PROCEEDINGS B

rspb.royalsocietypublishing.org

Research

Cite this article: Bear A, Kagan A, Rand DG. 2017 Co-evolution of cooperation and cognition: the impact of imperfect deliberation and context-sensitive intuition. *Proc. R. Soc. B* **284**: 20162326.
http://dx.doi.org/10.1098/rspb.2016.2326

Received: 1 November 2016
Accepted: 14 February 2017

Subject Category:
Evolution

Subject Areas:
evolution, behaviour

Keywords:
cooperation, evolutionary game theory, dual-process

Author for correspondence:
Adam Bear
e-mail: adam.bear@yale.edu

Electronic supplementary material is available online at https://dx.doi.org/10.6084/m9.figshare.c.3703627.

THE ROYAL SOCIETY PUBLISHING

Co-evolution of cooperation and cognition: the impact of imperfect deliberation and context-sensitive intuition

Adam Bear[1], Ari Kagan[1] and David G. Rand[1,2,3]

[1]Psychology Department, [2]Economics Department, and [3]School of Management, Yale University, New Haven, CT 06511, USA

AB, 0000-0002-8714-9144

How does cognitive sophistication impact cooperation? We explore this question using a model of the co-evolution of cooperation and cognition. In our model, agents confront social dilemmas and coordination games, and make decisions using intuition or deliberation. Intuition is automatic and effortless, but relatively (although not necessarily completely) insensitive to context. Deliberation, conversely, is costly but relatively (although not necessarily perfectly) sensitive to context. We find that regardless of the sensitivity of intuition and imperfection of deliberation, deliberating undermines cooperation in social dilemmas, whereas deliberating can increase cooperation in coordination games if intuition is sufficiently sensitive. Furthermore, when coordination games are sufficiently likely, selection favours a strategy whose intuitive response *ignores* the contextual cues available and cooperates across contexts. Thus, we see how simple cognition can arise from active selection for simplicity, rather than just be forced to be simple due to cognitive constraints. Finally, we find that when deliberation is imperfect, the favoured strategy increases cooperation in social dilemmas (as a result of reducing deliberation) as the benefit of cooperation to the recipient increases.

1. Introduction

Cooperation is a pervasive feature of human behaviour, yet its adaptive origins pose an evolutionary puzzle. Why would 'selfish' evolution lead people to pay personal costs to benefit non-kin? A wide variety of game-theoretic research over the past several decades has provided numerous answers to this question [1–5], many of which involve demonstrating how cooperating can actually be in one's long-term self-interest because of mechanisms such as direct or indirect reciprocity and signalling (i.e. cooperation can be 'strategic') [6–15].

Nevertheless, empirical study of humans' actual behaviour consistently finds that people cooperate even in contexts when these mechanisms are absent and cooperation is not in one's material self-interest (e.g. in social dilemmas where the interaction is anonymous and one-shot; for a review, see [16]). To explain 'pure' cooperation of this sort, recent research has turned to the proximate psychological mechanisms underlying cooperative behaviour. One prominent strain of this research has adopted a dual-process framework, distinguishing intuitive and deliberative cognitive processes [17,18]. Relatively speaking, intuitive processes are conceptualized as cognitively cheap, fast and automatic, but inflexible (i.e. insensitive to details of one's particular circumstances); whereas deliberative processes are conceptualized as cognitively expensive, slow and controlled, but flexible (i.e. sensitive to these details) [19–21].

Viewed through this dual-process lens, people's 'irrational' cooperative behaviour in social dilemmas can be explained as resulting from a 'social heuristic'—an intuitive (and, therefore, inflexible) judgement that typically does what is adaptive, but occasionally makes errors [17,22–25]. If cooperation is *typically* long-run payoff maximizing, because of mechanisms such as reciprocity and signalling, it can be advantageous to (sometimes) avoid the cost of deliberating about the specific situation one is facing and instead just intuitively cooperate. Importantly, this cost need not to be intrinsic to deliberation itself (as stipulated, e.g. by the theory of ego depletion [26]), but can also be operationalized as an opportunity cost that arises from the extra time and attentional demands that come from deliberating, which limit one's ability to engage in other potentially rewarding tasks (see [27]).

This argument has been articulated in a specifically dual-process way by the 'Social Heuristics Hypothesis' ((SHH) [22]; formalized in [25]). The SHH predicts that people should be more cooperative in social dilemmas when they are induced to rely more on intuitive cognitive processing relative to deliberative cognitive processing—a prediction that has garnered considerable empirical support, as shown by a recent large-scale meta-analysis [28]. Moreover, in further support of the SHH, this meta-analysis found no effect of intuitive versus deliberative cognitive processing on cooperation in coordination contexts in which cooperation is payoff-maximizing (e.g. repeated interactions) and thus it could be in one's self-interest to deliberatively cooperate.

Implicit in this theory, however, is an extreme dichotomy between intuitive and deliberative cognition: deliberation is assumed to be perfectly flexible, whereas intuition is assumed to be completely insensitive to the strategic nature of the situation one is facing [22,25]. While this extreme schema may apply to many situations in the real world, there is good reason to believe that it is too extreme in other cases.

With regard to deliberation, there are two primary ways in which this type of cognitive process could sometimes make mistakes and fail to accurately tailor its behaviour to context. First, sometimes the information available to an individual is simply limited. For example, it is sometimes unknown whether you will see a person again and therefore suffer consequences for defecting on them. Thus, even if a person carefully deliberates and analyses all of the information she has available to her in this situation, she may nevertheless come to the wrong conclusion about whether it is worth cooperating. Similarly, a person might get confused while deliberating and reach the wrong conclusion despite having full information available (e.g. experimental participants who mistakenly conclude that mutual cooperation is the highest payoff outcome in a social dilemma [29]). Second, deliberation itself sometimes merely helps people rationalize their intuitive desires, in which case deliberation will merely commit the same mistakes that intuition would [30].

With regard to intuition, an entire research programme separate from the heuristics and biases programme has emphasized how heuristics can be 'fast and frugal,' picking up on salient cues which can inform decision-making [31,32]. In other words, in many settings, intuition may prove quite flexible and accurate, even sometimes surpassing the capacities of deliberation. This may especially be true in domains in which a person has had considerable prior experience or in which there are learnable statistical regularities that distinguish one context from another (e.g. whether or not a person looks like they are from one's in-group).

In this paper, we use a formal evolutionary game-theoretic model to explore the consequences of imperfect deliberation and flexible intuition on the coevolution of cooperation and cognition. We find that adding imperfect deliberation has little qualitative effect on the evolutionary dynamics. Allowing intuition to be somewhat sensitive, on the other hand, can lead natural selection to favour a qualitatively different type of strategy—one for which deliberation increases cooperation in settings where cooperating can be payoff-maximizing. These findings extend our understanding of the cognitive underpinnings of human decision-making, and generate new, testable empirical predictions.

2. Model

We begin with a basic model framework (adapted from [25]) in which agents are placed in an environment consisting of a mix of social dilemmas (which occur with probability $1 - p$) and cooperative coordination games (which occur with probability p).

In the social dilemmas, agents can cooperate by paying a cost c to give a benefit b to their partner, or they can defect by doing nothing. Thus, in this type of situation, it is always payoff maximizing to defect. By contrast, it can be payoff-maximizing to cooperate in the coordination games, which represent cooperative interactions involving future consequences (e.g. arising from repeated interactions, reputation effects, signalling or sanctions). When exploitation occurs (i.e. when one player defects while the other cooperates), the possibility of future consequences has two implications. First, it reduces the benefit to the defector (due to, e.g. lost future cooperation, damaged reputation or material punishment). Second, it reduces the cost to the cooperator (owing to, e.g. switching to defection, improved reputation or material rewards). As a result, it becomes payoff-maximizing to cooperate if one's partner also cooperates (thereby creating a coordination structure). For simplicity, we focus on the limiting case where when one player cooperates and the other defects, both receive zero payoffs.

Within this environment, we examine the evolution of agents with dual-process cognition. We do so by allowing agents to decide to cooperate or defect on the basis of one of two types of cognitive process: intuition or deliberation. Deliberation is more sensitive to game type than intuition (as described below), but is also costlier. Specifically, in the real world, deliberation typically takes more time than intuition, requires more cognitive resources (e.g. those involved in flexible domain-general reasoning), and depends on inhibition of more automatic forms of cognition (e.g. those that govern emotions). Each of these factors can impose material costs on the deliberating agent, for example by causing them to miss out on opportunities to act and by preventing them from devoting cognitive resources to other (non-cooperation related) tasks [27].

Importantly, the magnitude of these costs vary based on the specific context of a particular decision (e.g. how much time the actor has to make a decision or what other tasks the actor had to complete at the same time). Thus, in our model the degree to which deliberation is costlier than intuition varies stochastically from interaction to interaction. Specifically, deliberating in a given interaction requires agents to incur a fitness cost d^* (stochastically sampled for each interaction from the uniform distribution $[0, d]$). Each agent then has a metacognitive threshold T that indicates their maximum willingness-to-pay for deliberation, such that they deliberate in interactions where the deliberation cost is sufficiently small, $d^* \leq T$, but act intuitively when deliberation is too costly, $d^* > T$. (This means that in any given interaction, an agent with threshold T deliberates with probability T/d and uses intuition with probability $1-T/d$.)

In prior work [25], intuition was assumed to be completely inflexible to game type, so there was a single intuitive response that governed whether an agent would cooperate or defect across all situations. By contrast, when an agent deliberated, she perfectly tailored her strategy to the game type she faced in the current interaction. Here, we instead implement more nuanced characterizations of intuition and deliberation, allowing each to vary in the degree of sensitivity that they have to detect context (social dilemma versus coordination game). With probability y_d, deliberation accurately infers its current interaction type, and with probability $1-y_d$ it makes the opposite (incorrect) judgement (i.e. confuses a social dilemma for a coordination game or vice versa). Likewise, with probability y_i, intuition accurately infers its current interaction type, and with probability $1-y_i$ it makes the opposite judgement. Both sensitivity parameters y_d and y_i range from 0.5 to 1. A value of 0.5 corresponds to total insensitivity to context: the judgement of social dilemma versus coordination game is totally random (correct with probability 0.5). A value of 1, conversely, is perfectly sensitive, making no errors. Thus, the extreme dichotomy between intuitive and deliberative cognition used in prior work is a special case of the model we present here, with totally insensitive intuition ($y_i = 0.5$) and perfectly sensitive deliberation ($y_d = 1$).[1]

Therefore, in addition to their metacognitive threshold T, agents have four additional strategy parameters which specify their cooperation choices. When responding using intuition, agents cooperate with probability S_{iD} when intuition indicates they are in a social dilemma and with probability S_{iC} when intuition indicates they are in a coordination game. Conversely, when responding using deliberation, agents cooperate with probability S_{dD} when deliberation indicates they are in a social dilemma and with probability S_{dC} when deliberation indicates they are in a coordination game. Figure 1 summarizes the parametrization of our model.

Our goal is to analyse the evolutionary outcomes of these five quantitates (T, S_{iD}, S_{iC}, S_{dD}, S_{dC}), which together specify an agent's genotype. To do so, we determine the set of Strict Nash Equilibria for our model. A given strategy is Strict Nash Equilibrium if, conditional on the other player using that strategy, you would earn a lower payoff by switching to any other strategy. The evolutionary logic behind this condition is that in a population where all agents use such a strategy, any mutant would be at a disadvantage relative to the rest of the population, and thus no new strategies could invade the population.

For parameter sets where multiple equilibria exist, we use the 'risk-dominance' condition to determine which equilibrium will be favoured by natural selection. Equilibrium strategy X risk-dominates equilibrium strategy Y if X earns a higher payoff than Y in a population composed of half X and half Y (this condition implies that X has a larger basin of attraction than Y, and has been shown to characterize evolutionary outcomes [33]).

We present only the results of these equilibrium and risk-dominance calculations in the main text, and give the derivations in the electronic supplementary material (ESM). In the ESM, we also present calculations showing that the results of these equilibrium analyses align with the outcomes of stochastic evolutionary dynamics.

3. Results

(a) Baseline case

We begin by examining the baseline case of completely insensitive intuition ($y_i = 0.5$) and perfectly sensitive deliberation ($y_d = 1$) explored in [25]. In this limiting case, there are two possible equilibrium strategies. Throughout, we use the following naming convention for strategies. The first word ('intuitive' versus 'dual-process') indicates whether the strategy is completely intuitive (i.e. only uses intuition, $T = 0$) or dual-process (i.e. sometimes deliberates, $T \gg 0$). The second word ('defector' versus 'cooperator' versus 'attender') indicates what the intuitive response is—i.e. whether intuition always defects, always cooperates or attends to information about game type (cooperating when intuition perceives a coordination game and defecting when intuition perceives a social dilemma).

The first equilibrium strategy is the intuitive defector (ID), which is an equilibrium for all values of p, b and c. ID intuitively defects regardless of perceived game type ($S_{iC} = S_{iD} = 0$) and never deliberates ($T = 0$). (As ID never deliberates, its deliberative strategy choices S_{dC} and S_{dD} are undefined.)

The second possible equilibrium strategy is the dual-process cooperator (DC), which is an equilibrium when $p > c/b$ (i.e. when coordination games are sufficiently likely). DC intuitively cooperates regardless of how intuition perceives the game type ($S_{iC} = S_{iD} = 1$). When deliberating, DC overrides the intuitively cooperative response in favour of defection when deliberation perceives the game to be a social dilemma ($S_{dD} = 0$), but maintains the intuitively cooperative response when deliberation perceives the game to be a coordination game ($S_{dC} = 1$). Finally, the maximum cost DC is willing to pay is $T = c(1-p)$. This is the optimal deliberation threshold for DC agents because the benefit for deliberating for DC is the ability to avoid cooperating (and, thus, avoid incurring a cost c) in the fraction $1-p$ of interactions that are social dilemmas. In other words, $c(1-p)$ is DC's expected payoff gain from deliberating, and so deliberation is disadvantageous when it is costlier than this value, leading to $T = c(1-p)$.

A risk-dominance calculation provides the condition for DC to be favoured by selection (see ESM for exact condition). The favoured strategy for the baseline case, as well as the corresponding equilibrium value of T, is shown as a function of p in figure 2*a*. We see that once coordination games become sufficiently likely ($p = 0.3$ for the parameters used in figure 2*a*), DC overtakes ID and T jumps from zero to a high level, $c(1 - p)$. Then as the likelihood of coordination games p increases further, T decreases (as social dilemmas become more rare and therefore deliberation becomes less useful).

