation based on those decisions. Choosers participate in only one round. They decide to trust the actor based on reputation (in the literature on repeated games, such choosers and actors would be sometimes referred to as short-lived and long-lived players, respectively, see Ref. (39)). This setup separates two adaptive problems. For choosers, the challenge is trust—predicting whether a partner will reciprocate trust if given the opportunity. For the actor, the challenge is reputation management—acting in ways that secure the trust of future partners.

To capture uncertainty about how the actor manages their reputation, we introduce heterogeneous time preferences. Before any interaction occurs, the actor is assigned a discount factor δ between zero and one. This discount factor measures how much the actor values future payoffs relative to immediate ones. More patient actors (higher δ) are more willing to incur the immediate costs of cooperation to maintain their reputation with choosers. In contrast, impatient actors (those with lower δ) prioritize short-term gains and are more prone to cheat. For simplicity, we assume a binary distribution: with probability p , the actor is future-oriented ($\delta = \delta_F$), and with probability $1 - p$, they are present-oriented ($\delta = \delta_P < \delta_F$). The actor privately observes their discount factor before the game begins; choosers only know the distribution.

In each round, a new chooser decides whether to trust or distrust the actor. Trust costs the chooser γ and gives the actor a benefit b (we use greek and latin letters to denote payoff components related to the chooser and actor, respectively). If trusted, the actor decides whether to reciprocate or cheat. Reciprocation costs the actor c_H and gives the chooser a benefit β . An actor’s choices affects their reputation. Initially, their reputation is *unknown* until they are first trusted, it is *good* if they last reciprocated, and it is *bad* if they last cheated. This reputation persists until the next time the actor is trusted, regardless of how many rounds pass without trust.

To make their decisions, both kinds of players implement a strategy. A chooser’s strategy needs to specify whether or not to trust the actor, depending on the actor’s reputation. An actor’s strategy, on the other hand, specifies whether to reciprocate trust. This decision may depend on the actor’s discount factor and their current reputation. For all our model variations, we ask under which conditions there is a (Perfect Bayesian Nash) equilibrium along which choosers at least trust occasionally, and some actors decide to reciprocate. In that case, we call the equilibrium *cooperative*.

The baseline model of course admits a trivial equilibrium, in which choosers never trust and actors never reciprocate. However, it also allows for a unique cooperative equilibrium. In this equilibrium, choosers trust actors with an unknown or good reputation but distrust those with a bad reputation. Whatever their reputation, future-oriented actors reciprocate choosers’ trust, while present-oriented actors always cheat. This equilibrium exists if and only if three conditions are satisfied,

$$c_H \leq \delta_F b, \quad (\text{B.1})$$

$$c_H > \delta_P b, \quad (\text{B.2})$$

$$\gamma \leq p \beta. \quad (\text{B.3})$$

Condition (B.1) ensures that future-oriented actors find it worthwhile to reciprocate: they pay the cost c_H now to secure the trust of the next chooser, worth $\delta_F b$. Condition (B.2) ensures that present-oriented actors prefer to cheat: they avoid the cost c_H , even though this means losing the trust of the next chooser, worth only $\delta_P b$. This behavioral variation is essential—without it, choosers would have no reason to condition trust on reputation. Were, for example, both types to consistently reciprocate, it would be beneficial for choosers to trust even given bad reputation. As a result, trust would be guaranteed for the actor, removing any incentive to reciprocate in the first place. Finally, condition (B.3) ensures that choosers are willing to trust someone with unknown reputation despite the risk, paying cost γ for an expected benefit $p \beta$.

From this point forward, we focus on scenarios where condition (B.1) is violated ($c_H > \delta_F b$). In that case, reciprocation is prohibitively costly even for future-oriented actors, and cooperation unravels. These are the situations where cooperation requires capacity building: partners unfamiliar with conflict resolution, untrained professionals, or newcomers to a community all initially face high cooperation costs.

Investment model

When reciprocation is prohibitively costly, the baseline model no longer allows for cooperation. Thus we extend the model by allowing the actor to build up their cooperative capacity. Before any interaction with choosers,

