Supporting Information

Delton et al. 10.1073/pnas.1102131108

SI Text

In the main text, we give examples showing that organisms undergoing nothing but a selective regime for direct reciprocity evolve to cooperate even in the presence of strong evidence that they are in one-shot interactions. Here, we detail the simulations used to generate these results, give a complete summary of the results, and develop an analytic model that complements the results. The main result to note is that, across very broad ranges of parameter space, cooperation in one-shot interactions is a robust outcome. It is also worth noting that the relatively small regions in which one-shot cooperation does not evolve heavily overlap with the regions where cooperation itself will not evolve, because repeat interactions are too few in relation to the benefits of cooperation for cooperation to be evolutionarily favored.

Simulation Details. To assess the evolutionary dynamics of one-shot encounters when interactions can either be one-shot or repeated, we conducted a series of agent-based simulations. Each simulation run consisted of a population of 500 agents allowed to evolve through 10,000 generations. For each generation, the population is randomly sorted, irrespective of genotype, into dyads. Dyads are then randomly assigned, irrespective of genotype, to interact in one-shot or indefinitely repeated prisoner's dilemmas (PDs). The base rate of one-shot dyads was varied across simulation runs at 10%, 30%, 50%, 70%, and 90%. These base rates for one-shot interactions were selected because they are high and, therefore, work against our hypothesis. Indeed, for our band-living hunter–gatherer ancestors, even 10% is likely an order of magnitude too high. These values serve as worst-case scenarios; the greater the base rate of one-shot interactions, the more decision architectures will be selected to respond adaptively to them, creating conditions more likely to favor defection.

Each member of a dyad independently (and randomly with respect to their own genotype) draws a cue summary from the appropriate distribution (Fig. 2). That is, each agent assigned to a one-shot interaction dyad independently draws a cue summary from the distribution for one-shot interactions, and each agent assigned to a repeated interaction dyad independently draws a cue summary from the distribution for repeated interactions.

In both types of dyads, the interaction consists of at least one round of interaction. If the interaction was one-shot, there was no additional interaction. If the interaction was indefinitely repeated, the two agents moved onto another round of interaction with probability w, with values of w varying across simulation runs from 0.5, 0.8, 0.9, 0.95 to 0.99. Whether a dyad moves from (e.g.) round 1 to round 2 is computed independently from whether it moves from (e.g.) round 2 to round 3. Thus, the average number of rounds for a repeated dyad is $1/(1-w)$, with values varying across simulation runs from 2, 5, 10, 20 to 100 average rounds of interaction in repeated dyads. Although we contrast indefinitely repeated interactions with one-shot interactions, because the latter are the simplest and most extreme cases, what is really at issue is the distinction between interactions with indefinite endpoints and interactions with finite endpoints. In the latter situation, the payoff-maximizing choice is for agents to defect in every round of interaction (1). Of all interactions with finite endpoints, one-shot interactions have the least potential to be profitable (e.g., they have the lowest potential for generating gains through mistaken cycles of mutual cooperation). Using one-shot encounters is, therefore, the most difficult test case for the evolution of cooperation in finite interactions.

During each round of their PD, each agent can give a benefit b to the other agent at a cost c to the self. There are gains in trade such that $b > c > 0$. For the simulations, we hold c constant at one, meaning b also represents the benefit to cost ratio. The fitness that an agent earns in each interaction is added to a constant baseline fitness. Because we are interested in the effects of selection, we set baseline fitness to 10 to ensure that, given the strategies and parameters used, no agent's fitness will ever be less than 1. The values of b varied across simulation runs from 1.5, 2, 3, 4, 5, 6, 7, 8, 9 to 10.

Based on their decision rules and the cue summary that they observe from their partner (both described below), agents either play a strategy of always defecting (ALLD) or the well-known TIT-for-TAT strategy (TFT) (2, 3). TFT cooperates on the first round of interaction and thereafter only cooperates when its partner cooperated in the previous round; otherwise, TFT defects. Note that our goal is to model a choice between cooperation and defection, not to determine exactly which cooperative strategy would perform best in the world modeled. There are many variants and alternatives to TFT, many of which outperform it. However, at least when playing themselves, all of these variants are fundamentally cooperative strategies. Compared with the alternatives, TFT has the additional benefit of being familiar to most readers. (Indeed, our own view is that the evolved human cooperative psychology plays a strategy or metastrategy that is far more complex than existing formalized cooperative strategies; our selection of a simple cooperative strategy is merely illustrative.)

After being sorted into one-shot or repeated dyads, agents are assigned a cue summary that is probabilistically associated with the type of dyad the agent belongs to. These cue summaries are drawn independently for both agents in a dyad and are drawn independently of the organisms' genotype (i.e., independently of their heritable strategy). In real life, a variety of cues could be probabilistically associated with whether the dyad is one-shot or repeated, such as prior history of interaction, community of residence, etc. As for many phenomena that involve summing multiple error-prone quantities, cue summaries are modeled as a single normal distribution.

Agents can see their partner's cue summary but cannot see their own. Each agent's cue summary is independently drawn from a normal distribution with a SD of 1 and a mean of $-d/2$ or $d/2$ for the one-shot and repeated dyads, respectively. Cue summaries discriminate between the two types of dyads probabilistically (rather than perfectly); that is, there is some overlap between the distribution of cue summaries associated with one-shot interactions and the distribution associated with repeated interactions. The ease of discrimination is defined by the parameter d (the distance between the means of the two distributions), which took on the values one, two, and three—corresponding to a proportional overlap between the distributions of 61.7%, 31.7%, and 13.4%, respectively.

We conducted two simulation sets, each with a different cooperative rule. One simulation set is designed to test the evolutionary dynamics of decision-making components concerned with representing states of the world (cognitive components). The other simulation set is designed to test the evolutionary dynamics of decision-making components that transform such representations into decisions about what actions to take (motivational components).