(b) Imperfect deliberation

We now examine the effect of weakening deliberation's accuracy at distinguishing social dilemmas from coordination games, such that $0.5 < y_d < 1$. (For tractability, we focus on the case where intuition remains completely insensitive, $y_i = 0.5$; in the ESM, we show that allowing intuition to be sensitive does not qualitatively alter the consequences of imperfect deliberation.) How do errors in deliberation

(*a*)

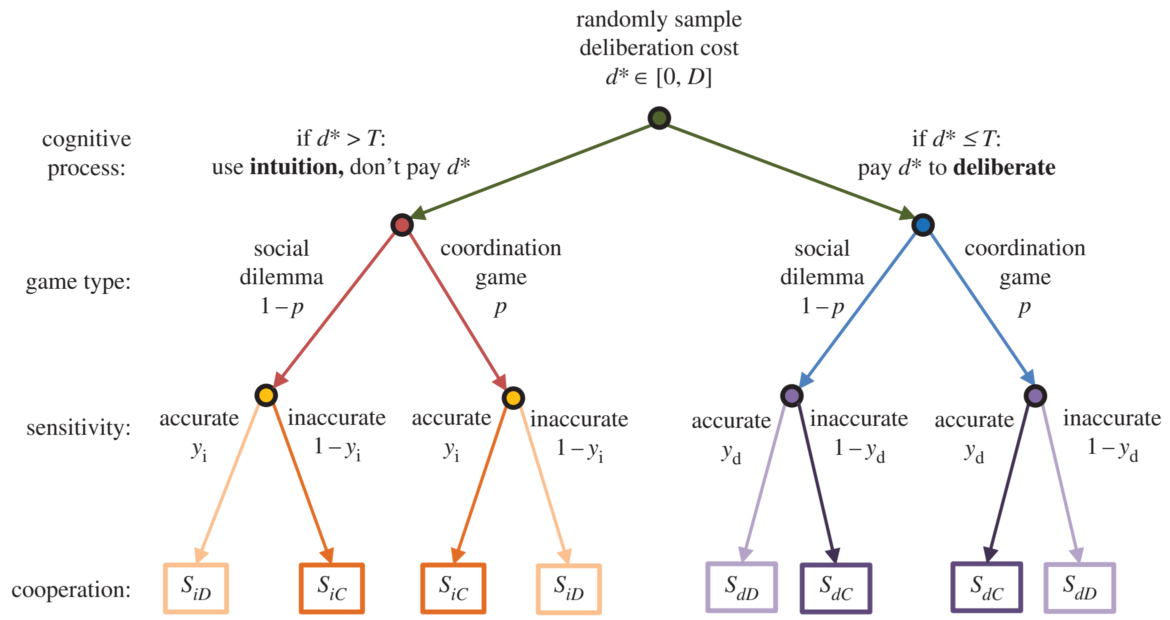

(*b*)

	probability of choosing cooperative strategy				cognitive style
	intuition		deliberation		
	social dilemma S_{iD}	coordination game S_{iC}	social dilemma S_{dD}	coordination game S_{dC}	maximum deliberation cost T
intuitive defector (**ID**)	0	0	–	–	0
intuitive cooperator (**IC**)	1	1	–	–	0
dual-process cooperator (**DC**)	1	1	0	1	>>0 (see ESM)
dual-process attender (**DA**)	0	1	0	1	>>0 (see ESM)

Figure 1. (*a*) Summary of an agent's decision process in the model. First, a deliberation cost is sampled, and on the basis of this and the agent's maximum willingness to pay to deliberate, the agent either uses intuition or deliberation. The agent then tries to implement her desired cooperation strategy for the given interaction, but sometimes makes errors and uses the opposite intuitive or deliberative strategy, depending on the sensitivity of the cognitive process being used. (*b*) Summary of the four possible equilibrium strategies discussed in the main text. (Online version in colour.)

influence the strategies that are selected by evolution and the amount of deliberation that agents are willing to engage in?

As in the baseline case, ID and DC are the only two possible equilibria across the entire $0.5 < y_d < 1$ range. The value of y_d does, however, have a substantial quantitative impact on the maximum cost T that DC is willing to pay to deliberate. This optimal T is influenced by y_d in two ways.

First, because deliberation is sometimes inaccurate, it becomes less beneficial to deliberate in order to avoid wasting the cost of cooperation when in social dilemmas. Specifically, in the $1-y_d$ fraction of cases that deliberation erroneously perceives that a social dilemma is a coordination game (and prescribes continued cooperation), nothing is gained by deliberating. As such, the benefit of deliberation from the baseline model $c(1-p)$ is transformed into $c(1-p)y_d$ here—agents only receive the benefits of defection in social dilemmas in the y_d fraction of cases in which deliberation is accurate.

Second, deliberation errors can lead agents to miss out on the benefits of cooperation in coordination games. Specifically, if a deliberating agent incorrectly perceives a coordination game to be a social dilemma (which happens with probability $1-y_d$), she will defect. Assuming that the other agent avoids making the same error and cooperates (which happens with probability $1-(T/d)(1-y_d)$), the first agent will miss out on the $b-c$ benefit of coordination in the p per cent of interactions that are coordination games.

Putting this all together allows us to calculate the overall expected value of (imperfect) deliberation, and therefore the corresponding optimal deliberation threshold T for the DC strategy. Unlike in the baseline case, however, this expected value can be negative (for example, when y_d is sufficiently small or p is sufficiently large). As an agent's deliberation threshold T is constrained to be non-negative, this means that the equilibrium value of T for the DC strategy is equal to the expected value of deliberation when that value is non-negative, and 0 when it is negative.

In this latter case where $T = 0$, the DC agent simply becomes an intuitive cooperator (IC) who is no longer

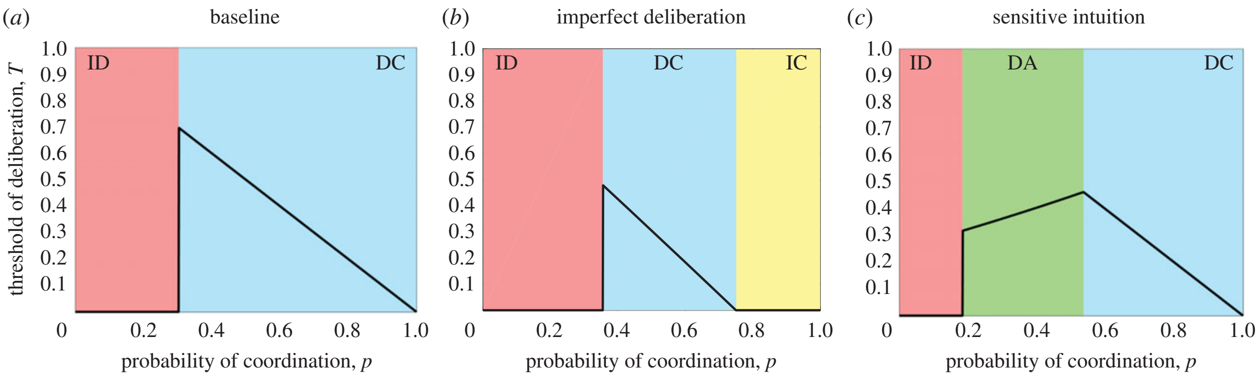

Figure 2. Summary of key results from the model for $b = 4$, $c = 1$ and $d = 1$. Each panel shows the equilibrium strategy and associated threshold of deliberation T, as a function of the probability p that any given interaction is a cooperative coordination game (as opposed to a social dilemma). (*a*) Results from the baseline model in which deliberation is perfectly sensitive to context ($y_d = 1$) and intuition is completely insensitive to context ($y_i = 0.5$). (*b*) Results when deliberation is imperfectly sensitive to context ($y_d = 0.9$) and intuition completely insensitive to context ($y_i = 0.5$). (*c*) Results when intuition is somewhat sensitive to context ($y_i = 0.75$) and deliberation perfectly sensitive to context ($y_d = 1$). (Online version in colour.)

'dual-process'. For clarity, we refer to the DC equilibrium with $T = 0$ as IC, although this is not technically a different equilibrium from DC (just a special case of the DC equilibrium). The non-deliberative IC strategy always plays the cooperative strategy because the risk of accidentally defecting in a coordination game is too great—which is the same 'error management' logic [34,35] used by Delton *et al.* [24] to explain cooperation in one-shot games (an example of a social dilemma). This is illustrated in figure 2*b*, which shows the favoured strategy and corresponding equilibrium value of T for $y_d = 0.9$, as a function of p. A more general characterization of the evolutionary outcomes over the full $[p \times y_d]$ space is shown in figure 3.

Another interesting difference in DC's deliberation threshold T between the baseline case and the case with imperfect deliberation is its dependence on b. When deliberation is perfect ($y_d = 1$), the equilibrium value of $T = c(1-p)$ does not depend on b: because the only consequence of deliberating is being able to avoid the cost of cooperation in social dilemmas, the extent to which cooperation is beneficial is irrelevant. This formulation makes the prediction that, within the region where DC is the risk-dominant equilibrium, making cooperation more beneficial in a social dilemma will have no effect on cooperation.

When $y_d < 1$, however, the equilibrium deliberation threshold T does contain b. Specifically, T decreases as b increases. This is because when deliberation is imperfect, deliberating can be costly by leading to defection in coordination games—and the larger b is, the worse it is to miss out on the benefit of cooperating in coordination games. The consequence of this fact is that as b increases, DC agents become less inclined to deliberate, and therefore become more likely to cooperate in social dilemmas.

In sum, the extent to which deliberation is imperfect does not affect the set of possible equilibrium strategies. Making deliberation inaccurate does, however, (i) reduce the extent to which DC agents deliberate (as one would expect, given that the more imperfect deliberation is, the less useful it is), in the extreme leading to a version of DC which never deliberates and just intuitively cooperates all the time; and (ii) makes it so that DC agents are sensitive to the benefit of cooperation b in social dilemmas, becoming less likely to deliberate (and therefore more likely to cooperate) as b increases.

(c) Sensitive intuition

We next consider how improving intuition's sensitivity, such that $0.5 < y_i < 1$, affects the evolutionary outcomes. (For tractability, we focus on the case where deliberation is perfect, $y_d = 1$; in the ESM, we show that allowing deliberation to be imperfect does not qualitatively alter the consequences of sensitive intuition.) Recall that, as for deliberation, in our model agents have two intuitive responses S_{iD} and S_{iC} corresponding to separate behaviours when intuition perceives the situation it faces as a social dilemma and coordination game, respectively. So, for example, an agent might defect when intuition judges that the interaction is a social dilemma ($S_{iD} = 0$), but intuitively cooperate when it believes it is a coordination game ($S_{iC} = 1$). Alternatively, agents need not attend to the difference between these two situations: they may choose to always intuitively defect ($S_{iD} = 0$ and $S_{iC} = 0$) or always intuitively cooperate ($S_{iD} = 1$ and $S_{iC} = 1$).

As in the baseline case, we again find that ID and DC are possible equilibria, and DC's threshold for deliberation is unaffected by having $y_i > 0.5$: it remains $T = c(1 - p)$. Recall that these strategies do not attend to whether their intuition perceives the game they are facing to be a social dilemma or a coordination game—ID always intuitively defects ($S_{iD} = 0$ and $S_{iC} = 0$) and DC always cooperates when using intuition ($S_{iD} = 1$ and $S_{iC} = 1$).

When y_i is sufficiently large, however, we find the emergence of a qualitatively new equilibrium that we will call the 'dual-process attender' (DA). DA's intuitive response *does* discriminate on the basis of interaction type—intuitively defecting in interactions that intuition perceives as social dilemmas ($S_{iD} = 0$) and intuitively cooperating in interactions that intuition perceives as coordination games ($S_{iC} = 1$). Moreover, DA also engages in deliberation ($T \gg 0$; see ESM for full condition). As with DC, a deliberating DA agent cooperates when deliberation indicates it is playing a coordination game and defects when deliberation indicates it is playing a social dilemma, $S_{dD} = 0$ and $S_{dC} = 1$.

The DA strategy thus exhibits a qualitatively new function for deliberation. Although the deliberative responses for DC and DA are the same, the intuitions that deliberation sometimes overrides are different. Unlike DC, which only uses deliberation to override its intuitively cooperative response in social dilemmas (saving c in these instances by defecting), DA uses deliberation to sharpen the accuracy of an intuition

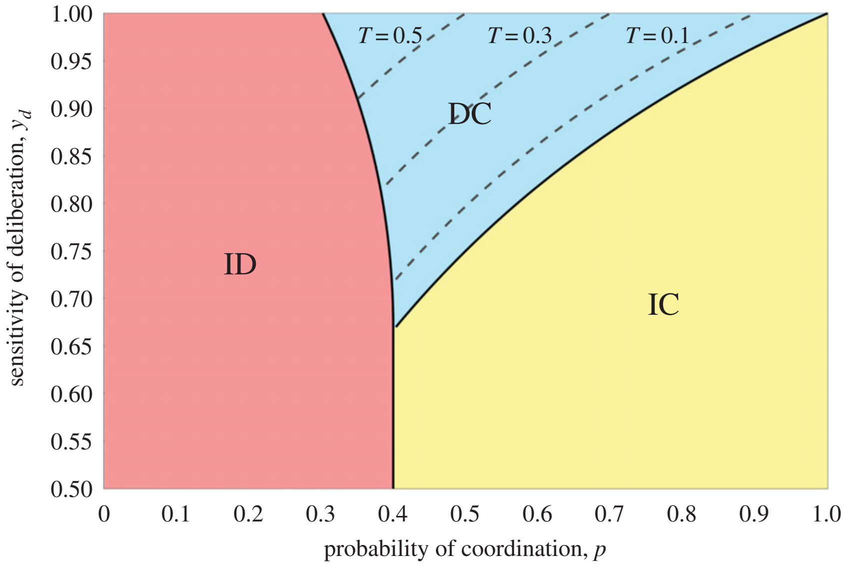

Figure 3. Risk-dominant strategies for $b = 4$, $c = 1$ and $d = 1$ as the probability of cooperative coordination games p and sensitivity of deliberation y_d are varied (with intuition fixed to be fully insensitive to context, $y_i = 0.5$). Dotted lines indicate isoclines for values of the deliberation threshold T within the DC region ($T = 0$ everywhere in the ID and IC regions). Note that IC and DC are not technically different equilibria: IC is merely the special case of DC where $T = 0$. (Online version in colour.)

that is already attending to game type. Consequently, the value of deliberation for DA is inversely related to the accuracy of intuition: as the sensitivity of intuition y_i improves, deliberation becomes less and less useful. By contrast, the value of deliberation for DC does not depend on y_i.

The influence of p, the fraction of interactions that are coordination games, on each of these strategies also importantly differs. For DC, because the value of deliberation lessens as there are fewer and fewer social dilemmas, T is strictly decreasing in p. By contrast, for DA, the value of deliberation results from correcting errors in intuition, which can either occur when intuition misperceives a social dilemma to be a coordination game (thereby leading the agent to waste c by accidentally cooperating) or when intuition misperceives a coordination game to be a social dilemma (thereby leading the agent to throw away $b - c$ by accidentally defecting, in cases where the other agent cooperates). As there is an asymmetry in the costliness of these intuitive errors, such that errors in coordination games are costlier than in social dilemmas ($b - c > c$; as per error management theory [24,34,35]), the overall benefit of deliberating increases as coordination games become more common. Thus, the equilibrium value of T is actually increasing in p for DA agents.

An interesting corollary of this result is that deliberation can have positive consequences for cooperation when agents play DA. For both DC agents and DA agents, deliberation undermines cooperation in social dilemmas by overriding intuitive cooperation. For DC agents, this was deliberation's only function, and so deliberation only worked to reduce cooperativeness. For DA agents, however, deliberation can also *promote* cooperation by overriding (mistaken) intuitive defection when in coordination games.

The favoured strategy and corresponding equilibrium value of T is shown as a function of p for $y_i = 0.75$ in figure 2*c*, and a more general characterization of the evolutionary outcomes over the full $[p \times y_d]$ space is shown in figure 4. Notably, DC (which ignores intuition's sensitivity and always intuitively cooperates) is the favoured equilibrium for high values of p, even when intuition is quite sensitive (y_i is quite large). This indicates that the conclusions of the baseline model are quantitatively unaffected by sensitive intuition, so long as p is sufficiently large. Furthermore, this illustrates an important point about the adaptive significance of cognitive ability: even if an agent has the cognitive capacity to attend to meaningful differences, it is not necessarily adaptive to do so. This follows from the asymmetry in costs between the errors of defecting in a coordination game and cooperating in a social dilemma. As the former is costlier, it can be adaptive for agents to err on the side of cooperation even if they are reasonably sure they could get away with defecting (as per [24]).

Also of note is the transition, for low values of p, from ID to DA as intuition becomes more sensitive (y_i increases). This transition may seem counterintuitive: improving intuition (increasing y_i) actually *reduces* the extent to which agents rely on intuition, leading a purely intuitive agent (ID) to be outcompeted by an only partially intuitive agent (DA). How could making intuition better lead to a decrease in its use?

The answer lies in considering why ID agents do not deliberate. It is only valuable for intuitively defecting agents to deliberate and switch to cooperation when (i) they are in coordination games and (ii) the other agent *also* deliberates and overrides their selfish response. This dynamics sets up a second-order coordination problem, where each agent's willingness to deliberate is linearly dependent on the other agent's willingness to deliberate. In most parameter regions, agents maximize their payoffs by deliberating slightly less than their partners, leading selection to an equilibrium in which nobody deliberates at all (ID). (See [25] for a detailed discussion of this race to the bottom in deliberativeness.)