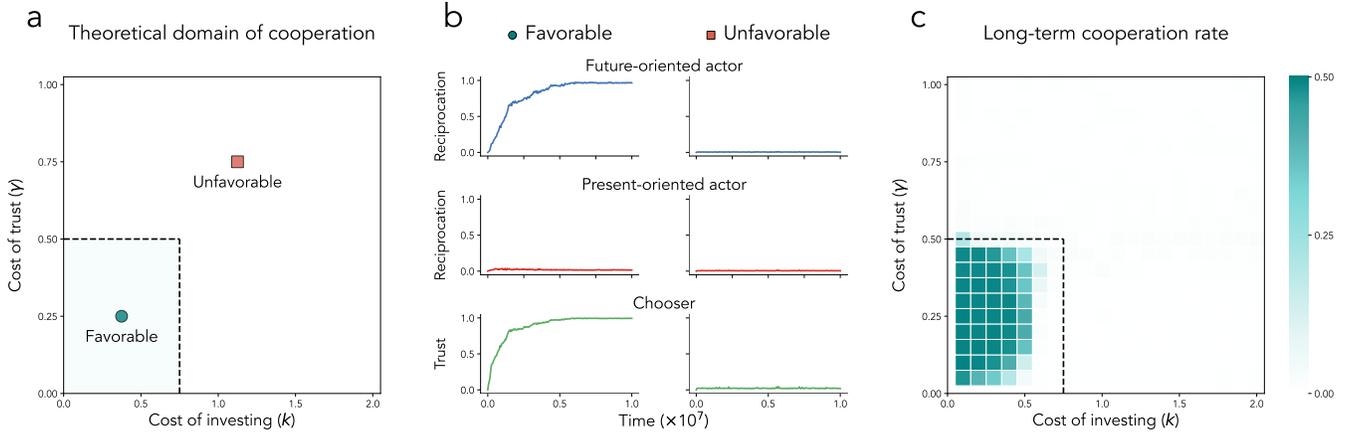


Fig. 2 | Investment enables cooperation. **a**, Investment enables cooperation when both the cost of investing k and the cost of trust γ are sufficiently low, even where the baseline model fails ($c_H > \delta_F b$). Blue circle and red square indicate favorable ($k = 3/8$, $\gamma = 1/4$) and unfavorable ($k = 11/8$, $\gamma = 3/4$) parameter sets, respectively. **b**, Simulation trajectories under both parameter sets: reciprocation rates for future-oriented actors (top) and present-oriented actors (middle), and trust rates for choosers (bottom). Under favorable conditions, choosers learn to trust at high rates and future-oriented actors learn to reciprocate. **c**, Long-term cooperation rates across the parameter space after 10^7 simulation steps. Cooperation occurs within the predicted domain but vanishes before investment costs reach their theoretical threshold. Other parameters: $p = 0.5$, $\delta_F = 0.75$, $\delta_P = 0.25$, $b = 1.0$, $c_L = 0.5$, $c_H = 0.9$, $\beta = 1.0$. Panels **b,c** show averages over 100 simulations with $n = 50$ actors and choosers, mutation rate $\mu = 0.02$, and selection strength $s = 10.0$.

the actor may decide to invest, paying a one-time cost k to permanently reduce their cost of reciprocation from c_H to c_L (Fig. 1b). Here, we assume the actor’s investment decision is not observed by choosers. Moreover, for simplicity we assume that investment separates between the two types: it makes reciprocation worthwhile for future-oriented actors ($c_L \leq \delta_F b$) but not for present-oriented ones ($c_L > \delta_P b$). Now, an actor’s strategy needs to specify whether to initially invest based on their discount factor, in addition to specifying reciprocation behavior as in the baseline model. The choosers’ strategies are unchanged.

Again, there is a unique cooperative equilibrium, despite the high initial reciprocation costs ($c_H > \delta_F b$). In this equilibrium, choosers trust actors with unknown or good reputations but distrust those with bad reputations. Future-oriented actors initially invest and subsequently reciprocate the trust of all their partners, while present-oriented actors opt out of investing and always cheat. This reflects how upfront efforts—such as building relationship understanding, professional training, or learning local customs—can enable cooperation with committed individuals. The equilibrium exists if and only if

$$k \leq \frac{\delta_F}{1 - \delta_F} (\delta_F b - c_L), \quad (\text{I.1})$$

$$\gamma \leq p\beta. \quad (\text{I.2})$$

Condition (I.1) ensures that future-oriented actors find investment worthwhile, incurring cost k to obtain the benefits of long-term cooperation. Present-oriented actors never invest, because reciprocation remains too costly for them even at the reduced rate ($c_L > \delta_P b$). Finally, condition (I.2) ensures choosers find initial trust worthwhile, as before. Fig. 2a depicts in which domain of the parameter space this cooperative equilibrium exists, varying the cost of investing k (x -axis) and the cost of trust γ (y -axis). Cooperation is feasible when both costs are sufficiently low, as per conditions (I.1-I.2).

So far, we have taken an equilibrium perspective. We explored whether cooperation can persist among rational individuals who fully comprehend the incentive structure of the game. Alternatively, we can also ask whether cooperation can emerge when strategies are adopted through social learning. To this end, we conduct evolutionary simulations, considering a population of choosers and actors, respectively. For simplicity, for these simulations we use a more conventional repeated-game setup. Choosers and actors are randomly matched to play indefinitely in a fixed pair (i.e., both players are now long-lived). The feasible strategies are the same as before. As a result of their interactions, choosers and actors accumulate a payoff, which other players in the respective population can observe. This allows players to imitate more successful

strategies. We model the resulting imitation dynamics as a pairwise-comparison process (40). We also allow for random mutations, which introduces novel strategies into the evolving population (see **Methods** for details).

We first focus on two parameter conditions, highlighted in **Fig. 2a**: a favorable condition within the cooperative domain (blue circle), and an unfavorable condition outside of it (red square). **Fig. 2b** shows how cooperation unfolds over time. Under favorable conditions, choosers learn to trust and future-oriented actors learn to reciprocate at high rates. Reciprocation from present-oriented actors remains rare, as expected. Under unfavorable conditions, trust and reciprocation remain consistently low.

To generalize these findings, we run simulations across a range of parameters. **Fig. 2c** shows the resulting long-term rate of mutual cooperation. Cooperation occurs where predicted, but it vanishes before investment costs reach their theoretical threshold (As before, there is a trivial equilibrium in which no ones trusts or reciprocates, in addition to the cooperative equilibrium. This trivial equilibrium exists for all parameter combinations).