The first simulation set (SS1) investigates the evolutionary dynamics of motivational components. In these simulations, the agents form beliefs using Bayesian updating—allowing them to have the most accurate beliefs possible—but then, translate those beliefs into behavior using evolvable motivational variables. In this simulation set, it is assumed that the agents have innate and perfect knowledge of the actual base rates of one-shot and repeated interactions in their world. Based on Bayesian updating, if the posterior probability that the dyad is one-shot is largest, then the agent believes that the dyad is one-shot. The agent then generates a random number on the uniform interval [0, 1] and compares that with an internal regulatory variable that represents the probability of cooperating given that the agent believes that the dyad is one-shot, $Probability\ Cooperation_{One\text{-}Shot}$. If the random number is less than $Probability\ Cooperation_{One\text{-}Shot}$, the agent cooperates by playing TFT; otherwise, the agent defects. If the posterior probability that the dyad is repeated is largest, the agent goes through a similar algorithm. In this case, the agent consults an internal regulatory variable that represents the probability of cooperating by playing TFT given a belief that the interaction is repeated, $Probability\ Cooperation_{Repeated}$. Over generations, the magnitudes of these internal regulatory variables are allowed to evolve. To work against our hypothesis, at the start of each simulation run, we set these evolvable variables at values that maximize consistency with belief: $Probability\ Cooperation_{One\text{-}Shot} \approx 0$ and $Probability\ Cooperation_{Repeated} \approx 1$ (see paragraph on mutation below).

Our second simulation set (SS2) investigates the evolutionary dynamics of cognitive components. In these simulations, we fix an agent's actions to be consistent with its belief—cooperating (by playing TFT) in repeated interactions and, otherwise, defecting. However, beliefs are not computed with Bayesian updating. Instead, agents compare their partner's cue summary—their only evidence about whether the interaction is one-shot or repeated—with an evolvable internal threshold to determine their belief. This threshold can take on the same range of values as the cue summary: any positive or negative number. If the partner's cue summary is greater than the threshold, then the agent believes that the interaction will be repeated (and therefore, cooperates by playing TFT); otherwise, the agent believes that the interaction is one-shot (and therefore, defects). The value of the threshold reflects an agent's default assumption about its interactions. As agents evolve increasingly positive thresholds, they demand more and more evidence that an interaction will be repeated before they believe that it is repeated and, therefore, cooperate (i.e., they require increasingly higher cue summaries). Conversely, as agents evolve increasingly negative thresholds, they become more credulous; they require less evidence before believing that they are having a repeated interaction (i.e., even low cue summaries will trigger the belief that the interaction is repeated). Again, to advantage the alternative hypothesis, we start the first generation with an average threshold that maximizes accuracy (see paragraph on mutation below).

For every simulation run, the population evolves through 10,000 generations. The probability that an agent in the current generation is the parent of a given member of the next generation is proportional to the agent's fitness. Specifically, the probability that an agent was the parent of a given member of the next generation was the ratio of its fitness to the total fitness of its entire generation. Algorithmically, agents were assigned intervals of probability space based on their proportion of the population fitness. Offspring were created by randomly sampling from this probability space until a complete successor generation was filled. Fitness only probabilistically influenced an agent's number of offspring, allowing the possibility of genetic drift. Offspring inherited their parent's genotype with a 5% probability of mutation, and, when mutations occurred, they modified the parent's genotype by a normally distributed random variable with a mean = 0 and SD = 0.025. For the first simulation set, mutations are independent for the two regulatory variables. However, because they represent probabilities, mutations are not allowed to move these values to greater than one or less than zero. To generate variance on which natural selection can act, in the initial generation of both simulation types, all agents are run through the mutation procedure described above, except that there is a 100% chance that a mutation takes place.

We factorially explored the parameter space of b, w, d, and the base rate of one-shot dyads; these parameters had 10, 5, 3, and 5 possible values, respectively. Four independent simulation runs were conducted for each instantiation of the parameters. Therefore, for each simulation type, there were $10 \times 5 \times 3 \times 5 \times 4 = 3{,}000$ runs. There were two types of decision rules, creating 6,000 runs. We then replicated this procedure introducing errors in agents' behavior (see below) for a total of 12,000 simulation runs.

Theoretical work has shown errors to have an important impact on the evolutionary dynamics of games by exposing hidden weaknesses of strategies, potentially destabilizing and leading to the extinction of nonrobust strategies (4). Therefore, we conducted a second set of simulations identical to those simulations described above, except that we incorporated errors. When an error occurred, an agent who intended to cooperate instead defected. The reverse was not true; agents never mistakenly cooperated when they intended to defect. The error rate was set at a relatively high proportion: 0.05. In other words, in 1 of 20 instances where an agent intends to cooperate, it defects instead. In these simulations, agents use GRIM, a variant of TFT. GRIM cooperates until it observes that its partner defected. After GRIM defects, it defects forever. When there are no errors, GRIM and TFT behave identically when playing themselves or ALLD. When there are errors, GRIM fails to effectively cooperate in situations where TFT would, thus lowering the average benefit of playing a cooperative strategy. Because our hypothesis depends on successful cycles of cooperation causing large benefits, using GRIM works against our hypothesis.

Simulation Results. The full results of the main simulations (i.e., those without errors) are presented in Figs. S1 and S3. These figures show how selection exclusively based on payoffs to direct reciprocity crafts organisms to be "irrationally" cooperative. As predicted, cooperation in one-shot encounters becomes increasingly likely as the within-round benefits of cooperation increase and the expected length of interactions increases. The inclusion of errors does not drastically change the evolutionary dynamics, although it does mitigate the magnitude of the effects (Figs. S2 and S4). In all cases, the results were computed by averaging the final 500 generations of the individual simulations. Because the evolvable thresholds from the second simulation set are not in an easily interpretable metric, for Figs. S3 and S4, we converted them to the expected probability that an agent with that threshold would believe the interaction is repeated and, therefore, would cooperate if it was in fact facing a one-shot interaction. This conversion is accomplished by simply determining the value at the threshold of a cumulative distribution function for a normal distribution with a mean $= -d/2$ and an SD = 1. Additionally, the main data were analyzed as multiple regressions to provide quantified estimates of the main effects (Tables S1 and S2).

As one might expect, greater ease of discrimination (i.e., higher values of d) selected for agents who were less willing to cooperate in one-shot interactions. However, the difficulty of discriminating one-shot from repeated dyads did not generally have a large effect on the evolution of regulatory variables: the amount of variance explained by d was 1.1% and 0.5% in SS1 and SS2, respectively, and was far less than that explained by each of the other variables, which ranged from 2% to 33% (because the four predictors are perfectly orthogonal, percentage of variance explained can be calculated by squaring the β of each variable). Thus, our summary graphs average over this parameter.