The situation changes, however, when intuition becomes sufficiently sensitive. Here, DA becomes the favoured equilibrium, and DA agents cooperate when intuition perceives that they are in a coordination game ($S_{iC} = 1$). Because of this, one's partner may cooperate even if they do not deliberate; and therefore, there is value in deliberating even if one's partner does not. This (at least partially) resolves the second-order coordination problem regarding deliberation, and leads to selection favouring non-zero deliberation. Thus, we see that making intuition more sensitive can, paradoxically, *increase* the value of deliberation for intuitively discriminating

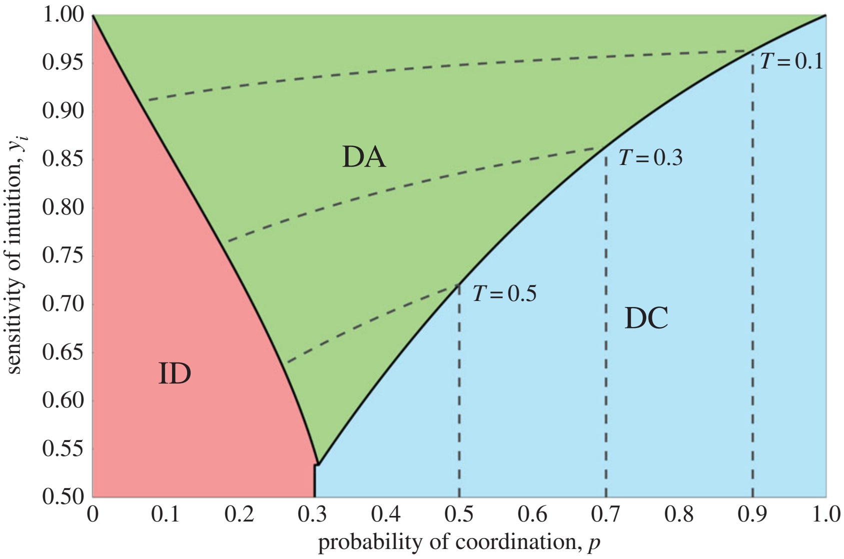

Figure 4. Risk-dominant strategies for $b = 4$, $c = 1$ and $d = 1$ as the probability of cooperative coordination games p and sensitivity of intuition y_i are varied (with deliberation fixed to be perfectly sensitive to context, $y_d = 1$). Dotted lines indicate isoclines for values of the deliberation threshold T ($T = 0$ everywhere in the ID region). (Online version in colour.)

Table 1. Deliberation's effects on cooperation for the four equilibrium strategies.

strategy	social dilemmas	coordination games
intuitive defector (ID)	no effect	no effect
intuitive cooperator (IC)	no effect	no effect
dual-process cooperator (DC)	reduces cooperation	no effect
dual-process attender (DA)	reduces cooperation	increases cooperation

agents—who want to ensure that they do not miss out on the $b - c$ payoff from mutual cooperation in coordination games.

4. Discussion

Here, we have explored the consequences of imperfect deliberation and sensitive intuition for the co-evolution of cognition and cooperation. We find that the following four observations hold regardless of how imperfect deliberation is, and how sensitive intuition is. First, when the probability of coordination games p is sufficiently small, selection will favour a strategy that has a totally insensitive intuition to defect in all situations, and that never deliberates (ID). Second, when p is sufficiently large, selection will favour a strategy whose intuition is totally insensitive and chooses cooperation regardless of the type of interaction it perceives (either DC that does deliberate or IC that does not, depending on how imperfect deliberation is). Third, deliberation only ever reduces cooperation in social dilemmas (for DC and DA) or has no effect (for IC and ID), but never increases social-dilemma cooperation. Fourth, deliberation only ever increases cooperation in coordination games (for DA) or has no effect (for IC, ID and DC), but never decreases coordination game cooperation.

From the first and second observations, we see that selection can favour intuitions that indiscriminately cooperate or defect across interactions, even when intuition is sensitive and therefore discrimination is possible. Although greater intelligence and cognitive complexity is often favoured by selection [36], simplicity can also be adaptive. Even when intuition can perceive context reasonably well without exerting any cognitive cost, evolution may still select for a simpler strategy that does not condition its behaviour on this information. This result illustrates the importance of studying cognitive processes not only at a proximate level, but also at an ultimate level: sometimes, simple cognition is actively built to be simple by design, and not because of cognitive constraints. In other words, inflexible intuitions that are observed in the real world may have gotten this way through adaptive selection, rather than cognitive compromise.

The third and fourth observations generate clear theoretical predictions about the consequences of experimentally manipulating the use of intuition versus deliberation (see table 1 for a summary). These predictions are borne out by empirical evidence from economic cooperation games. This can be seen in the results of a meta-analysis of 67 such experiments in which cognitive processing was manipulated [28]: in social-dilemma settings, deliberation was found to reduce cooperation, whereas in coordination settings, deliberation had no effect on cooperation.

While these four observations from our model do not depend on the extent of imperfect deliberation and sensitive intuition, other features do, lead to important consequences for the evolutionary outcomes. To start, although weakening deliberation's accuracy does not change the two possible equilibria that can be favoured by selection, it does affect some important features of the extent to which the DC strategy chooses to deliberate.

First, it makes DC's deliberation sensitive to the benefit of cooperation: the more beneficial it is to receive cooperation from one's partner, the less inclined non-perfect deliberators are to deliberate (out of fear of misperceiving a coordination game as a social dilemma and accidentally defecting; error management, cf. [24,34,35]). As a result, the DC strategy with imperfect deliberation is more likely to cooperate in social dilemmas as the benefit of cooperation increases, a pattern that fits empirical data (e.g. [37]).

Second, imperfect deliberation undermines deliberation's value, reducing the amount of deliberation that DC engages in. Even a fairly modest reduction in deliberation's accuracy can, under the right set of values for the other parameters, lead the DC equilibrium to collapse into a purely IC strategy, which places no value on the discrimination capabilities of deliberation (and thus never deliberates). This result can also be explained by error management: because missing out on the benefits of cooperation in coordination games is costlier than cooperating in social dilemmas, it can be payoff-maximizing to cooperate even when there is good reason to believe that one could get away with defecting. Given this, we might expect deliberation in the real world to be highly tuned to contextual factors since, otherwise, it would be of little use.

We now turn from the imperfection of deliberation to the sensitivity of intuition. As discussed above, selection can favour strategies that ignore sensitive intuition's ability to discriminate. However, once intuition becomes *sufficiently* sensitive, it starts being used in the same discriminating way that deliberation is used. (Recall that DA's intuitive response discriminates between social dilemmas and coordination games, and DA sometimes pays a cost to deliberate in order to improve the accuracy of this discrimination.) The emergence of this qualitatively new dual-process attender equilibrium, which differs from all of the other equilibria in that it makes use of intuition's sensitivity, has important implications for how people's reasoning affects cooperative behaviour in coordination games.

In social dilemmas, a discussed above, empirical work suggests that people become more selfish when they deliberate (for a meta-analysis, see [28]). Both the DC and DA equilibria predict this pattern of results. For DC, deliberation's only function is to override a cooperative response in social dilemmas. For DA, deliberation allows the agent to better tailor her strategy to the situation at hand, which leads her to defect more than intuition would in social dilemmas.

By contrast, DC and DA produce conflicting predictions regarding the role of deliberation in coordination games. DC agents, who always intuitively cooperate, never change their behaviour when learning that an interaction is a coordination game, whereas DA agents sometimes intuitively defect in these situations and thereby become more cooperative if they stop and deliberate.

This is an important result because in many real-world environments, intuition is likely to be reasonably good at discriminating (and thus, our model predicts that DA will be favoured). For example, simple forms of automatic and intuitive cognition are often quite adept at distinguishing ingroup members (likely potential beneficiaries of reciprocity) from outgroup members (unlikely potential beneficiaries of reciprocity) on the basis of cues like skin colour, accent and other salient perceptual features [38]. But stereotyping on the basis of these cues can lead to systematic inefficiencies, where people of different groups or ethnicities become outcasts in a community even when they could be contributing members of society [39,40]. In this class of situations, in which intuition attends to superficial features, deliberation can actually promote mutually beneficial cooperation by leading people to think more carefully about the situation they are in and whether these cues are valid for everyone's goals at the moment. Though some evidence already speaks to deliberation's role in reducing prejudice (e.g. studies on implicit bias among people who explicitly reject such bias [41]), more research is needed to directly assess this prediction and its dependence on the sensitivity of intuition.

Other domains, however, are not so rich with cues for intuition to pick up on. For example, in most economic game experiments, participants interact anonymously via the computer and thus have no visual or auditory cues regarding their partner's identity. In these situations, intuition should be relatively insensitive, and therefore our model predicts that DC should be favoured—such that the intuitive response favours cooperation regardless of game type. As a result, deliberation should not increase cooperation, as the intuitive and deliberative cooperative responses are perfectly aligned in coordination games (e.g. the repeated Prisoner's Dilemma). This prediction is supported by the experimentally observed lack of effect of deliberation on cooperation in 'strategic' coordination situations described above [28].

Finally, we note some limitations of this work. Only two types of games (social dilemma and coordination game) are considered, rather than a continuum of different possible games; the deliberation cost is sampled from a simple uniform distribution, rather than a more realistic cost distribution; and cognitive processing is unobservable to other agents, such that agents cannot condition on each other's use of intuition versus deliberation (as per [15,42]). Moreover, the flexibility of deliberation and intuition are assumed to be fixed parameters of the environment rather than dynamic features of individuals or populations, which could evolve or interact with other variables. We also do not consider the possibility that deliberation could alter beliefs, which could in turn lead to increased cooperation. Future work could augment the current model to test the robustness of these simplifications and explore how many of the key results about ease–flexibility trade-offs translate into other domains, involving, e.g. social norm compliance, intertemporal choice (as per [43,44]) or punishment.

In sum, by exploring the dynamics of imperfect deliberation and context-sensitive intuition, the model presented here brings us one step closer to linking proximate dual-process cognition with evolutionary game-theory. In doing so, we generate new predictions and help to clarify the ways in which deliberation can, and cannot, work to improve the greater good.

Ethics. This research was conducted without the participation of any human or animal subjects.

Data accessibility. No data were collected in order to conduct the research presented in this manuscript.

Authors' contributions. A.B., A.K. and D.G.R. designed research, performed research, analysed data and wrote the paper.

Competing interests. The authors declare that they have no competing interests.

Funding. The authors gratefully acknowledge funding from the Templeton World Charity Foundation (grant no. TWCF0209) and the Defense Advanced Research Projects Agency NGS2 programme (grant no. D17AC00005).

Acknowledgements. We thank Joshua Greene, Adam Morris, and Jonathan Phillips for helpful feedback on previous drafts.

Endnotes

[1]As we are interested in exploring the trade-off between ease and flexibility, we focus on cases where $y_d > y_i$—i.e. we assume that deliberation is more sensitive than intuition. In the real world, this condition may not always hold (e.g. for certain fast and frugal heuristics), but we do not consider these cases, as their evolutionary dynamics would be trivial, with selection clearly favouring intuition because there would be no ease–flexibility trade-off between intuition and deliberation.

References

1. Nowak MA. 2006 Five rules for the evolution of cooperation. *Science* **314**, 1560–1563. (doi:10.1126/science.1133755)
2. Lehmann L, Bargum K, Reuter M. 2006 An evolutionary analysis of the relationship between spite and altruism. *J. Evol. Biol.* **19**, 1507–1516. (doi:10.1111/j.1420-9101.2006.01128.x)
3. West SA, Griffin AS, Gardner A. 2007 Evolutionary explanations for cooperation. *Curr. Biol.* **17**, R661–R672. (doi:10.1016/j.cub.2007.06.004)
4. Sigmund K. 2010 *The calculus of selfishness*. Princeton, NJ: Princeton University Press.
5. Bowles S, Gintis H. 2003 Origins of human cooperation (ed. P Hammerstein). *Genetic Cult. Evol. Cooperation*, 429–443. Cambridge, MA: MIT Press.
6. Axelrod R, Hamilton WD. 1981 The evolution of cooperation. *Science* **211**, 1390–1396. (doi:10.1126/science.7466396)
7. Fudenberg D, Maskin ES. 1990 Evolution and cooperation in noisy repeated games. *Am. Econ. Rev.* **80**, 274–279.
8. Fudenberg D, Maskin ES. 1986 The folk theorem in repeated games with discounting or with incomplete information. *Econometrica* **54**, 533–554. (doi:10.2307/1911307)
9. Nowak MA, Sigmund K. 2005 Evolution of indirect reciprocity. *Nature* **437**, 1291–1298. (doi:10.1038/nature04131)
10. Gintis H, Smith EA, Bowles S. 2001 Costly signaling and cooperation. *J. Theor. Biol.* **213**, 103–119. (doi:10.1006/jtbi.2001.2406)
11. Jordan JJ, Hoffman M, Bloom P, Rand DG. 2016 Third-party punishment as a costly signal of trustworthiness. *Nature* **530**, 473–476. (doi:10.1038/nature16981)
12. Roberts G. 2008 Evolution of direct and indirect reciprocity. *Proc. R. Soc. B* **275**, 173–179. (doi:10.1098/rspb.2007.1134)
13. 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**, 1–11. (doi:10.1007/BF00167053)
14. van Veelen M, García J, Rand DG, Nowak MA. 2012 Direct reciprocity in structured populations. *Proc. Natl Acad. Sci. USA* **109**, 9929–9934. (doi:10.1073/pnas.1206694109)
15. Hoffman M, Yoeli E, Nowak MA. 2015 Cooperate without looking: why we care what people think and not just what they do. *Proc. Natl Acad. Sci. USA* **112**, 1727–1732. (doi:10.1073/pnas.1417904112)
16. Camerer CF. 2003 *Behavioral game theory: experiments in strategic interaction*. Princeton, NJ: Princeton University Press.
17. Rand DG, Greene JD, Nowak MA. 2012 Spontaneous giving and calculated greed. *Nature* **489**, 427–430. (doi:10.1038/nature11467)
18. Zaki J, Mitchell JP. 2013 Intuitive prosociality. *Curr. Directions Psychol. Sci.* **22**, 466–470. (doi:10.1177/0963721413492764)
19. Sloman SA. 1996 The empirical case for two systems of reasoning. *Psychol. Bull.* **119**, 3. (doi:10.1037/0033-2909.119.1.3)
20. Kahneman D. 2003 A perspective on judgment and choice: mapping bounded rationality. *Am. Psychol.* **58**, 697–720. (doi:10.1037/0003-066X.58.9.697)
21. Stanovich KE, West RF. 1998 Individual differences in rational thought. *J. Exp. Psychol. Gen.* **127**, 161–188. (doi:10.1037/0096-3445.127.2.161)
22. Rand DG, Peysakhovich A, Kraft-Todd GT, Newman GE, Wurzbacher O, Nowak MA, Green JD. 2014 Social heuristics shape intuitive cooperation. *Nat. Commun.* **5**, 3677. (doi:10.1038/ncomms4677)
23. Kiyonari T, Tanida S, Yamagishi T. 2000 Social exchange and reciprocity: confusion or a heuristic? *Evol. Hum. Behav.* **21**, 411–427. (doi:10.1016/S1090-5138(00)00055-6)
24. Delton AW, Krasnow MM, Cosmides L, Tooby J. 2011 Evolution of direct reciprocity under uncertainty can explain human generosity in one-shot encounters. *Proc. Natl Acad. Sci. USA* **108**, 13 335–13 340. (doi:10.1073/pnas.1102131108)
25. Bear A, Rand DG. 2016 Intuition, deliberation, and the evolution of cooperation. *Proc. Natl Acad. Sci. USA* **113**, 936–941. (doi:10.1073/pnas.1517780113)
26. Baumeister RF, Bratslavsky E, Muraven M, Tice DM. 1998 Ego depletion: is the active self a limited resource? *J. Pers. Soc. Psychol.* **74**, 1252. (doi:10.1037/0022-3514.74.5.1252)
27. Kurzban R, Duckworth A, Kable JW, Myers J. 2013 An opportunity cost model of subjective effort and task performance. *Behav. Brain Sci.* **36**, 661–679. (doi:10.1017/S0140525X12003196)
28. Rand DG. 2016 Cooperation, fast and slow: meta-analytic evidence for a theory of social heuristics and self-interested deliberation. *Psychol. Sci.* **27**, 1192–1206. (doi:10.1177/0956797616654455)
29. Burton-Chellew MN, El Mouden C, West SA. 2016 Conditional cooperation and confusion in public-goods experiments. *Proc. Natl Acad. Sci. USA* **113**, 1291–1296. (doi:10.1073/pnas.1509740113)
30. Haidt J. 2001 The emotional dog and its rational tail: a social intuitionist approach to moral judgment. *Psychol. Rev.* **108**, 814. (doi:10.1037/0033-295X.108.4.814)
31. Gigerenzer G, Goldstein DG. 1996 Reasoning the fast and frugal way: models of bounded rationality. *Psychol. Rev.* **103**, 650. (doi:10.1037/0033-295X.103.4.650)
32. Gigerenzer G, Todd PM, Group AR. 1999 *Simple heuristics that make us smart*. Oxford, UK: Oxford University Press.
33. Kandori M, Mailath GJ, Rob R. 1993 Learning, mutation, and long run equilibria in games. *Econometrica: J. Econ. Soc.* **61**, 29–56. (doi:10.2307/2951777)
34. Haselton MG, Buss DM. 2000 Error management theory: a new perspective on biases in cross-sex mind reading. *J. Pers. Soc. Psychol.* **78**, 81. (doi:10.1037/0022-3514.78.1.81)
35. Haselton MG, Nettle D. 2006 The paranoid optimist: an integrative evolutionary model of cognitive biases. *Pers. Soc. Psychol. Rev.* **10**, 47–66. (doi:10.1207/s15327957pspr1001_3)
36. McNally L, Brown SP, Jackson AL. 2012 Cooperation and the evolution of intelligence. *Proc. R. Soc. B* **279**, 3027–3034. (doi:10.1098/rspb.2012.0206)
37. Capraro V, Jordan JJ, Rand DG. 2014 Heuristics guide the implementation of social preferences in one-shot Prisoner's Dilemma experiments. *Sci. Rep.* **4**, 6790. (doi:10.1038/srep06790)
38. Fiske ST. 2002 What we know now about bias and intergroup conflict, the problem of the century. *Curr. Direc. Psychol. Sci.* **11**, 123–128. (doi:10.1111/1467-8721.00183)
39. Greene J. 2014 *Moral tribes: emotion, reason and the gap between us and them*. London, UK: Atlantic Books Ltd.
40. Stagnaro MN, Dunham Y, Rand DG. 2017 Profit versus prejudice: harnessing self-interest to reduce in-group bias. *Soc. Psychol. Person. Sci.* (doi:10.1177/1948550617699254)
41. Greenwald AG, Banaji MR. 1995 Implicit social cognition: attitudes, self-esteem, and stereotypes. *Psychol. Rev.* **102**, 4. (doi:10.1037/0033-295X.102.1.4)
42. Jordan JJ, Hoffman M, Nowak MA, Rand DG. 2016 Uncalculating cooperation is used to signal trustworthiness. *Proc. Natl Acad. Sci. USA* **113**, 8658–8663. (doi:10.1073/pnas.1601280113)
43. Tomlin DA, Rand DG, Ludvig E, Cohen JD. 2015 The evolution and devolution of cognitive control: The costs of deliberation in a competitive world. *Sci. Rep.* **5**, 11002.
44. Toupo DFP, Strogatz SH, Cohen JD, Rand DG. 2015 Evolutionary game dynamics of controlled and automatic decision-making. *Chaos* **25**, 073120. (doi:10.1063/1.4927488)