Observable investment model

In the previous model, investments were hidden from choosers. Their only effect was to reduce the actor’s cost of reciprocation. In the following extension, we allow choosers to observe the actor’s investment decision (**Fig. 1d**), introducing potential reputational benefits. In this model, the actor’s reputation reflects both their investment choice (invested or not) and their most recent cooperation history (unknown, good, or bad), resulting in six possible reputations. Both chooser and actor strategies depend on this expanded reputation set; the rest of the model is unchanged. We show that this extension admits two cooperative equilibria: a pooling equilibrium, where both types invest and remain indistinguishable, and an honest signaling equilibrium, where only future-oriented actors invest and thereby reveal their cooperative intent. We detail each in turn.

In the honest signaling equilibrium, future-oriented actors invest and always reciprocate, while present-oriented actors do not invest and never reciprocate. Choosers trust investors and distrust non-investors. Subsequently, trust depends on cooperation history: choosers trust investors with a good reputation but distrust those with a bad reputation. This equilibrium exists if and only if

$$k \leq \frac{\delta_F}{1 - \delta_F}(b - c_L), \quad (\text{HS.1})$$

$$k > \delta_P b, \quad (\text{HS.2})$$

$$\gamma \leq \beta. \quad (\text{HS.3})$$

Condition (HS.1) ensures that future-oriented actors find investment worthwhile, paying k upfront to secure

long-term cooperation with choosers. Because investment is now required to be trusted, it becomes worthwhile at higher costs than before, as reflected in the larger right-hand side of condition (HS.1) compared to (I.1). Condition (HS.2) ensures that present-oriented actors are deterred from investing: the benefit of exploiting the first chooser, worth $\delta_P b$, cannot cover the upfront cost k . As a result of this deterrence, an actor’s investment reliably indicates future cooperative behavior. Condition (HS.3) then ensures that choosers benefit from trusting such investors, paying cost γ to receive benefit β with certainty. Compared to the previous two models, cooperation is now sustainable for higher trust costs, where choosers had to initially trust without knowing the actor’s type. This captures how investments that are sufficiently demanding—such as extensive professional training or learning a community’s language—can provide credible signals that allow observers to identify trustworthy partners.

When costs are too low to deter present-oriented actors, we obtain a pooling equilibrium. In this equilibrium, both types invest, making it impossible to distinguish between present-oriented cheaters and future-oriented reciprocators. As in the honest signaling equilibrium, choosers initially trust investors and distrust non-investors. Afterwards, trust depends on cooperation history. This equilibrium exists if and only if

$$k \leq \delta_P b, \quad (\text{P.1})$$

$$\gamma \leq p\beta. \quad (\text{P.2})$$

Condition (P.1) ensures that present-oriented actors find it advantageous to initially invest, which guarantees that future-oriented actors do so as well. Since investment no longer provides information about actor type, choosers face the same uncertainty as in earlier models: condition (P.2) reverts to the baseline trust requirement.

Fig. 3a depicts the domain in parameter space where these two cooperative equilibria exist, varying the cost of investing k (x -axis) and the cost of trust γ (y -axis). Pooling occurs for low values of k and γ , covering a subset of the cooperative domain in the previous model (**Fig. 2a**). In contrast, honest signaling expands cooperation to a broader region, encompassing higher values of k and any $\gamma < \beta$. This expansion occurs because investment now provides a reputational benefit, making it worthwhile at higher costs, and because it fully separates actor types, enabling trust whenever cooperating with a trustworthy actor is beneficial.

We again use evolutionary simulations to examine whether these cooperative equilibria emerge when strategies evolve through social learning, using the same parameter combinations as before. **Fig. 3b** shows that unlike in the previous model, cooperation now emerges even in the unfavorable case, at low rates. Under favorable conditions, cooperation also arises, albeit at more