Finally, note that there are, nevertheless, some regions of parameter space where one-shot cooperation does not evolve. However, these regions are closely associated with areas where cooperation in repeated interactions will not evolve. Reciprocity will be selected against unless $w \times b > c$. Hence, even in a world where interactions are guaranteed the chance of being repeated, if the average length of repeated interactions is 2 ($w = 0.5$) and benefits are <2, then cooperation is not favored. Therefore, it is to be expected that these conditions, which select against cooperation in repeated interactions, also select against one-shot cooperation.

Strengths and limitations of each model. Each simulation set has different strengths and limitations. Many economists prefer models in which preferences (motivations) vary, but beliefs are formed rationally, as in SS1. By contrast, many cognitive psychologists doubt that beliefs are rationally formed using Bayesian updating and point to heuristics and biases that affect the judgments people make—a view more consistent with SS2. Evidence consistent with either model can be found depending on which area of the literature that a reader consults.

For example, some economists may view the evidence from experimental games as more consistent with SS1, given that (*i*) the experimenter informs subjects that the interaction was one-shot (or repeated, as the case may be) and (*ii*) most subjects verbally report that they believe what the experimenter told them (e.g., 5, 6). For the sake of argument, let us assume that these subjects are trying to report their beliefs truthfully (and not just being polite by telling the experimenter that they believe what he or she told them). These verbal reports, albeit face valid, would not settle the SS1 vs. SS2 issue for most psychologists because research in cognitive neuroscience (e.g., split brain patients), psychology (social, cognitive, developmental, and evolutionary), and even behavioral economics shows that behavior can be regulated by nonconscious or implicit variables that reflect assessments of the situation (rather than preferences) and that the content of verbal reports is often uncorrelated with these variables.

A verbal report that "X is true" can be associated with different degrees of confidence. What you say is true can be different from your gut feeling about what is true (a decision threshold), with consequential behaviors regulated by the decision threshold (as in SS2) rather than by a verbal statement (that was elicited by an experimenter's query). Indeed, given that speaking evolved in the context of navigating a complex social world, decisions about what to say, admit to, or advocate may be generated by different cognitive systems than decisions about what actions to take.

Psychologists—and many economists—would also point out that situational cues relevant to epistemic judgments about facts of the matter (is this situation one-shot? is it really anonymous?) importantly affect behavior in experimental games, even when these cues conflict with what the experimenter told the subjects (and with what the subjects report that they were told). For example, economists often find that verbal assurances of anonymity are not sufficient to elicit more selfish decisions in economic games, yet, generosity does decrease when experimenters create a transparent situation that guarantees anonymity (7) and it increases in the presence of stylized eyes—an ancestrally reliable cue that one is being observed (8). These situational cues could exert their effect by Bayesian updating (as in SS1) or by shifting the weight of evidence above a decision threshold (as in SS2). Whichever is true, the subject might report that they believe the situation to be anonymous—not because they are lying or being overly polite but because the system that generates these verbal judgments does not have any more access to the computations that regulate cooperative behavior than it does to the retinal smears and low-level line detectors that build our conscious (and verbally reportable) representation of the visual world.

So, should one prefer SS1 or SS2? Our own view is that natural selection has probably tinkered with both motivational and cognitive components of the architecture that regulate cooperative behavior. But, rather than prejudge the issue, we chose to model each component separately while holding the other constant. The results show that a disposition for generosity evolves in either case; future research will have to determine the precise design of the computational systems that actually evolved.

Analytic Model. Here, we develop a simple, nonevolutionary, best-response model showing that cooperating even in the face of an explicit one-shot belief is always the payoff-maximizing choice when the long-term net benefits of cooperation are large enough. This analytic model makes the same assumptions as the model of the main paper:

i) Interactions are structured as PDs such that, on each round of an interaction, each member of a dyad can choose whether to transfer a benefit b to their partner at a personal cost c. There are gains in trade (i.e., $b > c$).

ii) Some interactions are one-shot; this type of interaction occurs with probability P. Other interactions are repeated; this type of interaction occurs with probability $1 - P$.

iii) However, organisms can only probabilistically determine whether their own interaction is one-shot or repeated by observing their partner's cue summary. Cue summaries are drawn from two standard normal distributions that differ in mean value depending on whether the interaction is one-shot or repeated.

iv) Organisms have perfect, innate knowledge of the cue summary distributions, including their base rates (symbolized by P) and their relative distances from each other.

v) Given their partner's cue summary and this knowledge of the cue summary distributions, organisms use Bayesian updating to compute posterior probabilities of their current interaction being one-shot or repeated.

We can avoid representing the process of Bayesian updating explicitly and thus drastically simplify the model. Given that the parameters underlying the distribution of the cue summaries are fixed, then there are also fixed probabilities that an organism will come to believe its interaction is one-shot. We denote the probability that an organism comes to believe its interaction is one-shot when the interaction is, in fact, one-shot as Θ_{OS}, and we denote the probability that an organism comes to believe its interaction is one-shot when the interaction is, in fact, repeated as Θ_R. The complements of these probabilities denote the probabilities that an organism comes to believe its interaction is repeated. If an organism believes its interaction is one-shot, it cooperates with probability α_i, where the subscript i (equal to one or two) denotes which member of the dyad this value belongs to. If an organism believes that its interaction is repeated, it cooperates with probability β_i. These final two variables were previously referred to as $Probability\ Cooperation_{One\text{-}Shot}$ and $Probability\ Cooperation_{Repeated}$. However, given that here they are embedded within equations, we use a more concise notation.