Supplementary Material

Adam Bear, Ari Kagan, & David G. Rand

Contents

1 Strategy space and payoff function

1.1 Original model (Bear & Rand, 2016)

Our original model considered agents playing a mix of Social/Prisoner's Dilemma (PD) games and Coordination Games (CG); and responding using either a generalized intuition S_i or paying a cost d^* (stochastically sampled from the interval $[0, d]$) to deliberate and tailor their strategy such that they used strategy S_{dC} if the game was a CG and S_{dD} if the game was a PD. In each interaction, agents chose either to cooperate or defect. (Note that we use a slightly different notation here than the notation used in the original manuscript.)

An agent's strategy profile was specified by four variables: their (i) probability of intuitively cooperating S_i, (ii) probability of cooperating when they deliberated and faced a coordination game S_{dC}, (iii) probability of cooperating when they deliberated and faced a prisoner's dilemma S_{dD}, and (iv) maximum acceptable cost of deliberation T. Since we stipulated that the cost of deliberation was sampled uniformly from the interval $[0, d]$, an agent with threshold T deliberated with probability $\frac{T}{d}$ and on average payed a cost $\frac{T}{2}$ when deliberating.

Here we specify the expected payoff $\pi(x, y)$ of an agent with strategy profile $x = [S_i, S_{dD}, S_{dC}, T]$ playing against an agent with strategy profile $y = [S'_i, S'_{dD}, S'_{dC}, T']$. To do so, we calculate agent x's expected payoff from playing CGs with probability p and PDs with probability $1 - p$, over the cases in which (i) both agents deliberate (probability $\frac{TT'}{d^2}$), (ii) agent x deliberates and agent y decides intuitively (probability $\frac{T}{d}(1 - \frac{T'}{d})$), (iii) agent x decides intuitively and agent y deliberates (probability $(1 - \frac{T}{d})\frac{T'}{d}$), and (iv) both agents decide intuitively (probability $(1 - \frac{T}{d})(1 - \frac{T'}{d})$):

$$\pi(x, y) = \frac{TT'}{d^2}(\pi_{DD} - \frac{T}{2}) + \frac{T}{d}(1 - \frac{T'}{d})(\pi_{DI} - \frac{T}{2}) + (1 - \frac{T}{d})\frac{T'}{d}\pi_{ID} + (1 - \frac{T}{d})(1 - \frac{T'}{d})\pi_{II}$$

where π_{DD} is agent x's expected payoff when both agents deliberate, π_{DI} is agent x's expected payoff when agent x deliberates and agent y uses intuition, and so on.

These expected payoffs are calculated based on the payoff tables for PDs and CGs. In PDs, agents who cooperate pay a cost c to give a benefit b to their partner, while agents who defect do nothing. Thus, the payoff table for the PD is given by

PD Payoffs

	C	**D**
C	$b - c$	$-c$
D	b	0

where the row player's payoff is shown.

In CDs, as in PDs, agents who cooperate pay a cost c to give a benefit b to their partner, while agents who defect do nothing. However, when exploitation occurs, such that one player defects while the other cooperates, the benefit to the defector is reduced (due to, e.g., lost future cooperation, damaged reputation, or material punishment), as is the cost to the cooperator (due to, e.g., switching to defection, improved reputation, or material rewards). As a result, it becomes payoff-maximizing to cooperate if one's partner also cooperates (thereby creating a coordination structure). Our model takes the limiting case where when one player cooperates and the other defects, both receive zero payoffs.

CG Payoffs

	C	**D**
C	$b-c$	0
D	0	0

where $b, c > 0$.

Substituting in relevant payoff values yields

$$
\begin{aligned}
\pi_{DD} &= p(S_{dC}S'_{dC}(b-c)) + (1-p)(S_{dD}S'_{dD}(b-c) + S_{dD}(1-S'_{dD})(-c) + (1-S_{dD})S'_{dD}b) \\
\pi_{DI} &= p(S_{dC}S'_{i}(b-c)) + (1-p)(S_{dD}S'_{i}(b-c) + S_{dD}(1-S'_{i})(-c) + (1-S_{dD})S'_{i}b) \\
\pi_{ID} &= p(S_{i}S'_{dC}(b-c)) + (1-p)(S_{i}S'_{dD}(b-c) + S_{i}(1-S'_{dD})(-c) + (1-S_{i})S'_{dD}b) \\
\pi_{II} &= p(S_{i}S'_{i}(b-c)) + (1-p)(S_{i}S'_{i}(b-c) + S_{i}(1-S'_{i})(-c) + (1-S_{i})S'_{i}b)
\end{aligned}
$$

1.2 Current model

The present model expands upon the model above by supposing that (i) deliberation can sometimes incorrectly judge whether the agent is in a PD or a CG, and (ii) intuition can sometimes accurately judge whether the agent is in a PD or a CG. These changes require the introduction of two new parameters and one new strategy variable.

The parameter y_d, which ranges from .5 to 1, now tracks the accuracy/sensitivity of deliberation. When $y_d = .5$, deliberation accurately judges whether the agent is in a PD vs. CG with 50% probability and mistakenly gives the wrong judgment the other 50% of the time. In other words, deliberation is completely insensitive to context, like intuition in the original model. Conversely, when $y_d = 1$, deliberation always yields the correct answer, as in the original model.

Similarly, the parameter y_i, which also ranges from .5 to 1, now tracks the accuracy/sensitivity of intuition. When $y_i = .5$, intuition is completely insensitive to context, as in the original model. When $y_i = 1$, it is infallible and indistinguishable from deliberation. For the purposes of the model, we stipulate that it is always the case that $y_d > y_i$, i.e., that costly deliberation has some advantage over intuition.

Because intuition, like deliberation, can now sometimes discriminate PDs from CGs, we replace the old strategy variable S_i, which tracked agents' generalized intuitive response, with two new variables S_{iD} and S_{iC}, which track what an agent's intuitive response is when intuition dictates that the agent is in a PD vs. a CG, respectively. Thus, an agent's strategy profile x now consists of 5 variables: $x = [S_{iD}, S_{iC}, S_{dD}, S_{dC}, T]$.

These additions to the original model further complicate the calculation of expected payoffs for each agent. For instance, because deliberation now sometimes leads an agent to incorrectly judge that a

CG is a PD or vice versa (with probability $1 - y_d$), the expression for the payoff of a deliberating agent interacting with another deliberating agent becomes

$$\begin{aligned}\pi_{DD} &= y_d^2(p(S_{dC}S'_{dC}(b-c)) + (1-p)(S_{dD}S'_{dD}(b-c) + S_{dD}(1-S'_{dD})(-c) + (1-S_{dD})S'_{dD}b)\\ &+ (y_d(1-y_d))(p(S_{dC}S'_{dD}(b-c)) + (1-p)(S_{dD}S'_{dC}(b-c) + S_{dD}(1-S'_{dC})(-c) + (1-S_{dD})S'_{dC}b))\\ &+ ((1-y_d)y_d)(p(S_{dD}S'_{dC}(b-c)) + (1-p)(S_{dC}S'_{dD}(b-c) + S_{dC}(1-S'_{dD})(-c) + (1-S_{dC})S'_{dD}b)\\ &+ (1-y_d)^2(p(S_{dD}S'_{dD}(b-c)) + (1-p)(S_{dC}S'_{dC}(b-c) + S_{dC}(1-S'_{dC})(-c) + (1-S_{dC})S'_{dC}b).\end{aligned}$$

Similarly, the payoff of an intuitive agent interacting with another intuitive agent now depends on y_i in the following way:

$$\begin{aligned}\pi_{II} &= y_i^2(p(S_{iC}S'_{iC}(b-c)) + (1-p)(S_{iD}S'_{iD}(b-c) + S_{iD}(1-S'_{iD})(-c) + (1-S_{iD})S'_{iD}b)\\ &+ (y_i(1-y_i))(p(S_{iC}S'_{iD}(b-c)) + (1-p)(S_{iD}S'_{iC}(b-c) + S_{iD}(1-S'_{iC})(-c) + (1-S_{iD})S'_{iC}b))\\ &+ ((1-y_i)y_i)(p(S_{iD}S'_{iC}(b-c)) + (1-p)(S_{iC}S'_{iD}(b-c) + S_{iC}(1-S'_{iD})(-c) + (1-S_{iC})S'_{iD}b)\\ &+ (1-y_i)^2(p(S_{iD}S'_{iD}(b-c)) + (1-p)(S_{iC}S'_{iC}(b-c) + S_{iC}(1-S'_{iC})(-c) + (1-S_{iC})S'_{iC}b).\end{aligned}$$

The other payoffs, π_{DI} and πID are calculated in the exact same manner.

2 Nash equilibrium calculations

2.1 Setup

To facilitate Nash equilibria calculations, we consider a strategy space that is simplified relative to the main model in two ways: (i) agents' intuitive responses, S_{iD} and S_{iC}, are limited to being either 0 (never cooperate) or 1 (always cooperate); and (ii) agents' deliberative responses are fixed to be $S_{dD} = 0$ and $S_{dC} = 1$; i.e., always defecting when deliberating and facing a PD, and always cooperating when deliberating and facing a CG. As in the main model, agents specify a maximum cost of deliberation T ($0 \leq T \leq d$) that they are willing to pay in order to deliberate, and this determines when they deliberate.

We divide our Nash calculations into two parts. First, we explore the effects of making deliberation imperfect by considering what happens when we fix $y_i = .5$ (i.e., assume that intuition is completely insensitive to context, as in the original model) and vary y_d in the range $(.5, 1)$. For the sake of this set of calculations, S_{iD} and S_{iC} are functionally identical, and therefore we stick to the simpler notation of referring to both of these intuitive responses as S_i.

Second, we explore what happens when we stipulate that deliberation is perfect ($y_d = 1$), as in the original model, and vary y_i in the range $(.5, 1)$. This allows us to explore the implications of imperfect intuition on the strategies that evolution selects.

In these calculations, we make one further assumption that was not made in Bear & Rand's (2016) original model. We suppose that $d > p(b-c)$, i.e., that the maximum possible cost of deliberation d is sufficiently large to prevent an equilibrium strategy that *always* deliberates and never uses intuition. This simplification is appropriate here in light of the fact that we are exploring what happens when deliberation becomes relatively less effective, compared to the effectiveness of intuition. We are, therefore, less concerned about what happens when deliberation is very cheap and effective, but an extensive discussion of this purely deliberating equilibrium is included in Bear & Rand (2016).

Thus, an agent's strategy profile is specified by three variables: 1) a binary variable S_{iD} indicating whether or not the agent intuitively plays the cooperative strategy when intuition dictates that it is a PD; 2) a binary variable S_{iC} indicating whether or not the agent intuitively plays the cooperative strategy when intuition dictates that it is a CG; and 3) a continuous variable T indicating the agent's maximum cost they are willing to pay to deliberate. We denote a strategy profile for this reduced strategy space as $x = [S_{iD}, S_{iC}, T]$. However, when considering situations in which $y_i = .5$ (i.e., when we are varying the effectiveness of deliberation while assuming that intuition is purely insensitive to context), we denote a strategy in terms of just two variables $x = [S_i, T]$, since S_{iD} and S_{iC} are mathematically indistinguishable.

A strategy profile x is a Nash equilibrium if no strategy profile y is able to get a higher payoff against x than x gets against itself. That is,

$$\forall y : \pi(x, x) \geq \pi(y, x).$$

Given our restricted strategy space, the set of possible strategy profiles that an agent can adopt can be thought of as four continuous sets: 1) the set of strategy profiles that always intuitively defect ($S_{iD} = 0$ $S_{iC} = 0$) and have threshold $0 \leq T \leq d$; 2) the set of strategy profiles that always intuitively cooperate ($S_{iD} = 1$ $S_{iC} = 1$) and have threshold $0 \leq T \leq d$; 3) the set of strategy profiles that intuitively defect ($S_{iD} = 0$) when intuition thinks the interaction is a PD and intuitively cooperate when intuition thinks the opposite ($S_{iC} = 1$) and have threshold $0 \leq T \leq d$; and 4) strategy profile (3) with the intuitive responses reversed ($S_{iD} = 1$ $S_{iC} = 0$). However, when $y_i = .5$, only strategy profiles (1) and (2) are possible.

2.2 Imperfect Deliberation

We begin by exploring what happens when we fix $y_i = .5$ and allow the accuracy of deliberation to vary such that $.5 < y_d < 1$. We break the calculation up into two parts, in which the intuitive response is either to defect ($S_i = 0$) or to cooperate ($S_i = 1$).