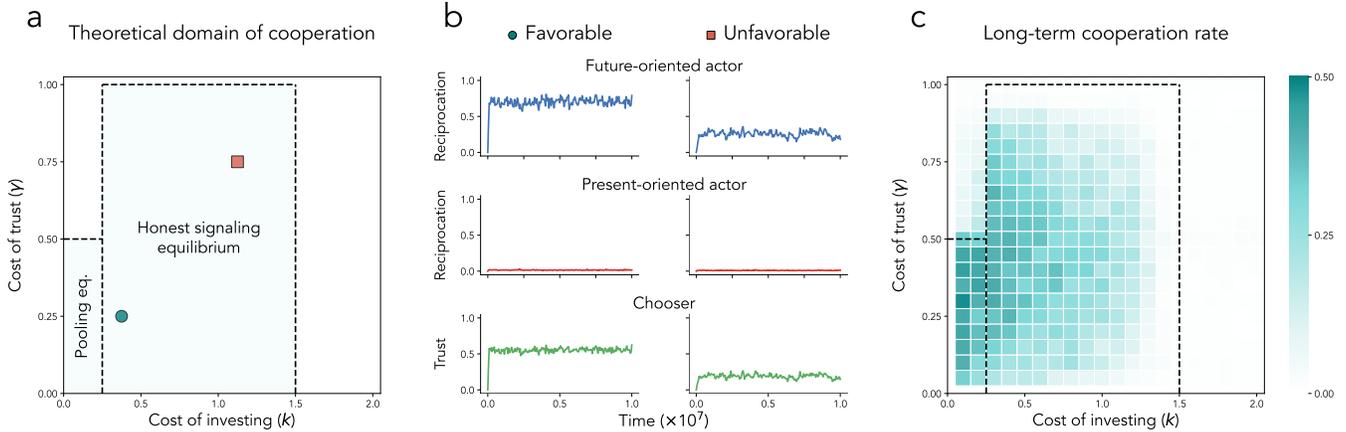


Fig. 3 | Observation expands cooperation. **a**, When investment costs deter present-oriented actors ($k > \delta_P b$), an honest signaling equilibrium is possible, enabling cooperation beyond Figure 2—at higher investment costs and any trust cost below the benefit of receiving reciprocation ($\gamma < \beta$). Blue circle and red square show the same parameter sets as before. **b**, Simulation trajectories under both parameter sets: reciprocation rates for future-oriented actors (top) and present-oriented actors (middle), and trust rates for choosers (bottom). Cooperation between future-oriented actors and choosers now emerges even in the unfavorable case, at low rates. Under favorable conditions, cooperation also arises, though at more moderate levels than before. **c**, Long-term cooperation rates across the parameter space after 10^7 simulation steps. Cooperation occurs within predicted regions but at moderate rates, vanishing before reaching theoretical boundaries. Other parameters: same values as in Figure 2.

moderate levels than before. Fig. 3c shows long-term cooperation across the entire parameter space. While cooperation evolves for a far larger set of parameter values, the evolving cooperation rates remain somewhat moderate. Especially as investment and trust costs increase, cooperation rates drop, vanishing before they reach their theoretical thresholds.

A general model with arbitrary time preferences

For simplicity, the previous model assumes two distinct types of actors, who either are present-oriented (with $\delta = \delta_P$) or future-oriented (with $\delta = \delta_F$). However, our results extend beyond this binary case. In the **Supplementary Information**, we develop a model in which the actors’ time preferences are described by an arbitrary distribution over the space of possible time preferences $\delta \in (0, 1)$. Again, we provide necessary and sufficient conditions for the existence of cooperative equilibria (those where there is a positive probability to observe a chooser trust and an actor to reciprocate).

Fig. 4 shows results when time preferences are distributed according to a truncated normal distribution. We observe the same qualitative pattern. In the baseline model, cooperation is only feasible under rather restrictive conditions (blue). These conditions become easier to satisfy if we allow for cooperation investments (teal), especially when those investments are observable (dark

blue). With continuous distributions, cooperation is feasible in the baseline model whenever trust costs are sufficiently low, because patient reciprocating actors always exist (provided that $c_H < b$). The domain expanded by investment shows smoother boundaries than before, depending on the fraction of actor types for whom investment makes reciprocation worthwhile. The observable investment domain covers nearly the entire parameter space. For large investment costs, perfect separation occurs: only the most patient actors invest and reciprocate, while others do neither. This enables cooperation even when reciprocating types are vanishingly rare.

Discussion

People often invest time, effort, and resources to transform themselves into more effective cooperators. Herein, we introduce an evolutionary game-theoretic model of such cooperation investments to better understand when and why they occur, and how they facilitate cooperation. We proceed in three steps. First, we establish a baseline model where cooperation costs are fixed. Here, cooperation can emerge if individuals are willing to incur these costs for future reputational benefits. However, once cooperation is prohibitively costly, even the most patient individuals who value long-term relationships prefer to cheat. In a second step, we allow individuals to invest

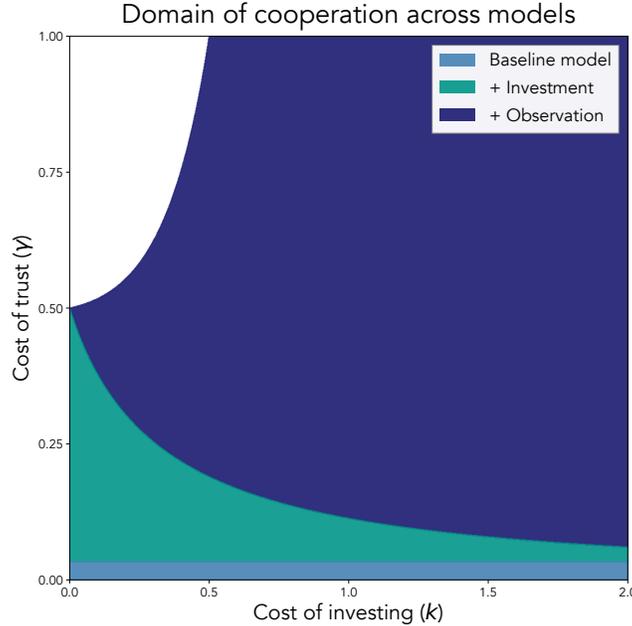


Fig. 4 | Generalization to continuous discount factors. Replacing the binary distribution from Figures 2-3 ($\delta_P = 0.25$, $\delta_F = 0.75$) with a continuous one (truncated normal, mean 0.5, sd 0.25) yields the same qualitative pattern: cooperation under restrictive conditions in the baseline model (blue), expanded by investment (teal), and further expanded by making investment decisions observable (dark blue). Other parameters as in Figures 2-3.

in reducing their cooperation costs. Patient individuals can now make cooperation worthwhile for themselves, enabling it where it would otherwise fail. Finally, we make investments observable to potential partners. When investments are sufficiently costly, we find that only individuals committed to long-term cooperation find them worthwhile. This allows observers to identify trustworthy partners, further expanding the conditions under which cooperation can emerge.