Given this specification, we derive payoff functions for the organisms. As an example of part of the payoff function, the expression $(1 - P)\ (1 - \Theta_R)\ \beta_1\Theta_R\alpha_2((b - c)/(1 - w))$ represents, going from left to right, the probability that the interaction is in fact repeated, multiplied by the probability that player 1 believes the interaction is repeated (given that it is in fact repeated), multiplied by the probability that player 1 cooperates given a repeated belief, multiplied by the probability that player 2 believes the interaction is one-shot (given that it is actually repeated), multiplied by the probability player 2 cooperates given a one-shot belief, multiplied by the final quantity enclosed in parentheses, which represent the net benefits of sustained mutual cooperation. The full payoff function for player 1, V_1, is given as Eq. **S1** (because players are symmetrical, this equation also applies to player 2, with appropriate subscript substitutions).

$$V_1 = P\left[\begin{array}{l}\Theta_{OS}\left[\begin{array}{l}\Theta_{OS}(\alpha_1\alpha_2(b-c)-\alpha_1(1-\alpha_2)\,c+(1-\alpha_1)\,\alpha_2 b)\\ +(1-\Theta_{OS})(\alpha_1\beta_2(b-c)-\alpha_1(1-\beta_2)\,c+(1-\alpha_1)\beta_2 b)\end{array}\right]\\ +(1-\Theta_{OS})\left[\begin{array}{l}\Theta_{OS}(\beta_1\alpha_2(b-c)-\beta_1(1-\alpha_2)\,c+(1-\beta_1)\alpha_2 b)\\ +(1-\Theta_{OS})(\beta_1\beta_2(b-c)-\beta_1(1-\beta_2)c+(1-\beta_1)\beta_2 b)\end{array}\right]\end{array}\right]+(1-P)\left[\begin{array}{l}\Theta_R\left[\begin{array}{l}\Theta_R\left(\alpha_1\alpha_2\dfrac{b-c}{1-w}-\alpha_1(1-\alpha_2)\,c+(1-\alpha_1)\,\alpha_2 b\right)\\ +(1-\Theta_R)\left(\alpha_1\beta_2\dfrac{b-c}{1-w}-\alpha_1(1-\beta_2)\,c+(1-\alpha_1)\beta_2 b\right)\end{array}\right]\\ +(1-\Theta_R)\left[\begin{array}{l}\Theta_R\left(\beta_1\alpha_2\dfrac{b-c}{1-w}-\beta_1(1-\alpha_2)\,c+(1-\beta_1)\alpha_2 b\right)\\ +(1-\Theta_R)\left(\beta_1\beta_2\dfrac{b-c}{1-w}-\beta_1(1-\beta_2)\,c+(1-\beta_1)\beta_2 b\right)\end{array}\right]\end{array}\right] \quad \textbf{[S1]}$$

To find the best response of player 1, we take the derivative of this function with respect to α_1. This quantity, with some algebraic rearrangement, is shown as Eq. **S2**.

$$\frac{b-c}{(1/w)-1}(1-P)\Theta_R(\Theta_R\alpha_2+(1-\Theta_R)\beta_2)-c(P\Theta_{OS}+(1-P)\Theta_R) \quad \textbf{[S2]}$$

This expression no longer contains any terms involving α_1. Therefore, we need simply to determine whether the total quantity is greater than or less than zero. If the derivative is less than zero, then the best response is to never cooperate with a one-shot belief. However, if the derivative is greater than zero, then the best response is to always cooperate with a one-shot belief. The total left-hand quantity represents the net payoffs of repeated cooperation (albeit discounted by an additional factor of w) multiplied by the probability that the focal agent's partner cooperates when the interaction is in fact repeated and the probability that the focal agent believes the interaction is one-shot. The total right-hand quantity represents the cost of cooperation when the partner does not cooperate multiplied by the probability of an agent believing that the interaction is one-shot. All parameters denoting probability (with the exception of α_2) can reasonably be assumed to be strictly greater than zero: We assume that either type of interaction can actually occur. Under this assumption, the probabilities of believing that an interaction is one-shot are necessarily strictly greater than zero. Given that, outside a narrow range of values cooperation in repeated interactions is favored, the probability of cooperating given a belief that the interaction is repeated should also be strictly greater than zero, possibly near one. At the very least, given the assumptions of the model, the range of parameters leading to nonzero probabilities of cooperation are less restrictive for β than for α. Given these assumptions, the total expression will be greater than zero as long as the benefits of cooperation are sufficiently large compared with the costs and as long as repeated interactions are sufficiently long.

This analysis confirms the reasoning and simulations already presented: When there are sizeable gains in trade to be made through mutually beneficial, long-lasting exchange, organisms should cooperate in one-shot encounters even when they have explicit beliefs that cooperation is one-shot.

1. Gibbons R (1992) *Game Theory for Applied Economists* (Princeton University Press, Princeton).
2. Axelrod R, Hamilton WD (1981) The evolution of cooperation. *Science* 211:1390–1396.
3. Nowak MA (2006) *Evolutionary Dynamics: Exploring the Equations of Life* (Belknap Press, Boston).
4. Panchanathan K, Boyd R (2003) A tale of two defectors: The importance of standing for evolution of indirect reciprocity. *J Theor Biol* 224:115–126.
5. Fehr E, Fischbacher U (2003) The nature of human altruism. *Nature* 425:785–791.
6. Fehr E, Henrich J (2003) Is strong reciprocity a maladaptation? On the evolutionary foundations of human altruism. In *Genetic and Cultural Evolution of Cooperation*, ed Hammerstein P (MIT Press, Cambridge, MA), pp 55–82.
7. Hoffman E, McCabe K, Shachat K, Smith V (1994) Preferences, property rights, and anonymity in bargaining games. *Games Econ Behav* 7:346–380.
8. Haley KJ, Fessler DMT (2005) Nobody's watching? Subtle cues affect generosity in an anonymous economic game. *Evol Hum Behav* 26:245–256.