2.2.1 Intuitively defecting equilibria

We first consider whether any strategy profile with $S_i = 0$ is a Nash. To do this, we calculate the expression for the payoff that an agent with $S_i = 0$ and $T = T$ gets against an agent with $S_i = 0$ and $T = T'$:

$$\begin{aligned}\pi([0,T],[0,T']) = &\frac{(T(\frac{T'}{d} - 1)(\frac{T}{2} + 2c(\frac{y_d}{2} - .5)(p-1)))}{d} \\ &- (TT'(\frac{T}{2} + (p-1)((b-c)(y_d-1)^2 - by_d(y_d-1) + cy_d(y_d-1)) \\ &- \frac{py_d^2(b-c)}{d^2})) - \frac{(2T'b(\frac{y_d}{2} - .5)(p-1) * (\frac{T}{d} - 1))}{d}\end{aligned}$$

Since the concavity of this function (with respect to T) is always negative ($\frac{\partial^2}{\partial T^2}\pi([0,T],[0,T']) < 0$), there is a unique best-response $[0, T_b]$ that maximizes one's payoff when playing against $[0, T']$, which can be found by asking what value of T satisfies the equation

$$\frac{\partial}{\partial T}\pi([0,T],[0,T']) = 0.$$

Doing so yields

$$T_b = \frac{cdp - cd + cdy_d - cdpy_d + T'bpy_d^2 - T'cpy_d^2}{d}.$$

Since a strategy profile must be a best response against itself in order to be Nash, it must be the case that $T_b = T'$ in the above equation for T' to be Nash. That is, this is the unique case in which T' maximizes its payoff by playing itself. Solving for T' yields a T' that is always less than or equal to 0 (regardless of the values of any of the parameters). Since $T = 0$ is the lowest possible amount of deliberation an agent can use, $[0, 0]$, a strategy that never deliberates and always defects, is a unique best response to itself.

For the strategy $[0, 0]$ to be a Nash, however, it must also be the case that no intuitively cooperative strategy can beat it. This follows straightforwardly. The payoff that strategy $[0, 0]$ gets against itself is 0 (since neither player is paying a cost of cooperation to benefit the other or paying a cost to deliberate). Any intuitively cooperative strategy, on the other hand, is going to incur a cooperation cost c on the fraction of interactions that it cooperates. Moreover, since the $[0, 0]$ agent is always defecting, this intuitively cooperative strategy receives no benefit from the $[0, 0]$ agent. Thus, its payoff is always negative and it cannot invade the $[0, 0]$ strategy under any value of p. As a result, $[0, 0]$ (referred to as the "Intuitive defector (ID)" strategy profile in the main text) is always a Nash equilibrium.

2.2.2 Intuitively cooperating equilibria

We next consider whether any intuitively cooperative strategy profile is a Nash. Following the procedure used above, we calculate the expression for the payoff that an intuitively cooperative agent with strategy profile $[1, T]$ gets against an intuitively cooperative agent with strategy profile $[1, T']$:

$$\begin{aligned}
\pi([1,T],[1,T']) =& \frac{p(b-c) - (b-c)(p-1))(\frac{T}{d}-1)(\frac{T'}{d}-1)}{d^2} \\
&- \frac{T(\frac{T'}{d}-1)(((y_d-1)(b-c) - by_d)(p-1) - \frac{T}{2} + py_d(b-c))}{d^3} \\
&- \frac{T'(((y_d-1)(b-c) + cy_d)(p-1) + py_d(b-c))(\frac{T}{d}-1)}{d^3} \\
&- \frac{(TT'(\frac{T}{2} + (p-1)((b-c)(y_d-1)^2 - by_d(y_d-1) + cy_d(y_d-1)) - py_d^2(b-c))}{d^2}
\end{aligned}$$

We then find the best-response T_b by solving for when the partial derivative of this expression with respect to T is 0, yielding

$$T_b = \frac{p(b-c)(1-y_d) - cy_d(1-p)}{p(b-c)(1-y_d)^2 - d} d.$$

In other words, an intuitively cooperative agent's best response against another intuitively cooperative agent is to deliberate only in cases where the cost of deliberation is not greater than T_b.

It's important to note that the value of T_b here can become negative. Since an agent cannot deliberate a negative amount, though, this strategy profile collapses into an intuitively cooperative strategy $[1, 0]$ when

$$y_b \leq \frac{p(b-c)}{c + bp - 2cp}.$$

Thus, as described in the main text, the possible intuitively cooperative equilibrium strategy is either a "Dual-Process Cooperator (DC)", which sometimes deliberates, or an "Intuitive Cooperator (IC)", which doesn't deliberate at all, depending on the accuracy and usefulness of deliberation.

In order to test whether the strategy profile $[1, T_b]$ is Nash, we must also consider whether any intuitively defective strategy profile $[0, T']$ can beat it. To do this, we find the intuitively defective strategy profile that is a best response against the optimal intuitively cooperative strategy profile $[1, T_b]$ by solving $\frac{\partial}{\partial T}\pi([0, T], T_b) = 0$ for T. We find that the strategy $[0, T_b']$ is the intuitively defecting strategy that performs best against the intuitively cooperative strategy $[1, T_b]$, where

$$T_b' = \frac{c^2p^2 - c^2p - cd + 3c^2py_d - 2c^2py_d^2 - 3c^2p^2y_d + bcp + cdp + cdy_d + 2c^2p^2y_d^2 - bcp^2 - 3bcpy_d + bdpy_d - 2cdpy_d + 2bcpy_d^2 + 3bcp^2y_d - 2bcp^2y_d^2}{d - bp + cp - bpy_d^2 + cpy_d^2 + 2bpy_d - 2cpy_d}.$$

We then find the conditions under which the optimal intuitively cooperative strategy profile does better against itself than the best response intuitive defecting strategy profile does,

$$\pi([1, T_b], [1, T_b]) \geq \pi([0, T_b'], [1, T_b]) \quad \text{if} \quad y_b > \frac{p(b-c)}{c + bp - 2cp} \tag{1}$$

$$\pi([1, 0], [1, 0]) \geq \pi([0, T_b'], [1, 0]) \quad \text{otherwise.} \tag{2}$$

in order to find out when DC ($[1, T_b]$) or IC ($[1, 0]$) is Nash.

An analysis of these conditions finds that the DC strategy profile $[1, T_b]$ always becomes an equilibrium when it is still the case that $T_b > 0$, so we only need to evaluate equation (1) above. Calculations find that DC becomes an equilibrium when

$$p = \frac{c^2 - bc - bd + 3bcy_d - 3c^2y_d - 2bcy_d^2 + 2c^2y_d^2 + \sqrt{(bc - c^2 + bd - 3bcy_d + 3c^2y_d + 2bcy_d^2 - 2c^2y_d^2)^2 + 4cd(-bc + c^2 + 3bcy_d - 3c^2y_d - 2bcy_d^2 + 2c^2y_d^2)}}{2(-bc + c^2 + 3bcy_d - 3c^2y_d - 2bcy_d^2 + 2c^2y_d^2)}.$$

2.2.3 Summary

The present analyses explored how the accuracy of deliberation (y_d) influences what Nash equilibria are possible in our dual-process model. We find that the same two equilibria from Bear & Rand's (2016) original model are selected here: an intuitive defecting equilibrium that never deliberates (ID) and a dual-process cooperating equilibrium (DC) that sometimes deliberates when the value of deliberation outweighs the cognitive cost. However, the specific threshold of cost that DC agents are willing to deliberate is sensitive to the accuracy of deliberation, and the value of deliberation drops off steeply as y_d decreases. Thus, to a greater extent than what was observed in the original model, DC agents tend to become purely intuitive cooperators, when the optimal amount of deliberation for an intuitive cooperator hits floor at 0.

2.3 Context-sensitive intuition

Now we explore what happens when we fix $y_d = 1$ and allow the accuracy of intuition to vary such that $.5 < y_i < 1$. Recall that, because intuition is now partially sensitive to whether the interaction is a PD or CG, there are two intuitive responses: S_{iD} and S_{iC}, respectively. We thus break the calculation up into three parts, in which the intuitive response is 1) to always defect ($S_{iD} = 0$ and $S_{iC} = 0$); 2) to always cooperate ($S_{iD} = 1$ and $S_{iC} = 1$); or 3) to discriminate such that intuition defects when it believes the interaction is a PD ($S_{iD} = 0$) and cooperates otherwise ($S_{iC} = 1$). We ignore the situation in which intuition cooperates when it believes the interaction is a PD ($S_{iD} = 1$) and defects otherwise ($S_{iC} = 0$), since such a strategy is obviously maladaptive and clearly could never outperform strategies of types (1)-(3).

2.3.1 Intuitively defecting equilibria

Using the same method from section 2.2.1, we first find the intuitively defecting strategy ($S_{iD} = 0$ and $S_{iC} = 0$) that has the best response against another intuitively defecting strategy. Consistent with previous results, we find that this strategy profile never deliberates: $[0, 0, 0]$ (strategy ID from the main text).

To assess whether/when this is an equilibrium, we consider whether any intuitively cooperative ($S_{iD} = 1$ and $S_{iC} = 1$) or intuitively discriminating ($S_{iD} = 0$ and $S_{iC} = 1$) strategies can do better against ID than ID gets against itself. Because $\pi(ID, ID) = 0$, this amounts to asking whether an intuitively cooperative or intuitively discriminating strategy can get a positive payoff against ID.

The intuitively cooperative strategy profile that performs best against ID is $[1, 1, c(1 - p)]$, where $0 < c(1 - p) < d$. Calculations confirm that $\pi([1, 1, c(1 - p)], ID) \leq 0$ for all parameters that meet these conditions, and therefore no intuitively cooperative strategy can beat ID.

The intuitively discriminating strategy profile that performs best against ID is $[0, 1, c(1-p)(1-y)]$, where $0 < c(1 - p)(1 - y) < d$. Once again, calculations confirm that this strategy can never get a positive payoff against ID.

Thus, ID is always a Nash.

2.3.2 Intuitively cooperating equilibria

Once again, the results of this calculation are similar to those from past calculations. As in Bear & Rand (2016), the best intuitively cooperative strategy profile against another intuitively cooperative strategy profile is the "Dual-process Cooperator" (DC), $[1, 1, c(1-p)]$, where $c(1-p)$ is the expected gain from defecting in the PD.

First, we assess when an intuitively defecting strategy can beat DC. The strategy profile $[0, 0, p(b-c)]$ is the best such strategy profile to compete with DC. Consistent with previous results, $\pi([0, 0, p(b-c)], DC) \leq \pi(DC, DC)$ when $p > \frac{c}{b}$. (Note that y_i is irrelevant here since neither strategy relies on intuition's discriminatory capabilities.)

We also must consider whether any intuitively discriminatory strategy can compete with DC, however. We find that the best such strategy profile is $[0, 1, T']$, where

$$T' = (1 - y_i)(c + bp - 2cp).$$

This intuitively discriminating profile gets a lower or equal payoff than DC gets against itself just in case

$$p > \frac{cy_i}{b - c - by_i + 2cy_i} \qquad \text{and} \qquad y_i > \frac{bc - c^2 - bd + cd}{bc - c^2 - bd + 2cd}.$$

Thus, DC is an equilibrium when the following three conditions are met:

$$\begin{aligned} &(1) \qquad p > \frac{c}{b} \\ &(2) \qquad p > \frac{cy_i}{b - c - by_i + 2cy_i} \\ &(3) \qquad y_i > \frac{bc - c^2 - bd + cd}{bc - c^2 - bd + 2cd}. \end{aligned}$$

2.3.3 Intuitively attending equilibria

By supposing that y_i can vary between .5 and 1, this new model allows for the possibility of an additional equilibrium that has an intuition that conditions its cooperative decision on whether the agent intuitively believes the interaction is a PD or a CG. We call such a strategy an intuitively attending strategy, since it attends to the information intuition is giving it. Specifically, an agent with this strategy profile who uses intuition defects when intuition tells it that the interaction is a PD (S_{iD}) and cooperates when intuition tell sit that the interaction is a CG (S_{iC}).

Note that, for this type of strategy, deliberation responds in the same way as intuition, but has higher discrimination accuracy ($y_d > y_i$). Thus, the function of deliberation is to improve an agent's ability to discriminate context, but not qualitatively change the behavior it is exhibiting.

We first consider the best attending strategy profile $[0, 1, T]$ against other attending strategies profiles. This strategy has a deliberation threshold

$$T = \frac{d(1-y_i)(c - cp + bpy_i - cpy_i)}{d - bp + cp - bpy_i^2 + cpy_i^2 + 2bpy_i - 2cpy_i}.$$

This is the strategy we call the "Dual-process Attender" (DA) from the main text.

We next consider when an intuitively defecting strategy $[0, 0, T']$ or an intuitively cooperative strategy $[0, 0, T*]$ can beat DA. The best intuitively defecting strategy has a deliberation threshold

$$T' = \frac{p(b-c)(c - cp - 2cy_i + dy_i + cy_i^2 - cpy_i^2 + 2cpy_i)}{d - bp + cp - bpy_i^2 + cpy_i^2 + 2bpy_i - 2cpy_i}.$$

The best intuitively cooperative strategy against DA is simply DC ($[1, 1, c(1-p)]$).

Because DA (and the intuitively defecting strategy that can potentially beat it) are considerably more complex strategies than we have encountered in earlier calculations, a precise characterization of the conditions of the DA equilibrium is intractable to solve for generic parameters b, c, and d. We, therefore, present calculations from the example in the main text, where $b = 4$, $c = 1$, and $d = 1$. For this set of parameters, DA is an equilibrium just in case

$$\frac{4 - 7y_i + 6y_i^2}{6(-1+y_i)^2} - \frac{1}{6}\sqrt{\frac{4 - 20y_i + 61y_i^2 - 72y_i^3 + 36y_i^4}{-1+y_i)^4}} \leq p \leq \frac{3 - 2y_i}{6(-1+y_i)^2} - \frac{1}{6}\sqrt{\frac{9 - 24y_i + 28y_i^2 - 12y_i^3}{(-1+y_i)^4}}.$$

2.3.4 Summary

In this subsection, we explored what equilibria were present when we allowed intuition to be somewhat sensitive to context, such that $.5 < y_i < 1$. In addition to uncovering the same two equilibria from the original model (ID $[0, 0, 0]$ and DC $[1, 1, c(1-p)]$), we found the emergence of a qualitatively new, intuitively attending equilibrium (DA in the main text). This strategy is an equilibrium for intermediate values of p (see Figure 4 in main text) and uses intuition in the same way it uses deliberation, conditioning its cooperative decision on the information it is given.

3 Risk dominance calculations

Given that we have identified the game's three possible Nash equilibria, we are now interested in identifying when each of these equilibria will be favored by natural selection. For parameters where ID is the only Nash, it is clearly predicted that evolution will lead to ID. When DC or DA becomes Nash, however, ID also remains Nash. Thus knowing when a strategy becomes Nash is not enough to know when selection will favor it.

Risk-dominance, which is a stricter criterion than Nash, has been shown to answer this question in situations in which multiple equilibria exist. In these cases evolution will favor the risk-dominant equilibrium [1].

One Nash x risk-dominates another Nash y when x earns a higher expected payoff than y when there is a 50% chance of playing against either of the two strategies. Or, in population dynamic terms, the risk-dominant strategy profile is the one that fares better in a population where both are equally common. More precisely, x risk-dominates y just in case

$$\frac{1}{2}\pi(x,y)+\frac{1}{2}\pi(x,x)>\frac{1}{2}\pi(y,y)+\frac{1}{2}\pi(y,x).$$

3.1 Imperfect deliberation

We first consider conditions of risk dominance for situations in which the accuracy of deliberation y_d is varied, with y_i fixed at .5. The calculations above confirm that there are only two equilibria in such situations: ID and DC. However, DC becomes IC when deliberation loses its value and the optimal T therefore becomes 0. Thus, there are two risk-dominance calculations to consider: (i) When does DC risk-dominate ID?; and (ii) When does IC risk-dominate ID? Note that we need not consider when IC risk-dominates DC since these two strategies are never both equilibria simultaneously.

Since these calculations are, again, intractable for generic variables b, c, and d, we present results from the example in the main text, where $b = 4$, $c = 1$, and $d = 1$. Under these conditions, we find that DC risk-dominates ID when

$$p>\frac{10y_d^2-18y_d+11-\sqrt{144y_d^4-432y_d^3+480y_d^2-252y_d+73}}{22y_d^2-36y_d+12},$$

and IC risk-dominates ID when $p > .4$.