Overall, we show that cooperation investments serve a dual function. On the one hand, they allow individuals to build capacity, enabling cooperation at a lower cost. Second, cooperative investments often serve as credible signals of cooperative intent. In the right parameter range, the actors who invest in their cooperative abilities are exactly the actors who in fact cooperate subsequently. In this way, the model sheds light on a number of interesting questions related to cooperation investments.

Who invests in cooperation? Our model predicts that individuals will invest in becoming more effective cooperators when they are patient—when they value the long-term benefits of a relationship, role, or community, enough to justify upfront costs. When investments are observable, they provide an additional advantage: they signal commitment to partners, making trust more likely. These two functions are typically inseparable. Consider open source software (41). Widely-used tools like Linux

exist thanks to the cooperation of skilled contributors, who take the time to learn coding standards, master review protocols, and absorb community norms. Because these efforts are publicly visible on platforms like GitHub, contributors not only become better collaborators; they also enhance their standing with fellow programmers and potential employers.

Why do we attend to others’ investments? Our model also explains why people pay close attention to others’ cooperation investments – because they can provide clues about others’ intentions and their incentives to act socially. We support immigrants’ permanent residency based on their efforts to learn our language quickly (38)—efforts that demonstrate commitment to community membership while building cultural competence. Similarly, we are more likely to make personal sacrifices for invested partners (see also 42), such as those who strive to understand us (43) or with whom we have worked through conflict (37). Our model also offers an additional lens on commitment displays. Hazing (44), loyalty tests (45), and restrictive religious prohibitions (46) are typically understood as signals of commitment (47) or as screening devices that weed out free-riders (48). They may also build cooperative capacity, or be perceived to do so by those who enact them. Tech interviews push candidates to develop relevant skills in advance; initiations can instill shared norms and cultural knowledge (49, 50);

and shared ordeals may be perceived to strengthen interpersonal bonds (51, 52).

The structure of human cooperation. Cooperation investments are inherently selective. Individuals cannot optimize their cooperative capacity across every relationship, profession, or community. While our model includes only one investment decision, it provides a microfoundational mechanism for structured cooperation: as people choose where to invest, they become better equipped to cooperate with certain partners or groups, and therefore more likely to cooperate with them. Future work could explore how individuals allocate effort across competing relationships, how observers vie for their time and commitment, how groups impose investment requirements, and how these dynamics jointly shape the structure of human cooperation.

Methods

Equilibrium analysis

For the general **baseline model**, we consider a repeated trust game between a long-run actor and a sequence of short-run choosers. Before play begins, the actor is assigned a discount factor $\delta \in (0, 1)$, according to a distribution with support Δ . Only the actor observes their discount factor, but the distribution is common knowledge. In the main text, we focus on the case of a binary distribution. The actor is future-oriented ($\delta = \delta_F$) with probability $p \in (0, 1)$ and present-oriented ($\delta = \delta_P < \delta_F$) with probability $1 - p$. In the **Supplementary Information**, we allow for any generic distribution (discrete, continuous, or mixed).

In every round $t \geq 1$, the actor is paired with a new chooser, who decides whether to trust the actor. A chooser who trusts incurs a cost $\gamma > 0$ to give the actor a benefit $b > 0$. If trusted, the actor can reciprocate, incurring a cost $c_H > 0$ to give the chooser a benefit $\beta > 0$ in return. Alternatively, the actor can decide to cheat, keeping the entire benefit b without reciprocating. These choices affect the actor’s subsequent reputation. We assume there are three different kinds of reputation an actor can have. The actor’s reputation is unknown until they are first trusted, good if they last reciprocated, and bad if they last cheated. Hence, the set of possible reputations is $\Omega \equiv \{\text{unknown, good, bad}\}$. Choosers hold shared beliefs $\mu(\cdot | \omega)$ over Δ for each reputation $\omega \in \Omega$, encoding how likely they deem an actor to have a certain discount factor, conditional on the actor’s observed reputation.

A strategy for the chooser is a function $\sigma_{\text{ch}} : \Omega \rightarrow \{\text{trust, distrust}\}$. It specifies whether to trust the actor based on the actor’s reputation. An actor strategy, on the other hand, is represented by a function $\sigma_{\text{ac}} : \Delta \times \Omega \rightarrow \{\text{reciprocate, cheat}\}$. It specifies whether to reciprocate a

chooser’s trust based on the actor’s discount factor and current reputation.

For the **investment model**, we add a preliminary stage at $t = 0$, in which the actor may decide to invest. Such an investment incurs a one-time cost of $k > 0$, but it permanently lowers the reciprocation cost from c_H to c_L , where $0 < c_L \leq c_H$ and $c_L < b$. Choosers do not observe the actor’s investment decision. As a result, they continue to rely on the same set of reputations. An investment strategy, represented by a function $\sigma_{\text{inv}} : \Delta \rightarrow \{\text{invest, do not invest}\}$, specifies whether to invest based on discount factor. The actor’s full strategy profile, $(\sigma_{\text{inv}}, \sigma_{\text{ac}})$, comprises both the initial investment decision and subsequent behavior in trust games. The rest of the model is unchanged.