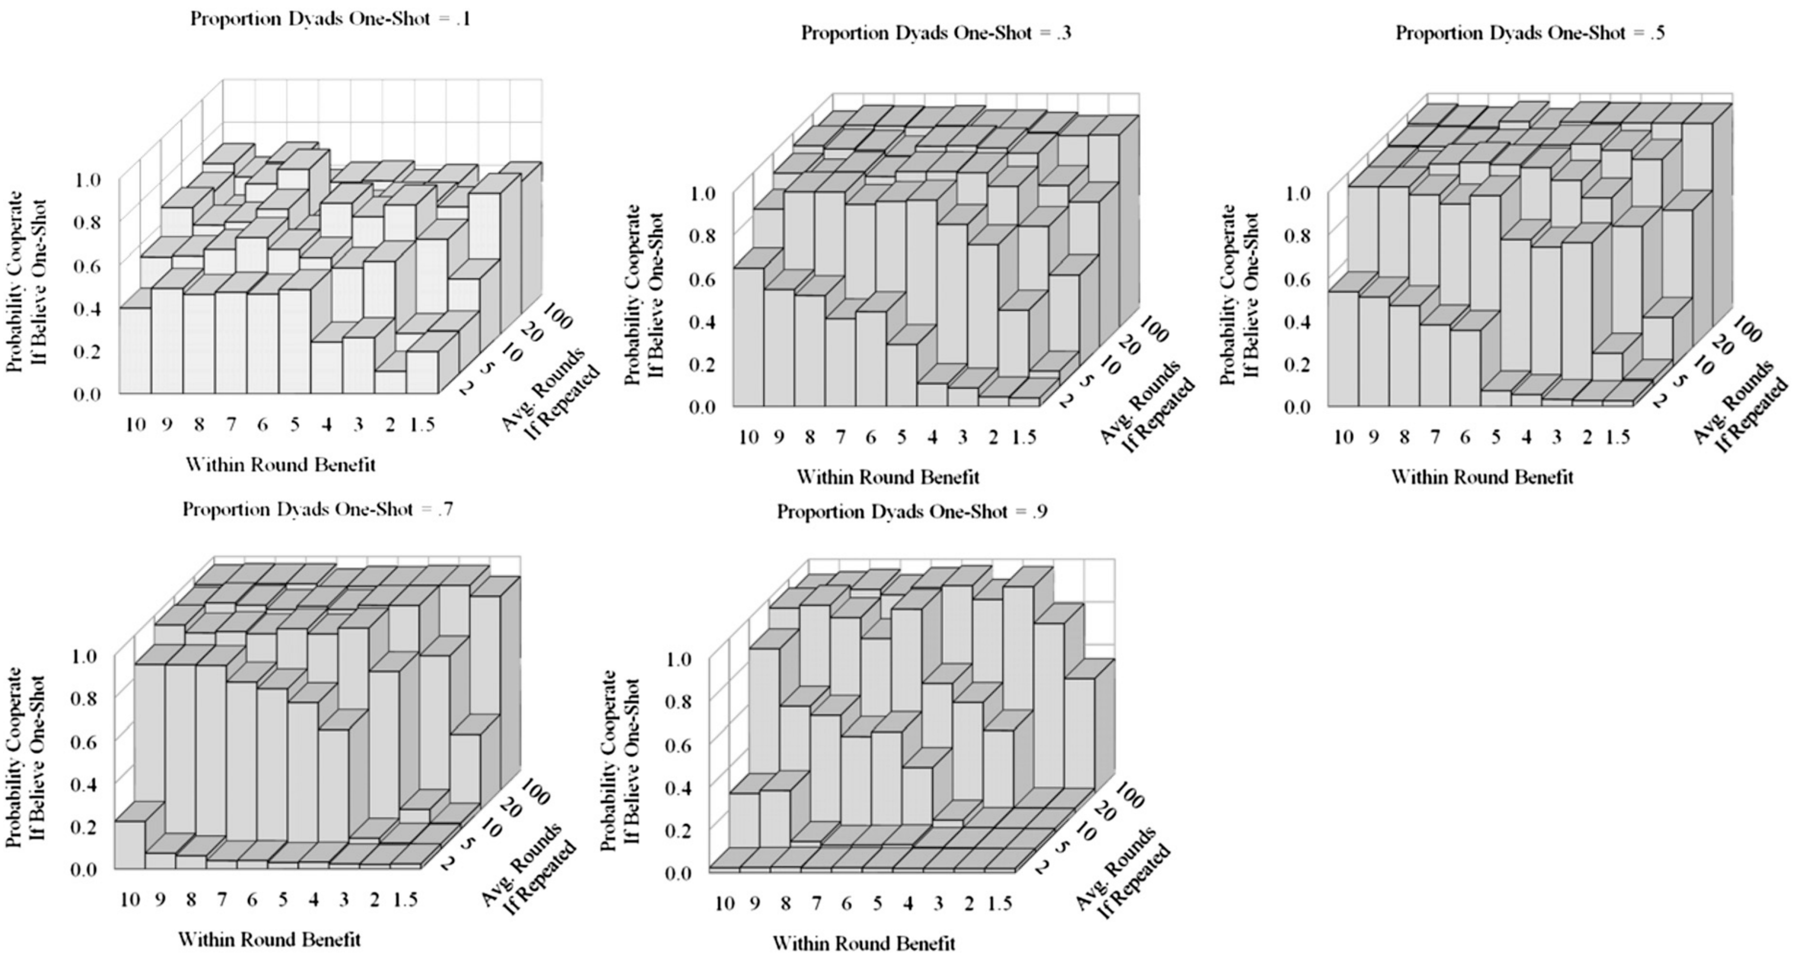

Fig. S1. The evolution of motivation (SS1; no errors possible). If repeated interactions last on average five or more rounds—surely a lower bound for human cooperative relationships that may span decades—even modest returns to cooperation generally cause agents to be very willing to cooperate when they believe that their interaction will be one-shot (SS1 simulations, using Bayesian updating). Graphs display final values of *Cooperation Probability*$_{One\text{-}Shot}$ when errors are not possible. Values are averaged over the final 500 generations within a simulation run and over distances between cue summary distributions (*d*).