3.2 Context-sensitive intuition

We now consider conditions of risk dominance for situations in which the accuracy of intuition y_i is varied, with y_d fixed at 1. According to the above calculations, we find that there are three possible equilibria in such situations: ID, DC, and DA. Thus, we need to consider three calculations: (i) When does DC risk-dominate ID?; (ii) When does DA risk-dominate ID?; and (iii) When does DC risk-dominate DA?

We present results from the example in the main text, where $b = 4$, $c = 1$, and $d = 1$. Note that competition between DC and ID are invariant to y_i since neither of these strategies condition their intuitive response on intuition's belief about whether the interaction is a PD or CG. DC simply risk-dominates ID when $p > .3028$. In contrast, the other two conditions are more complex because of DA's dependence on y_i. For this set of parameters, however, DA *always* risk-dominates DC when both are equilibria. Finally, we find that DA risk-dominates ID when

$$p>\frac{7y_i^2-10y_i+6-\sqrt{69y_i^4-180y_i^3+184y_i^2-80y_i+16}}{2(5y_i^2-10y_i+5)}.$$

4 Evolutionary dynamics

4.1 Basic setup

We now turn from Nash calculations to evolutionary dynamics. We study the transmission of strategies through an evolutionary process, which can be interpreted either as genetic evolution or as social learning. In both cases, strategies that earn higher payoffs are more likely to spread in the population, while lower payoff strategies tend to die out. Novel strategies are introduced by mutation in the case of genetic evolution or innovation and experimentation in the case of social learning.

We study a population of N agents evolving via a frequency dependent Moran process with an exponential payoff function [2]. In each generation, one agent is randomly selected to change strategy. With probability u, a mutation occurs and the agent picks a new strategy at random. With probability $(1-u)$, the agent adopts the strategy of another agent j, who is selected from the population with probability proportional to $e^{w\varphi_j}$, where w is the intensity of selection and φ_j is the expected payoff of agent j when interacting with agents picked at random from the population.

4.2 Limit of low mutation calculation method

In the low mutation limit, a mutant either goes to fixation or dies out before another mutant appears. Thus, the population makes transitions between homogeneous states, where all agents use the same strategy. Here the success of a given strategy depends on its ability to invade other strategies, and to resist invasion by other strategies. We use an exact numerical calculation to determine the average frequency of each strategy in the stationary distribution [3, 4, 5].

Let s_i be the frequency of strategy i, with a total of M strategies. We can then assemble a transition matrix between homogeneous states of the system. The transition probability from state i to state j is the product of the probability of a mutant of type j arising ($\frac{1}{M-1}$) and the fixation probability of a single mutant j in a population of i players, $\rho_{i,j}$. The probability of staying in state i is thus $1 - \frac{1}{M-1}\sum_k \rho_{k,i}$, where $\rho_{i,i} = 0$. This transition matrix can then be used to calculate the steady state frequency distribution s^* of strategies:

$$\begin{pmatrix} s_1^* \\ s_2^* \\ \vdots \\ s_M^* \end{pmatrix} = \begin{pmatrix} 1-\Sigma_j \frac{\rho_{j,1}}{M-1} & \frac{\rho_{1,2}}{M-1} & \cdots & \frac{\rho_{1,M}}{M-1} \\ \frac{\rho_{2,1}}{M-1} & 1-\Sigma_j \frac{\rho_{j,2}}{M-1} & \cdots & \frac{\rho_{2,M}}{M-1} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\rho_{M,1}}{M-1} & \frac{\rho_{M,2}}{M-1} & \cdots & 1-\Sigma_j \frac{\rho_{j,M}}{M-1} \end{pmatrix} \begin{pmatrix} s_1^* \\ s_2^* \\ \vdots \\ s_M^* \end{pmatrix}$$

The eigenvector corresponding to the largest eigenvalue (1) of this matrix gives the steady state distribution of the stochastic process.

Note that this method requires discretizing the strategy space, such that there is some finite number of strategies M that agents can select. We consider a strategy space in which: (i) agents' cooperation strategies S_{iD}, S_{iC}, S_{dD}, and S_{dC} are limited to being either 0 (never play the cooperative strategy) or 1 (always play the cooperative strategy); and (ii) agents' maximum cost of deliberation T $(0 \leq T \leq d)$ that they are willing to pay in order to deliberate is rounded to the nearest $\frac{d}{10}$ (so

T is selected from the set $\{0, d/10, 2d/10, \ldots, d\}$). Thus, the strategy space consists of a total of $2*2*2*2*11 = 176$ strategies.

Using the Moran process, the fixation probability $\rho_{B,A}$ (the probability that a single A mutant introduced into a population of B-players will take over) is calculated according to an exponential fitness function. In a population of i A-players and $N-i$ B-players, the fitness of an A-player f_i and B-player g_i are defined as

$$f_i = e^{w(a\pi(A,A)+(1-a)(\frac{i-1}{N-1}\pi(A,A)+\frac{N-i}{N-1}\pi(A,B)))}$$
$$g_i = e^{w(a\pi(B,B)+(1-a)(\frac{i}{N-1}\pi(B,A)+\frac{N-i-1}{N-1}\pi(B,B)))}$$

where $\pi(A,A)$ is the expected payoff of an A-player against an A-player, $\pi(A,B)$ is the expected payoff of an A-player against a B-player, etc.

The fixation probability of a single A-player in a population of B-players can then be calculated as follows:

$$\rho_{B,A} = \frac{1}{1+\sum_{k=1}^{N-1}\prod_{i=1}^{k}\frac{g_i}{f_i}}$$

As shown in Figures S1-S12, these evolutionary calculations are in quantitative agreement with the risk-dominance calculations across p, y_d and y_i values discussed in the main text.

5 Combining imperfect deliberation and sensitive intuition

In the main text, we explore the separate impact of imperfect deliberation and sensitive intuition on the co-evolution of cooperation and cognition. However, we do not consider the evolutionary dynamics that might arise if both deliberation is made to be imperfect and intuition is made to be (partially) sensitive to context. Varying both of these features at once is not analytically tractable, and therefore we cannot assess the exact conditions in which each of the evolutionary stable strategies risk-dominates the others. Nevertheless, we report here the results from steady state calculations in which both the accuracy of deliberation and sensitivity of intuition are varied. Our results suggest that combining imperfect deliberation and sensitive intuition together in evolutionary simulations does not qualitatively change the key results presented in the main text.

5.1 Varying the accuracy of deliberation when intuition is sensitive

First, we consider which strategies evolution favors when we fix the sensitivity of intuition y_i at 0.75 and vary y_d from 0.75 to 1 (recall that we only consider cases in which $y_d > y_i$ to ensure that there is always some ease-flexibility tradeoff). Figure S13 shows the results of steady-state calculations in p x y_d space for the amount of deliberation T (Figure S13a) and the intuitive responses (Figure S13b) using the same b, c, and d values from the main text. Compared to the result in which intuition is completely insensitive to context (Figure 3 from main text), we see that the value of deliberation once again drops precipitously as a function of y_d, leading the dual-process strategy

DC to collapse into a purely intuitively cooperative strategy (IC). Moreover, for low values of p, the purely intuitively selfish strategy ID is selected by evolution.

However, allowing intuition to be somewhat sensitive to context leads to the emergence of the DA strategy for intermediate values of p. Within this DA region, the amount of deliberation increases a function of both p and y_d. This is consistent with the function of DA discussed in the main text: because the value of deliberation for this strategy is to improve the accuracy of discriminating social dilemmas and coordination games, deliberation increases as errors become costlier (a consequence of increasing p) and deliberation becomes better able to correct these errors (a consequence of increasing y_d).

5.2 Varying the sensitivity of intuition when deliberation is imperfect

We next consider which strategies evolution favors when we fix the accuracy of deliberation y_d at the imperfect value of 0.9 and vary the sensitivity of intuition y_i between 0.5 and 0.9 (again maintaining that $y_d > y_i$). Figure S14 shows the results of steady-state calculations in p x y_i space for the amount of deliberation T (Figure S14a) and the intuitive responses (Figure S14b) using the same b, c, and d values from the main text. Compared to the result in which deliberation is perfect (Figure 5 from main text), we observe a qualitatively similar tripartite relationship among the ID, DA, and DC strategies discussed in detail in the main text. Making deliberation imperfect simply compresses the space in which these three strategies are selected by evolution and leads evolution to select the IC strategy for high values of p. Thus, introducing imperfect deliberation reduces the total amount that agents deliberate, but does not change deliberation's function.

References

[1] Michihiro Kandori, George J Mailath, and Rafael Rob. Learning, mutation, and long run equilibria in games. *Econometrica: Journal of the Econometric Society*, pages 29–56, 1993.

[2] Tibor Antal, Martin A Nowak, and Arne Traulsen. Strategy abundance in 2× 2 games for arbitrary mutation rates. *Journal of Theoretical Biology*, 257(2):340–344, 2009.

[3] Drew Fudenberg and Lorens A Imhof. Imitation processes with small mutations. *Journal of Economic Theory*, 131:251–262, 2006.

[4] Lorens A Imhof, Drew Fudenberg, and Martin A Nowak. Evolutionary cycles of cooperation and defection. *Proceedings of the National Academy of Sciences of the United States of America*, 102(31):10797–10800, 2005.

[5] David G Rand and Martin A Nowak. The evolution of antisocial punishment in optional public goods games. *Nature Communications*, 2:434, 2011.

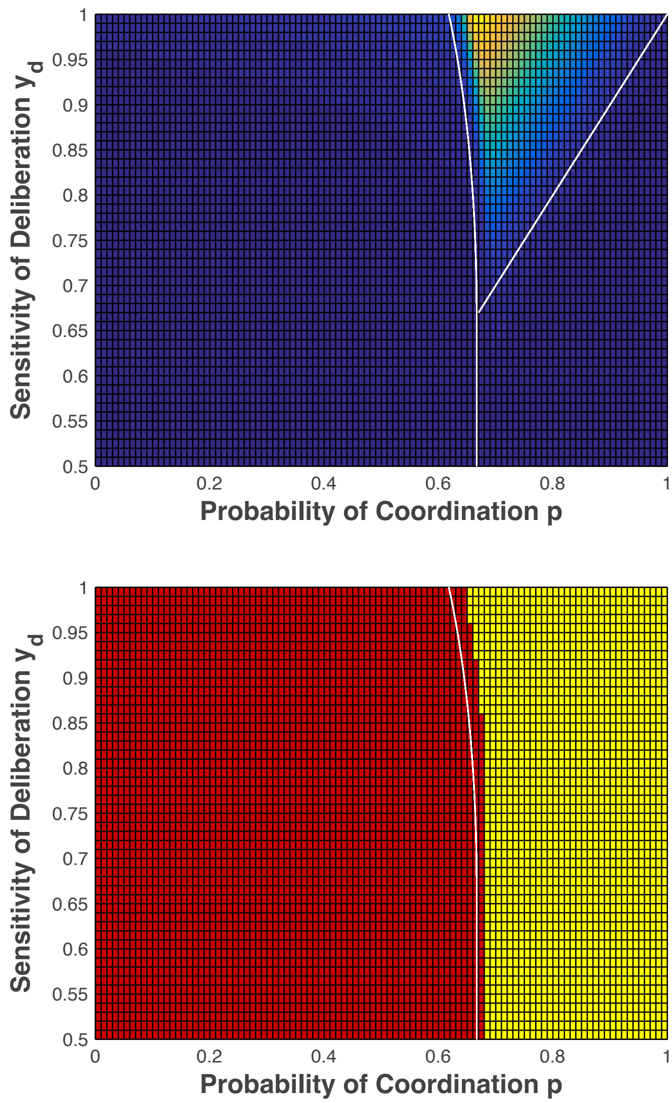

Figure S1: *Evolutionary calculations of the steady state distribution using* $N = 50$*,* $b = 2$*,* $c = 1$*,* $d = 1$*,* $w = 3$*, for various values of* p *and* y_d*. Shown are the average values of (a)* T*, where brighter colors indicate larger* T *values; and (b)* S_i*, rounded to 0 (red) or 1 (yellow). White lines indicate the results of analytic risk-dominance calculations distinguishing ID from DC (left line) and DC from IC (right line).*

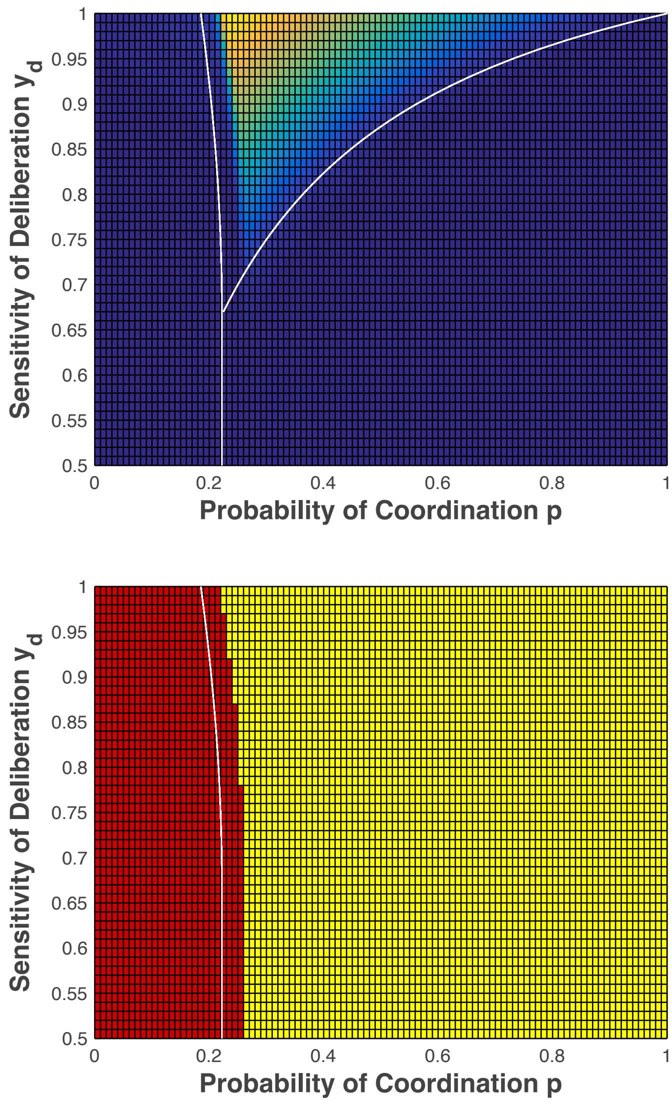

Figure S2: *Evolutionary calculations of the steady state distribution using* $N = 50$, $b = 4$, $c = 0.5$, $d = 1$, $w = 3$, *for various values of* p *and* y_d. *Shown are the average values of (a)* T, *where brighter colors indicate larger* T *values; and (b)* S_i, *rounded to 0 (red) or 1 (yellow). White lines indicate the results of analytic risk-dominance calculations distinguishing ID from DC (left line) and DC from IC (right line).*