In the **observable investment model**, choosers are able to observe the actor’s initial investment decision. As a result, the actor’s cost of reciprocation is now common knowledge; it is c_L if they invested, or c_H if they did not. The set of possible reputations expands to $\Omega' = \{\text{invested, did not invest}\} \times \Omega$, reflecting initial investment decision and behavior in trust games. All strategies and beliefs are now defined with respect to this expanded set. Trust and reciprocation strategies are defined over Ω' , and choosers hold beliefs for each $\omega \in \Omega'$.

In each case, we study the model’s Perfect Bayesian Equilibria (PBEs) in pure strategies. A PBE is a pair (σ, μ) where players have no profitable deviations at any history (on or off the equilibrium path), and beliefs are updated according to Bayes’ rule whenever possible (53). In the main text, we focus on the subset of cooperative equilibria—PBEs in which mutual cooperation occurs with positive probability: along the equilibrium path, some actor types are trusted and reciprocate that trust. In the Supplementary Information, we also study the model’s non-cooperative PBEs. For simplicity, our analysis assumes that players choose the cooperative action (invest, reciprocate, trust) when indifferent; these cases occur only at knife-edge parameter values.

Evolutionary simulations

To complement our general equilibrium analysis, we run evolutionary simulations for the simplified model with a binary distribution of discount factors. We simulate $n = 50$ actors and $n = 50$ choosers, run each simulation for 10^7 time steps, and report averages over 100 independent runs. Actors are future-oriented ($\delta = \delta_F$) with probability p and present-oriented ($\delta = \delta_P$) with probability $1 - p$.

In contrast to the model used for our mathematical analysis, the simulations are based on a conventional repeated game setup. Choosers are randomly matched with actors. Then they engage in an infinitely repeated game

(with choosers discounting the future with some fixed and uniform factor δ). This assumption makes it straightforward to compute a chooser’s expected payoff, which is necessary for strategy updating (see **Supplementary Information** Section 4.3).

We encode strategies as bit vectors. Actor strategies specify whether to invest and whether to reciprocate for both discount factors, yielding $2^4 = 16$ strategies in total. As we demonstrate in the **Supplementary Information**, actors do not condition their behavior on their own reputation in equilibrium. Based on this insight, we restrict attention to reputation-independent strategies. Chooser strategies specify whether to trust given the actor’s reputation, yielding $2^3 = 8$ strategies in the investment model (3 possible reputations) and $2^6 = 64$ in the observable investment model (6 possible reputations).

To model how individuals update their strategies over time, we use a pairwise comparison process (40). At each time step, a randomly selected individual revises their strategy through either random mutation (with probability μ) or payoff-based imitation. Under imitation, the individual compares their payoff π_{self} with that of a randomly selected member of their own population type, π_{other} . These payoffs depend on each individual’s strategy and the current distribution of strategies in both populations (formulas provided in the **Supplementary Information**). The focal individual adopts the other’s strategy with probability given by the Fermi function,

$$\mathbb{P}(\text{adopt}) = \frac{1}{1 + e^{-s(\pi_{\text{other}} - \pi_{\text{self}})}}.$$

The parameter s is the selection strength, indicating how strongly individuals are affected by payoff differences when deciding to imitate others.

For the simulations, we focus on cases where cooperation cannot be sustained in the baseline model, and run simulations for the investment and observable investment models. Unless noted otherwise, we use $p = 0.5$, $\delta_F = 0.75$, $\delta_P = 0.25$, $\delta = 0.5$, $b = 1.0$, $\beta = 1.0$, $c_H = 0.9$, and $c_L = 0.5$. With these values, reciprocation is prohibitively costly without investment but becomes viable for future-oriented actors once they invest. As evolutionary parameters, we use a selection strength of $s = 10.0$ and a mutation rate of $\mu = 0.02$.

Acknowledgments. JLP acknowledges funding from the Social Sciences & Humanities Research Council of Canada (SSHRC), Grant number 430287. CH acknowledges generous support by the European Research Council Starting Grant 850529 (E-DIRECT).

References

1. Hruschka, D. J. *Friendship: Development, Ecology, and Evolution of a Relationship* (University of California Press, Berkeley, CA, 2010).
2. Sorokowski, P. *et al.* Love as a universal phenomenon: Data from nine non-WEIRD societies. *International Journal of Intercultural Relations* **108**, 102211 (2025).
3. Jankowiak, W. R. & Fischer, E. F. A cross-cultural perspective on romantic love. *Ethnology* **31**, 149–155 (1992).
4. Gottman, J. M., Coan, J., Carrere, S. & Swanson, C. Predicting marital happiness and stability from newlywed interactions. *Journal of Marriage and the Family* **60**, 5–22 (1998).
5. Stinson, L. & Ickes, W. Empathic accuracy in the interactions of male friends versus male strangers. *Journal of Personality and Social Psychology* **62**, 787–797 (1992).
6. Wegner, D. M., Erber, R. & Raymond, P. Transactive memory in close relationships. *Journal of Personality and Social Psychology* **61**, 923–929 (1991).
7. Lightner, A. D., Heckelsmiller, C. & Hagen, E. H. Ethnoscience expertise and knowledge specialisation in 55 traditional cultures. *Evolutionary Human Sciences* **3**, e37 (2021).
8. Hagen, E. H., Garfield, Z. H. & Lightner, A. D. Headmen, shamans, and mothers: Natural and sexual selection for computational services. *Evolution and Human Behavior* **46**, 106651 (2025).
9. Alvard, M. S. & Nolin, D. A. Rousseau’s whale hunt? Coordination among big-game hunters. *Current Anthropology* **43**, 533–559 (2002).
10. Hooper, P. L., Demps, K., Gurven, M., Gerkey, D. & Kaplan, H. S. Skills, division of labour and economies of scale among Amazonian hunters and South Indian honey collectors. *Philosophical Transactions of the Royal Society B: Biological Sciences* **370**, 20150008 (2015).
11. Lave, J. & Wenger, E. *Situated Learning: Legitimate Peripheral Participation* (Cambridge University Press, Cambridge, 1991).
12. Amir, D. *et al.* The emergence of cooperative behaviors, norms, and strategies across five diverse societies 2024. https://osf.io/preprints/psyarxiv/f7ud4_v1.