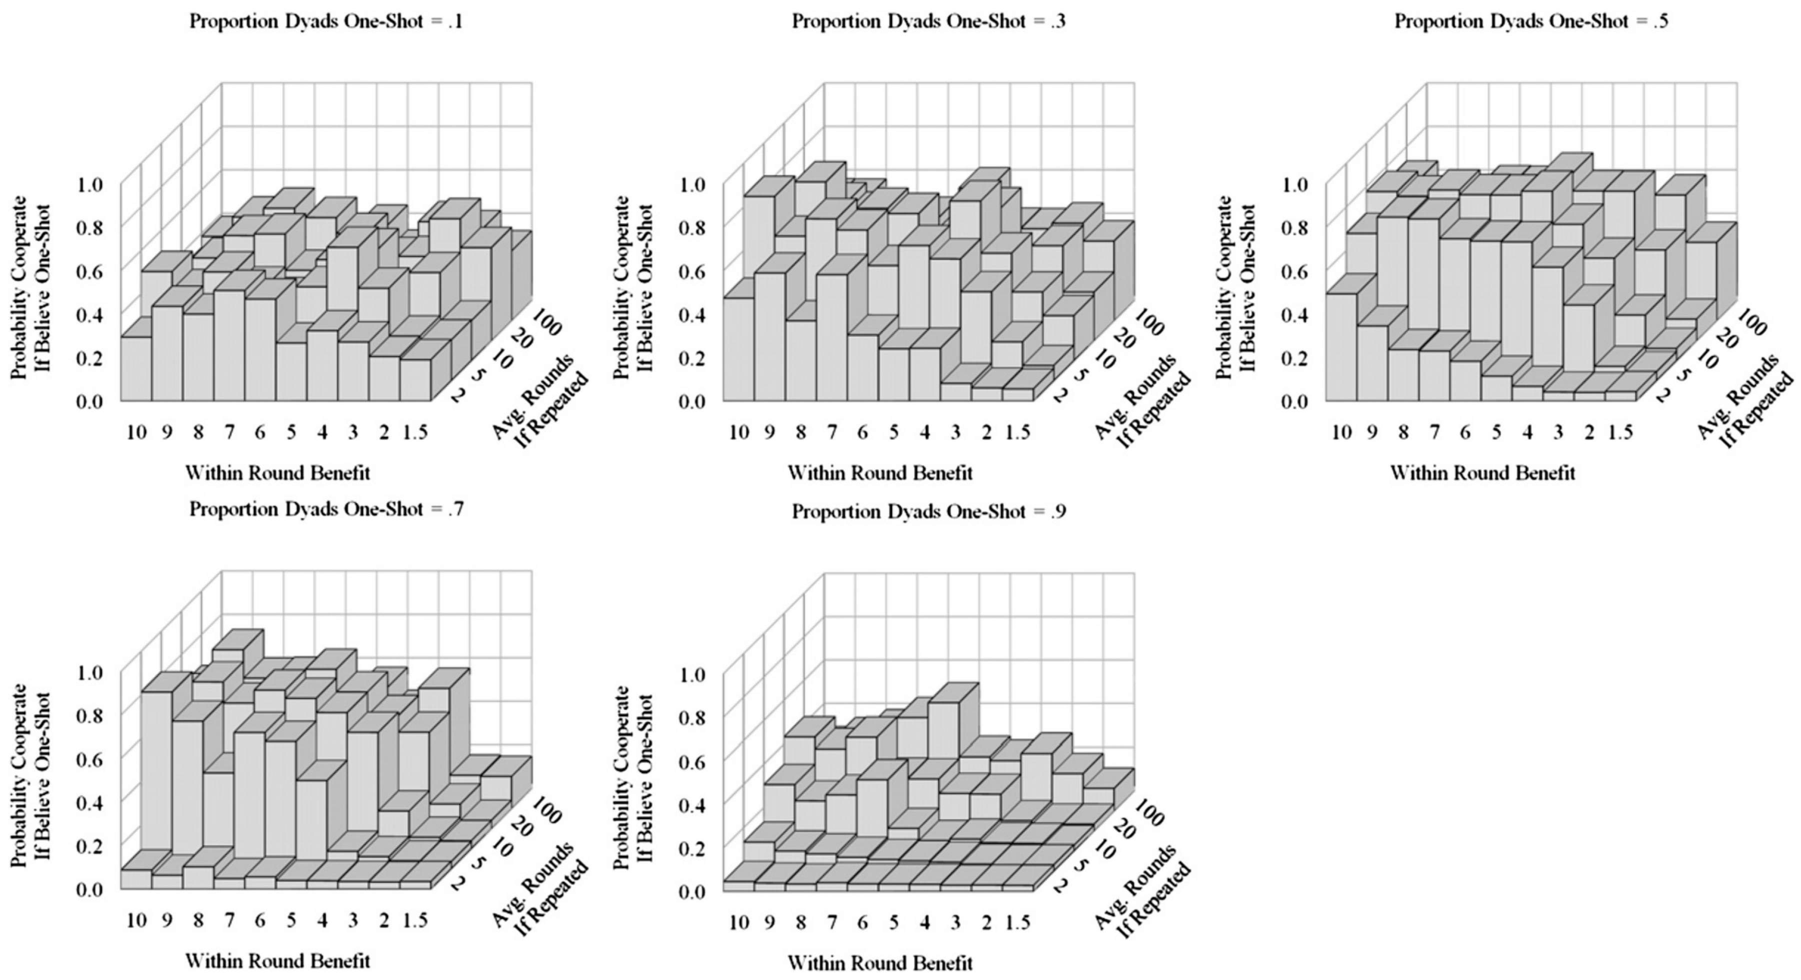

Fig. S2. The evolution of motivation (SS1; errors possible). Even with the possibility of errors, if repeated interactions last on average five or more rounds, agents still evolve to be willing to cooperate with one-shot partners (SS1 simulations, using Bayesian updating). Cooperation only seems substantially reduced by the inclusion of errors when 90% of interactions are one-shot, a situation unlikely to have characterized ancestral human environments. Graphs display final values of *Cooperation Probability*$_{One\text{-}Shot}$ when errors are possible. Values are averaged over the final 500 generations within a simulation run and over distances between cue summary distributions (*d*).