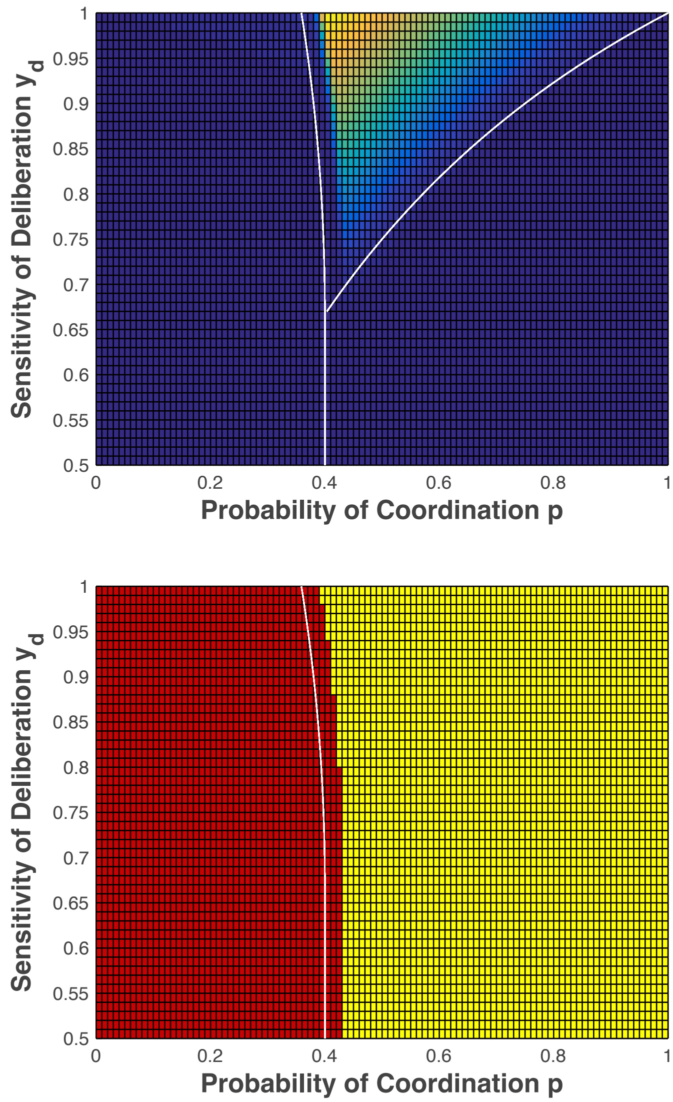

Figure S3: *Evolutionary calculations of the steady state distribution using* $N = 50$, $b = 4$, $c = 1$, $d = 2$, $w = 3$, *for various values of* p *and* y_d. *Shown are the average values of (a)* T, *where brighter colors indicate larger* T *values; and (b)* S_i, *rounded to 0 (red) or 1 (yellow). White lines indicate the results of analytic risk-dominance calculations distinguishing ID from DC (left line) and DC from IC (right line).*

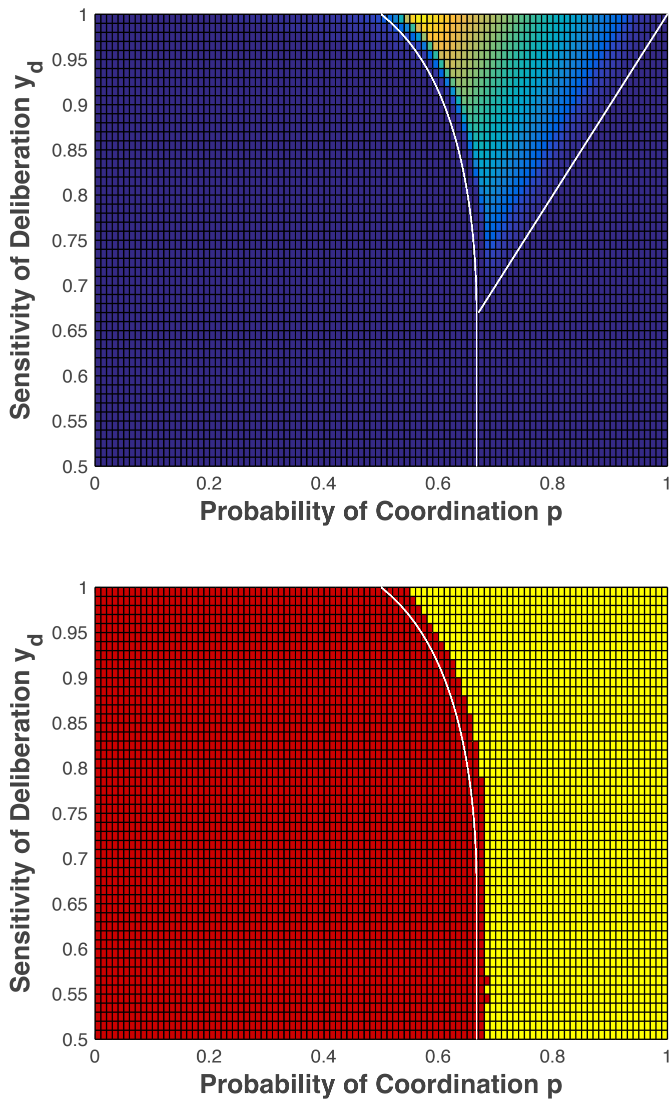

Figure S4: *Evolutionary calculations of the steady state distribution using* $N = 50$*,* $b = 4$*,* $c = 2$*,* $d = 1$*,* $w = 3$*, for various values of* p *and* y_d*. Shown are the average values of (a)* T*, where brighter colors indicate larger* T *values; and (b)* S_i*, rounded to 0 (red) or 1 (yellow). White lines indicate the results of analytic risk-dominance calculations distinguishing ID from DC (left line) and DC from IC (right line).*

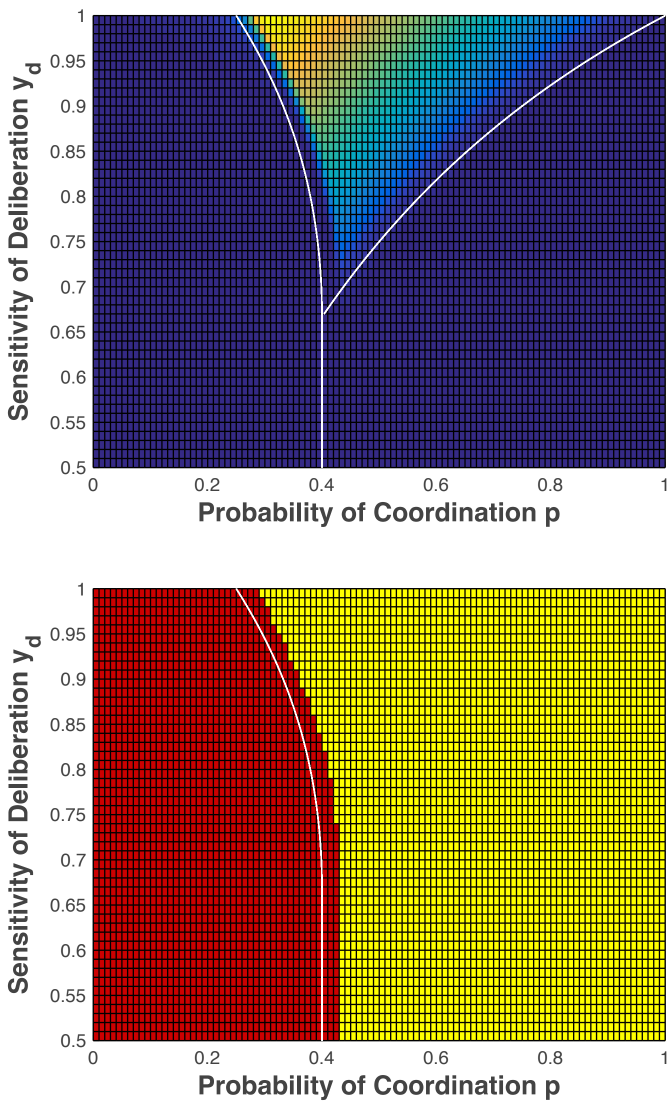

Figure S5: *Evolutionary calculations of the steady state distribution using* $N = 50$, $b = 4$, $c = 1$, $d = 0.75$, $w = 3$, *for various values of* p *and* y_d. *Shown are the average values of (a)* T, *where brighter colors indicate larger* T *values; and (b)* S_i, *rounded to 0 (red) or 1 (yellow). White lines indicate the results of analytic risk-dominance calculations distinguishing ID from DC (left line) and DC from IC (right line).*

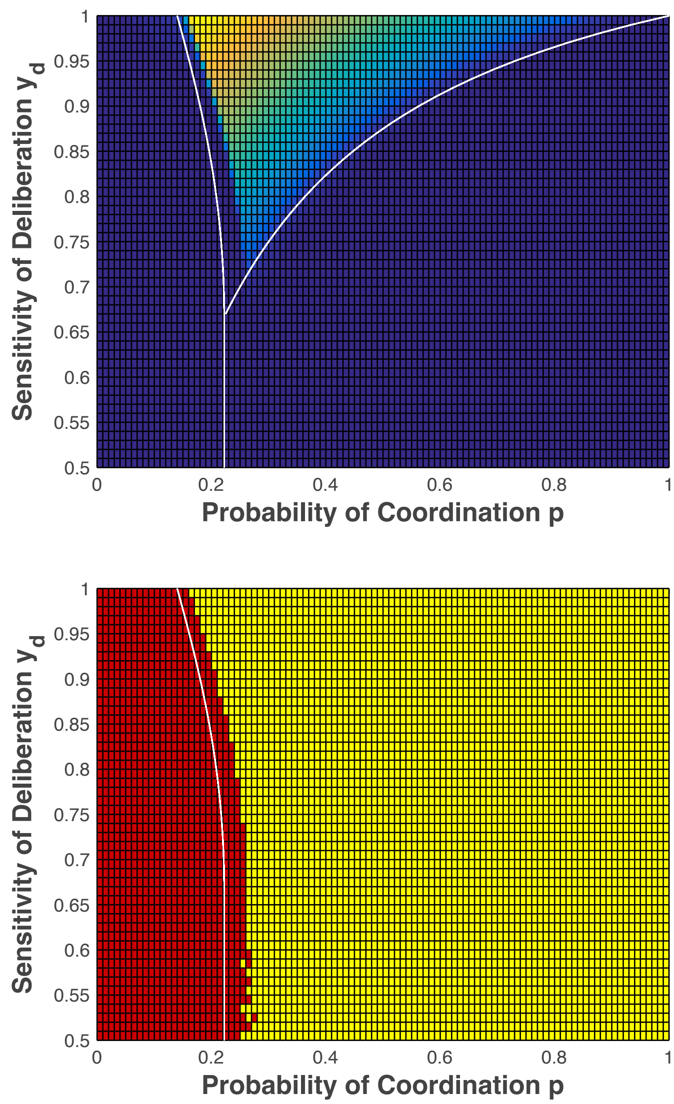

Figure S6: *Evolutionary calculations of the steady state distribution using* $N = 50$, $b = 8$, $c = 1$, $d = 1$, $w = 3$, *for various values of* p *and* y_d. *Shown are the average values of (a)* T, *where brighter colors indicate larger* T *values; and (b)* S_i, *rounded to 0 (red) or 1 (yellow). White lines indicate the results of analytic risk-dominance calculations distinguishing ID from DC (left line) and DC from IC (right line).*

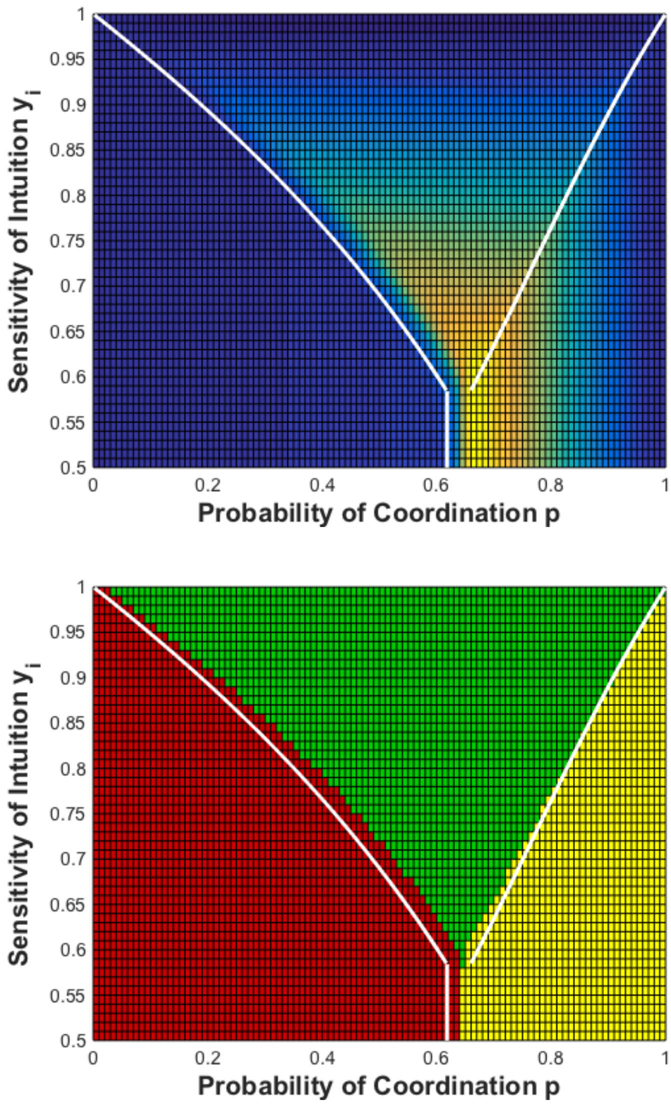

Figure S7: *Evolutionary calculations of the steady state distribution using* $N = 50$*,* $b = 2$*,* $c = 1$*,* $d = 1$*,* $w = 3$*, for various values of* p *and* y_i*. Shown are the average values of (a)* T*, where brighter colors indicate larger* T *values; and (b)* S_{iD} *and* S_{iC}*, both rounded to 0 or 1. In plot (b), red indicates a strategy that always defects when using intuition (*$S_{iD} = S_{iC} = 0$*), yellow a strategy that always cooperates when using intuition (*$S_{iD} = S_{iC} = 1$*), and green a strategy that intuitively discriminates (*$S_{iD} = 0$ *and* $S_{iC} = 1$*). White lines indicate the results of analytic risk-dominance calculations distinguishing ID from DA (top left line), DA from DC (top right line), and ID from DC (bottom line).*

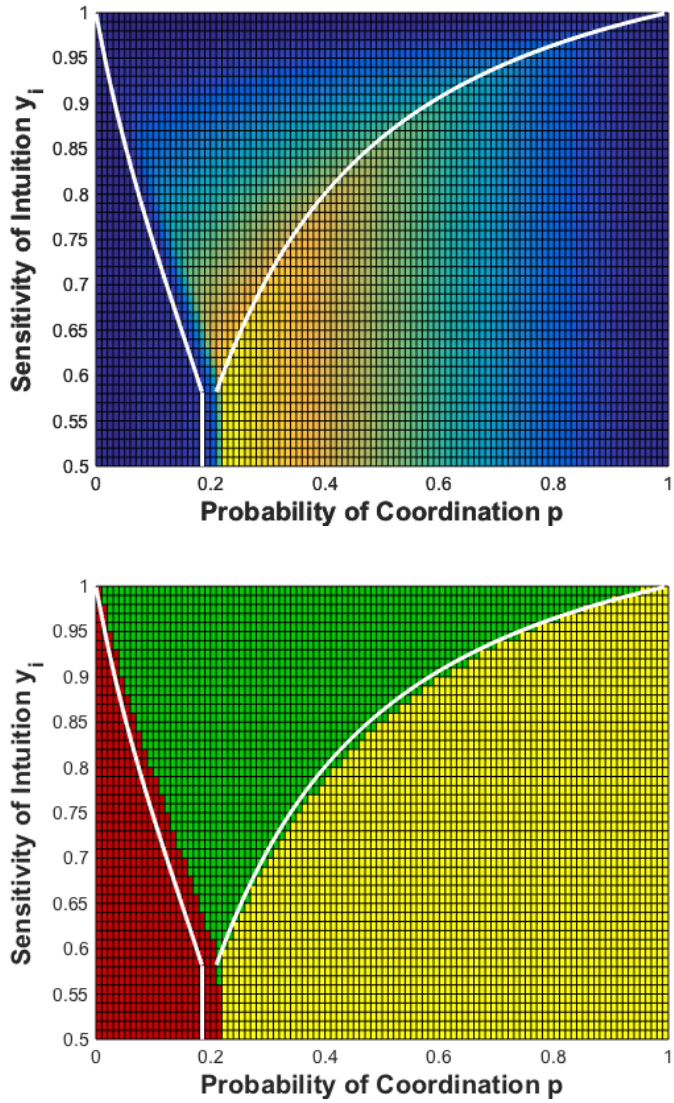

Figure S8: *Evolutionary calculations of the steady state distribution using* $N = 50$, $b = 4$, $c = 0.5$, $d = 1$, $w = 3$, *for various values of* p *and* y_i. *Shown are the average values of (a)* T, *where brighter colors indicate larger* T *values; and (b)* S_{iD} *and* S_{iC}, *both rounded to 0 or 1. In plot (b), red indicates a strategy that always defects when using intuition* ($S_{iD} = S_{iC} = 0$), *yellow a strategy that always cooperates when using intuition* ($S_{iD} = S_{iC} = 1$), *and green a strategy that intuitively discriminates* ($S_{iD} = 0$ *and* $S_{iC} = 1$). *White lines indicate the results of analytic risk-dominance calculations distinguishing ID from DA (top left line), DA from DC (top right line), and ID from DC (bottom line).*