13. Lew-Levy, S., Lavi, N., Reckin, R., Cristóbal-Azkarate, J. & Ellis-Davies, K. How do hunter-gatherer children learn social and gender norms? A meta-ethnographic review. *Cross-Cultural Research* **52**, 213–255 (2018).
14. House, B. R. *et al.* Universal norm psychology leads to societal diversity in prosocial behaviour and development. *Nature Human Behaviour* **4**, 36–44 (2020).
15. Lew-Levy, S., Reckin, R., Lavi, N., Cristóbal-Azkarate, J. & Ellis-Davies, K. How do hunter-gatherer children learn subsistence skills? A meta-ethnographic review. *Human Nature* **28**, 367–394 (2017).
16. Axelrod, R. & Hamilton, W. D. The evolution of cooperation. *Science* **211**, 1390–1396. ISSN: 0036-8075, 1095-9203. <https://science.sciencemag.org/content/211/4489/1390> (2021) (Mar. 27, 1981).
17. Nowak, M. A. & Sigmund, K. Evolution of indirect reciprocity by image scoring. *Nature* **393**, 573–577. ISSN: 1476-4687. <https://www.nature.com/articles/31225> (2020) (June 1998).
18. Panchanathan, K. & Boyd, R. A tale of two defectors: the importance of standing for evolution of indirect reciprocity. *Journal of Theoretical Biology* **224**, 115–126. ISSN: 00225193. <https://linkinghub.elsevier.com/retrieve/pii/S0022519303001541> (2020) (Sept. 2003).
19. Leimar, O. Reciprocity and communication of partner quality. *Proceedings of the Royal Society of London. Series B: Biological Sciences* **264**, 1209–1215. ISSN: 1471-2954. <http://www.royalsocietypublishing.org/doi/10.1098/rspb.1997.0167> (2019) (Aug. 22, 1997).
20. Leimar, O. & Hammerstein, P. Evolution of cooperation through indirect reciprocity. *Proceedings of the Royal Society of London. Series B: Biological Sciences* **268**, 745–753. <https://royalsocietypublishing.org/doi/abs/10.1098/rspb.2000.1573> (2020) (Apr. 7, 2001).
21. Roberts, G. Competitive altruism: from reciprocity to the handicap principle. *Proceedings of the Royal Society B: Biological Sciences* **265**, 427–431. ISSN: 0962-8452. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1688900/> (2025) (Mar. 7, 1998).
22. André, J.-B. The Evolution of Reciprocity: Social Types or Social Incentives? *The American Naturalist* **175**, 197–210. ISSN: 0003-0147, 1537-5323. <https://www.journals.uchicago.edu/doi/10.1086/649597> (2019) (Feb. 2010).
23. Lie-Panis, J. & André, J.-B. Cooperation as a signal of time preferences. *Proceedings of the Royal Society B: Biological Sciences* **289**, 20212266. <https://royalsocietypublishing.org/doi/10.1098/rspb.2021.2266> (2022) (Apr. 27, 2022).
24. Barclay, P., Bliege Bird, R., Roberts, G. & Számádó, S. Cooperating to show that you care: costly helping as an honest signal of fitness interdependence. *Philosophical Transactions of the Royal Society B: Biological Sciences* **376**, 20200292. ISSN: 0962-8436, 1471-2970. <https://royalsocietypublishing.org/doi/10.1098/rstb.2020.0292> (2021) (Nov. 22, 2021).
25. Santos, F. P., Santos, F. C. & Pacheco, J. M. Social norm complexity and past reputations in the evolution of cooperation. *Nature* **555**, 242–245. ISSN: 0028-0836, 1476-4687. <https://www.nature.com/articles/nature25763> (2024) (Mar. 2018).
26. Kessinger, T. A., Tarnita, C. E. & Plotkin, J. B. Evolution of norms for judging social behavior. *Proceedings of the National Academy of Sciences* **120**, e2219480120. <https://www.pnas.org/doi/10.1073/pnas.2219480120> (2024) (June 13, 2023).
27. Michel-Mata, S. *et al.* The evolution of private reputations in information-abundant landscapes. *Nature*, 1–7. ISSN: 1476-4687. <https://www.nature.com/articles/s41586-024-07977-x> (2024) (Sept. 25, 2024).
28. Murase, Y. & Hilbe, C. Computational evolution of social norms in well-mixed and group-structured populations. *Proceedings of the National Academy of Sciences of the United States of America* **121**, e2406885121. ISSN: 1091-6490 (Aug. 13, 2024).
29. Colnaghi, M., Santos, F. P., Van Lange, P. A. M. & Balliet, D. Adaptations to infer fitness interdependence promote the evolution of cooperation. *Proceedings of the National Academy of Sciences* **120**, e2312242120. ISSN: 0027-8424, 1091-6490. <https://pnas.org/doi/10.1073/pnas.2312242120> (2024) (Dec. 12, 2023).
30. Pal, S., Hilbe, C. & Glynatsi, N. E. *The co-evolution of direct, indirect and generalized reciprocity* Nov. 5, 2024. arXiv: [2411.03488\[q-bio\]](https://arxiv.org/abs/2411.03488). <http://arxiv.org/abs/2411.03488> (2025).
31. Glynatsi, N. E., Akin, E., Nowak, M. A. & Hilbe, C. Conditional cooperation with longer memory. *Proceedings of the National Academy of Sciences* **121**, e2420125121. <https://www.pnas.org/doi/abs/10.1073/pnas.2420125121> (2024) (Dec. 10, 2024).