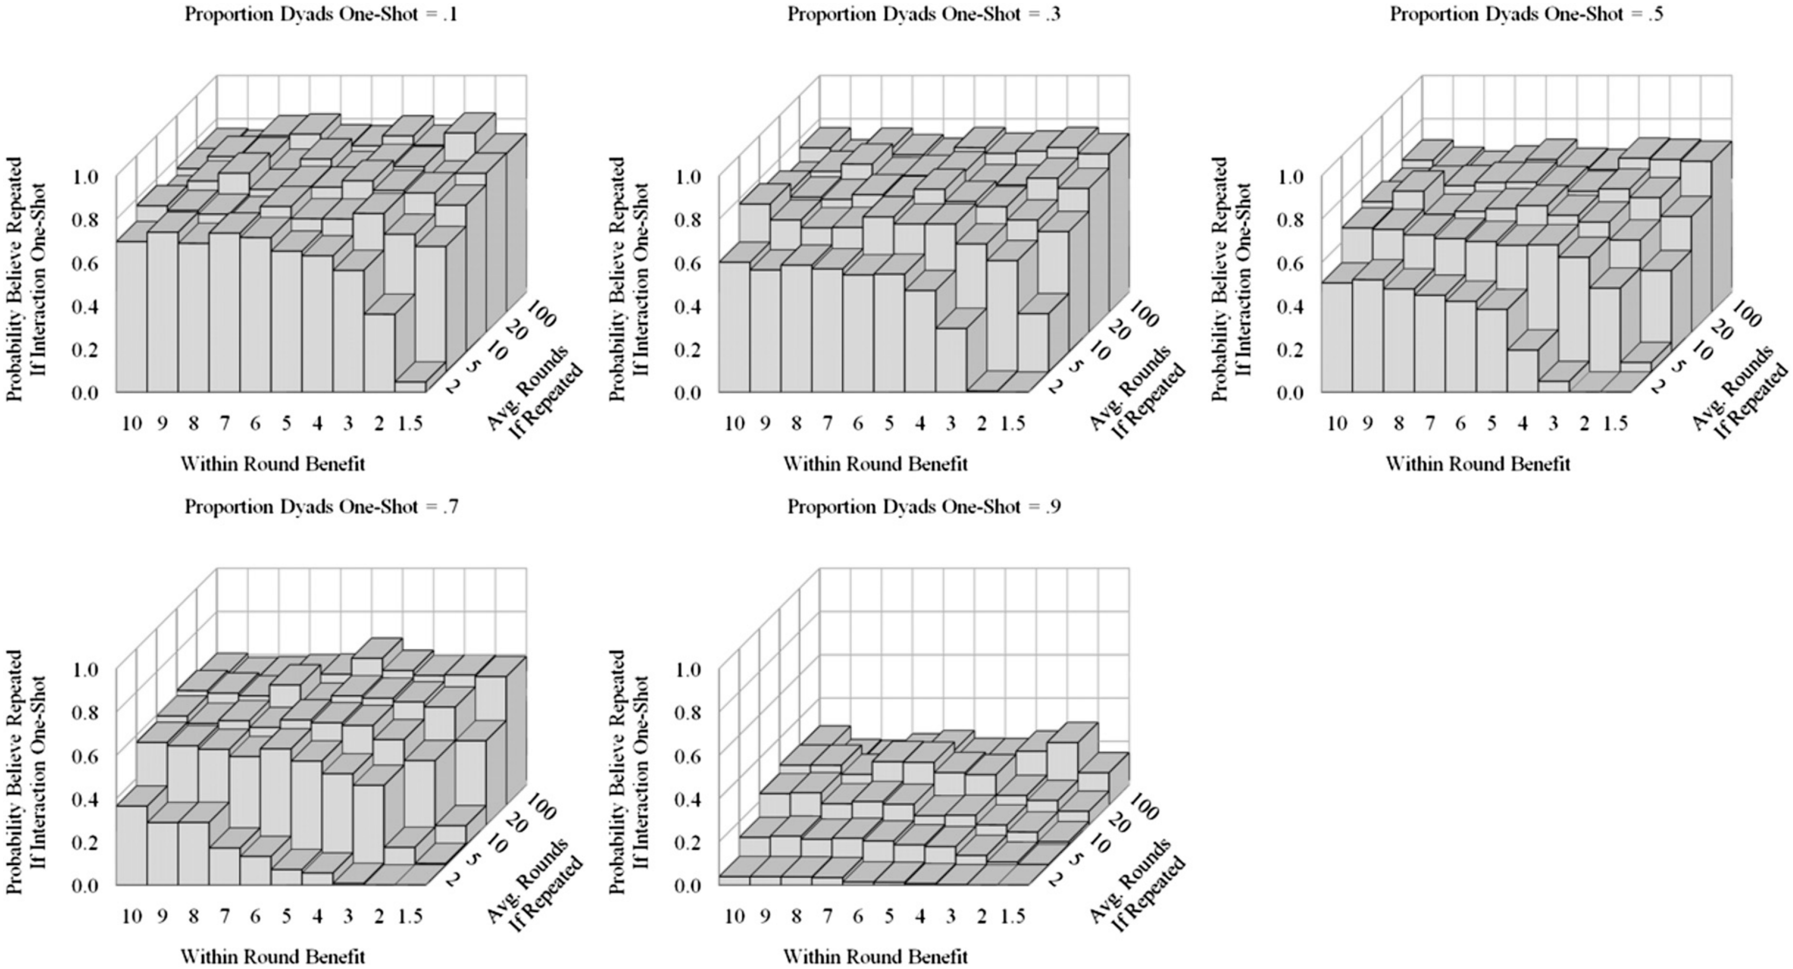

Fig. S3. The evolution of cognition (SS2; no errors possible). If repeated interactions last on average five or more rounds—likely a lower bound for human cooperative relationships—even modest returns to cooperation generally cause agents to require a great deal of evidence to believe that their interactions are one-shot (SS2 simulations, without Bayesian updating). Given these high thresholds, agents often believe their one-shot interactions are repeated and choose to cooperate. Graphs display the expected probability that agents will believe that their interactions are repeated when they are in fact one-shot. Values are averaged over the final 500 generations within a simulation run and over distances between cue summary distributions (*d*).

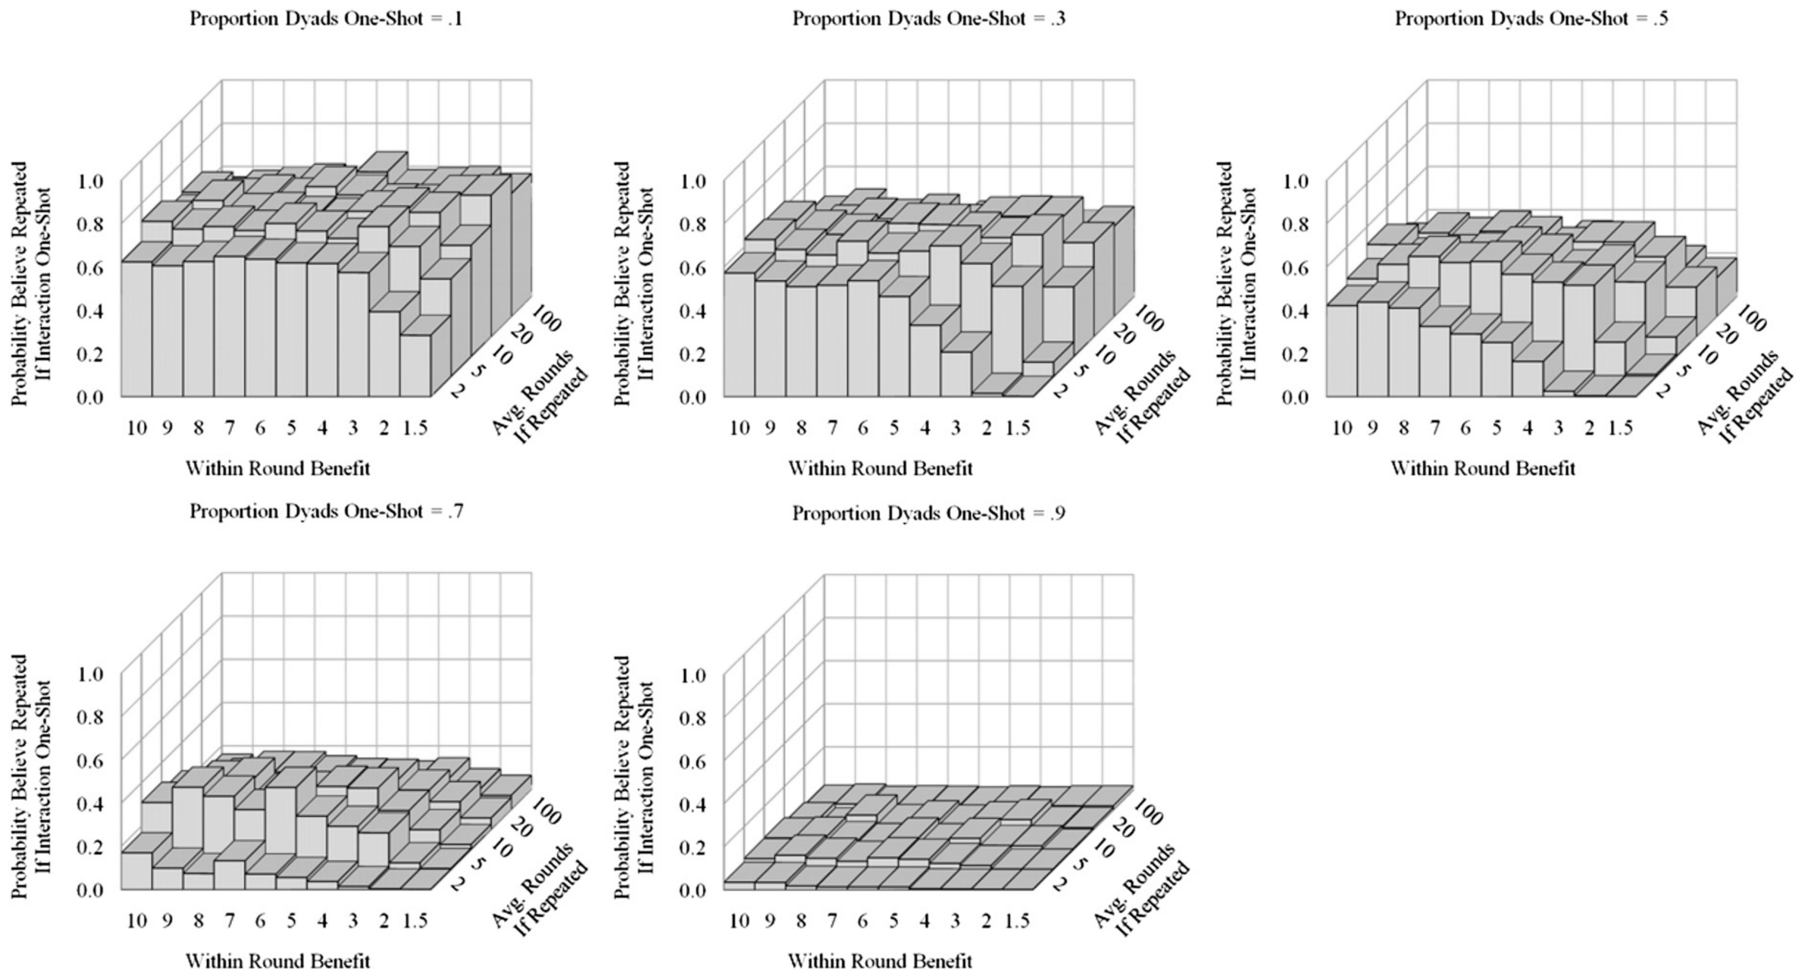

Fig. S4. The evolution of cognition (SS2; errors possible). Even with the possibility of errors, if repeated interactions last on average five or more rounds, agents still evolve to require a great deal of evidence to believe that their interactions are one-shot (SS2 simulations, without Bayesian updating). With errors present, agents generally still believe that their one-shot interactions are repeated and choose to cooperate. Cooperation only seems substantially reduced by the inclusion of errors when 90% of interactions are one-shot, a situation unlikely to have characterized ancestral human environments. Graphs display the expected probability that agents will believe that their interactions are repeated when they in fact are one-shot. Values are averaged over the final 500 generations within a simulation run and over distances between cue summary distributions (*d*).

Table S1. The evolution of motivation (SS1) regression equation (adjusted model $R^2 = 0.456$)

Predictor	*b*	SE	β	*t*	*p*
Intercept	−0.532	0.032		−16.370	$<10^{-57}$
Within-round benefit (*b*)	0.044	0.002	0.307	22.802	$<10^{-105}$
Probability of repeated interactions continuing (*w*)	1.315	0.031	0.577	42.847	$<10^{-300}$
Distance between cue summary distributions (*d*)	−0.051	0.007	−0.104	−7.693	$<10^{-13}$
Proportion of dyads that were one-shot	−0.197	0.019	−0.139	−10.323	$<10^{-23}$

Table S2. The evolution of cognition (SS2) regression equation (adjusted model $R^2 = 0.573$)

Predictor	*b*	SE	β	*t*	*p*
Intercept	−1.032	0.103		−9.980	$<10^{-22}$
Within-round benefit (*b*)	0.138	0.006	0.268	22.443	$<10^{-102}$
Probability of repeated interactions continuing (*w*)	3.418	0.098	0.418	35.006	$<10^{-224}$
Distance between cue summary distributions (*d*)	−0.124	0.021	−0.070	−5.899	$<10^{-8}$
Proportion of dyads that were one-shot	−2.890	0.061	−0.568	−47.591	$<10^{-300}$