32. Kawakatsu, M., Kessinger, T. A. & Plotkin, J. B. A mechanistic model of gossip, reputations, and cooperation. *Proceedings of the National Academy of Sciences* **121**, e2400689121. <https://www.pnas.org/doi/10.1073/pnas.2400689121> (2024).
33. Hilbe, C., Schmid, L., Tkadlec, J. & Nowak, M. A. Indirect reciprocity with private, noisy, and incomplete information. *Proceedings of the National Academy of Sciences* **115**, 12241–12246. <https://www.pnas.org/doi/10.1073/pnas.1810565115> (2018).
34. Murase, Y. & Hilbe, C. Indirect reciprocity under opinion synchronization. *Proceedings of the National Academy of Sciences* **121**, e2418364121. <https://www.pnas.org/doi/10.1073/pnas.2418364121> (2024).
35. Pal, S. & Hilbe, C. Reputation effects drive the joint evolution of cooperation and social rewarding. *Nature Communications* **13**, 5928 (2022).
36. Hübner, V., Schmid, L., Hilbe, C. & Chatterjee, K. Stable strategies of direct and indirect reciprocity across all social dilemmas. *PNAS Nexus* **4**, pgaf154. ISSN: 2752-6542. <https://doi.org/10.1093/pnasnexus/pgaf154> (2025) (May 1, 2025).
37. Lange, P. *et al.* Willingness to Sacrifice in Close Relationships. *Journal of personality and social psychology* **72**, 1373–95 (June 1, 1997).
38. Sijlmasi, A., Mercier, H., Safra, L. & Chevallier, C. *Information about Immigrants' Deservingness Reduces Misperceptions and Opposition to Immigration* Dec. 13, 2024. https://osf.io/48xmv_v1 (2025).
39. Mailath, G. J. & Samuelson, L. *Repeated games and reputations* (Oxford Univ. Press, Oxford, UK, 2006).
40. Traulsen, A., Nowak, M. A. & Pacheco, J. M. Stochastic dynamics of invasion and fixation. *Physical Review E* **74**, 011909. <https://link.aps.org/doi/10.1103/PhysRevE.74.011909> (2025) (July 17, 2006).
41. Von Krogh, G., Haefliger, S., Spaeth, S. & Wallin, M. W. Carrots and Rainbows: Motivation and Social Practice in Open Source Software Development. *MIS Quarterly* **36**, 649–676. <https://www.jstor.org/stable/41703476> (2012).
42. Rusbult, C. E., Martz, J. M. & Agnew, C. R. The Investment Model Scale: Measuring commitment level, satisfaction level, quality of alternatives, and investment size. *Personal Relationships* **5**, 357–387. ISSN: 1475-6811. <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1475-6811.1998.tb00177.x> (2025) (1998).
43. Cohen, S., Schulz, M. S., Weiss, E. & Waldinger, R. J. Eye of the beholder: The individual and dyadic contributions of empathic accuracy and perceived empathic effort to relationship satisfaction. *Journal of Family Psychology* **26**, 236–245 (2012).
44. Cimino, A. The evolution of hazing: Motivational mechanisms and the abuse of newcomers. *Journal of Cognition and Culture* **11**, 241–267 (2011).
45. Gambetta, D. *Codes of the underworld: How criminals communicate* (Princeton University Press, Princeton, 2009).
46. Iannaccone, L. R. Sacrifice and stigma: Reducing free-riding in cults, communes, and other collectives. *Journal of Political Economy* **100**, 271–291 (1992).
47. Power, E. A. Discerning devotion: Testing the signaling theory of religion. *Evolution and Human Behavior* **38**, 82–91 (2017).
48. Sosis, R. & Bressler, E. R. Cooperation and commune longevity: A test of the costly signaling theory of religion. *Cross-Cultural Research* **37**, 211–239 (2003).
49. Alcorta, C. S. & Sosis, R. Ritual, emotion, and sacred symbols: The evolution of religion as an adaptive complex. *Human Nature* **16**, 323–359 (2005).
50. Turner, V. W. *The Ritual Process: Structure and Anti-Structure* (Aldine Publishing, Chicago, 1969).
51. Durkheim, É. *The Elementary Forms of Religious Life* (Alcan, Paris, 1912).
52. Xygalatas, D. *et al.* Extreme rituals promote prosociality. *Psychological Science* **24**, 1602–1605 (2013).
53. Fudenberg, D. & Tirole, J. *Game Theory* (MIT Press, Cambridge, MA, 1991).