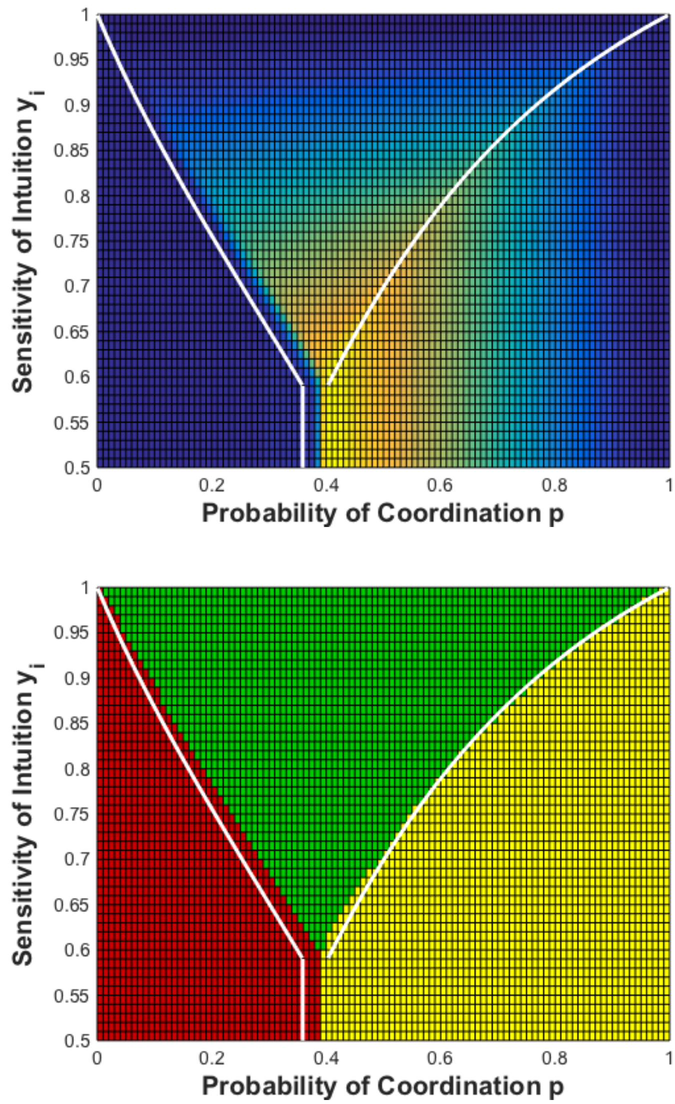

Figure S9: *Evolutionary calculations of the steady state distribution using* $N = 50$, $b = 4$, $c = 1$, $d = 2$, $w = 3$, *for various values of* p *and* y_i. *Shown are the average values of (a)* T, *where brighter colors indicate larger* T *values; and (b)* S_{iD} *and* S_{iC}, *both rounded to 0 or 1. In plot (b), red indicates a strategy that always defects when using intuition (*$S_{iD} = S_{iC} = 0$*), yellow a strategy that always cooperates when using intuition (*$S_{iD} = S_{iC} = 1$*), and green a strategy that intuitively discriminates (*$S_{iD} = 0$ *and* $S_{iC} = 1$*). White lines indicate the results of analytic risk-dominance calculations distinguishing ID from DA (top left line), DA from DC (top right line), and ID from DC (bottom line).*

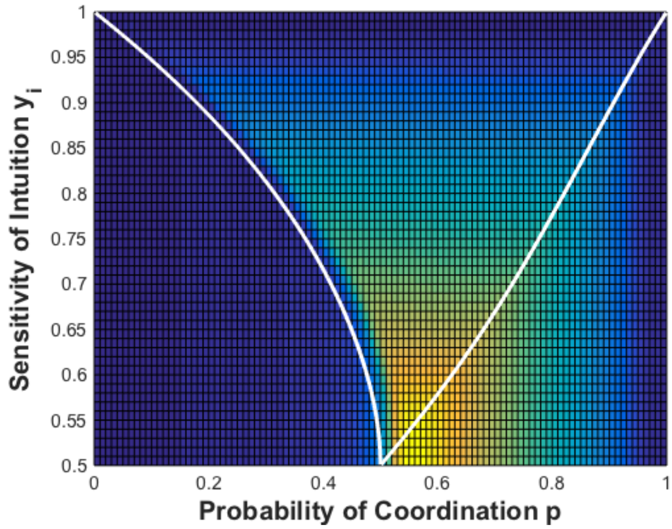

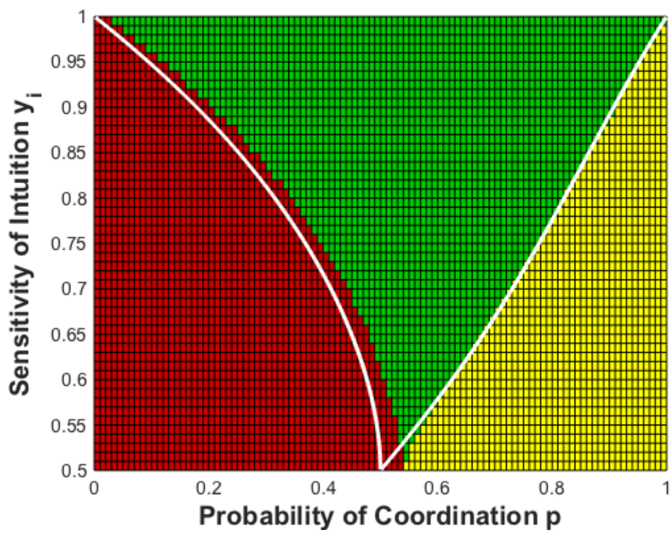

Figure S10: *Evolutionary calculations of the steady state distribution using* $N = 50$*,* $b = 4$*,* $c = 2$*,* $d = 1$*,* $w = 3$*, for various values of* p *and* y_i*. Shown are the average values of (a)* T*, where brighter colors indicate larger* T *values; and (b)* S_{iD} *and* S_{iC}*, both rounded to 0 or 1. In plot (b), red indicates a strategy that always defects when using intuition (*$S_{iD} = S_{iC} = 0$*), yellow a strategy that always cooperates when using intuition (*$S_{iD} = S_{iC} = 1$*), and green a strategy that intuitively discriminates (*$S_{iD} = 0$ *and* $S_{iC} = 1$*). White lines indicate the results of analytic risk-dominance calculations distinguishing ID from DA (top left line) and DA from DC (top right line).*

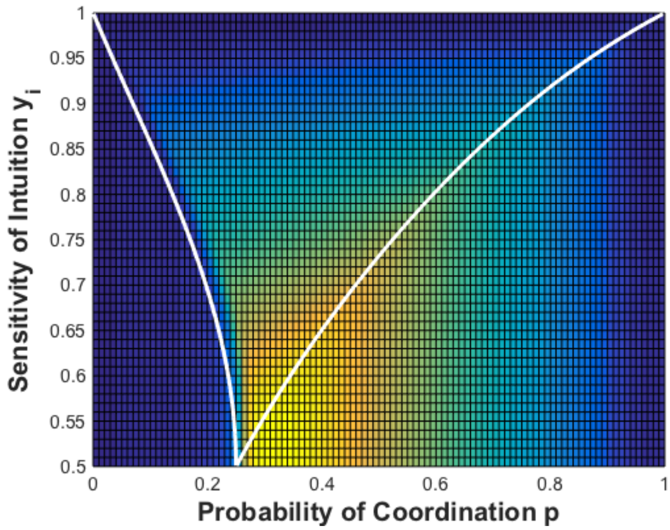

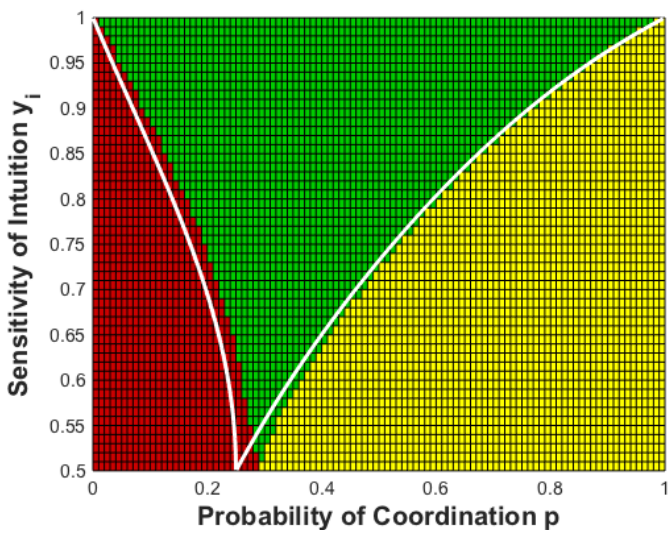

Figure S11: *Evolutionary calculations of the steady state distribution using* $N = 50$, $b = 4$, $c = 1$, $d = 0.75$, $w = 3$, *for various values of* p *and* y_i. *Shown are the average values of (a)* T, *where brighter colors indicate larger* T *values; and (b)* S_{iD} *and* S_{iC}, *both rounded to 0 or 1. In plot (b), red indicates a strategy that always defects when using intuition (*$S_{iD} = S_{iC} = 0$*), yellow a strategy that always cooperates when using intuition (*$S_{iD} = S_{iC} = 1$*), and green a strategy that intuitively discriminates (*$S_{iD} = 0$ *and* $S_{iC} = 1$*). White lines indicate the results of analytic risk-dominance calculations distinguishing ID from DA (top left line) and DA from DC (top right line).*

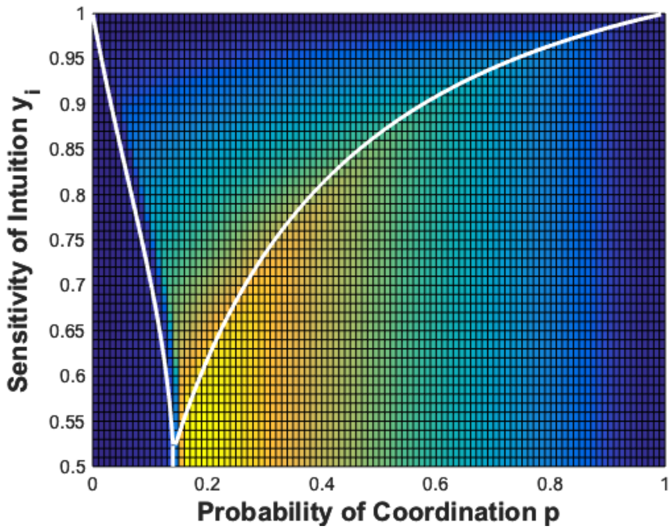

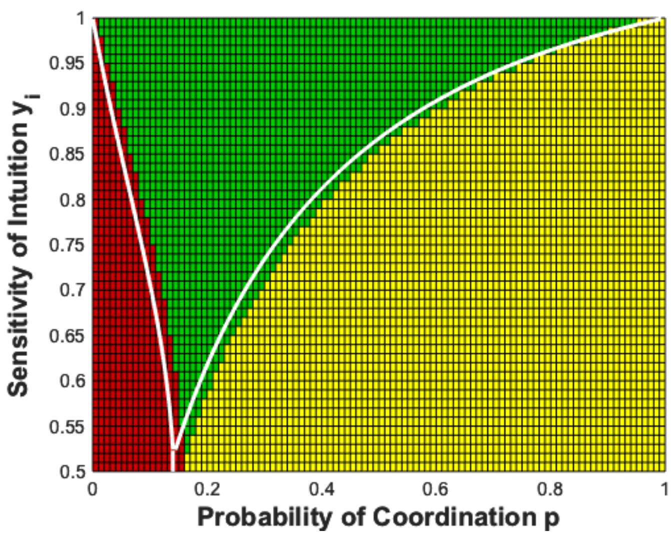

Figure S12: *Evolutionary calculations of the steady state distribution using $N = 50$, $b = 8$, $c = 1$, $d = 1$, $w = 3$, for various values of p and y_i. Shown are the average values of (a) T, where brighter colors indicate larger T values; and (b) S_{iD} and S_{iC}, both rounded to 0 or 1. In plot (b), red indicates a strategy that always defects when using intuition ($S_{iD} = S_{iC} = 0$), yellow a strategy that always cooperates when using intuition ($S_{iD} = S_{iC} = 1$), and green a strategy that intuitively discriminates ($S_{iD} = 0$ and $S_{iC} = 1$). White lines indicate the results of analytic risk-dominance calculations distinguishing ID from DA (top left line), DA from DC (top right line), and ID from DC (bottom line).*

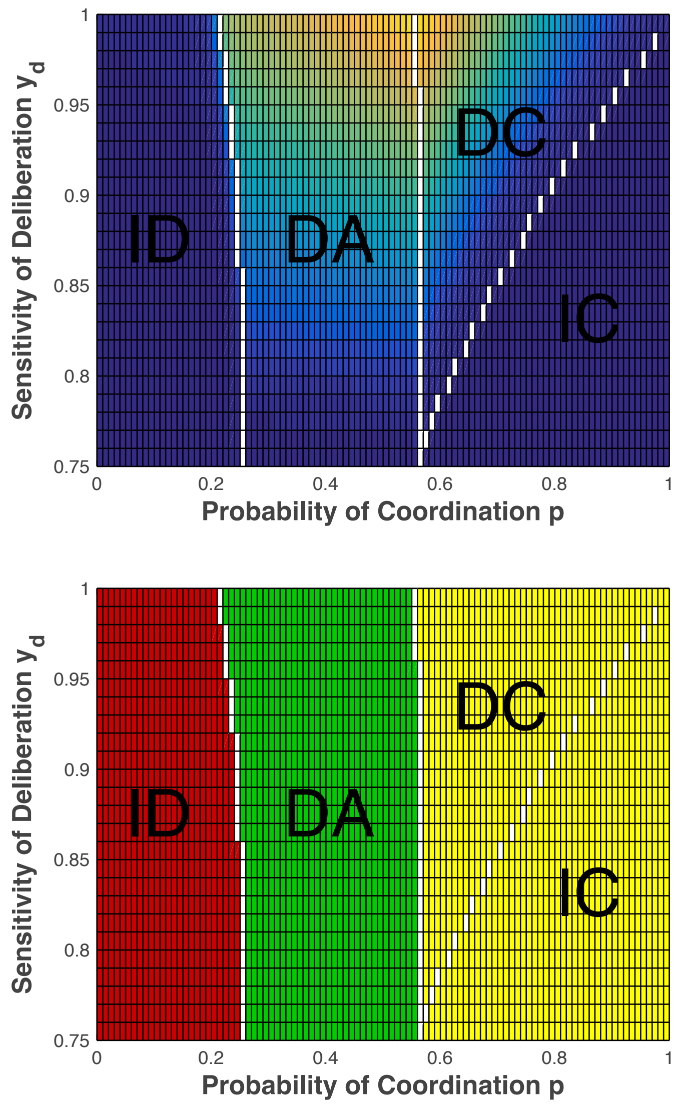

Figure S13: *Evolutionary calculations of the steady state distribution using* $N = 50$, $b = 4$, $c = 1$, $d = 1$, $w = 3$, *and* $y_i = .75$ *for various values of* p *and* y_d. *Shown are the average values of (a)* T, *where brighter colors indicate larger* T *values; and (b)* S_{iD} *and* S_{iC}, *both rounded to 0 or 1. In plot (b), red indicates a strategy that always defects when using intuition* ($S_{iD} = S_{iC} = 0$), *yellow a strategy that always cooperates when using intuition* ($S_{iD} = S_{iC} = 1$), *and green a strategy that intuitively discriminates* ($S_{iD} = 0$ *and* $S_{iC} = 1$). *White lines indicate approximate boundaries between the 4 equilibrium strategies ID, DA, DC, and IC, calculated based on the steady-state results.*

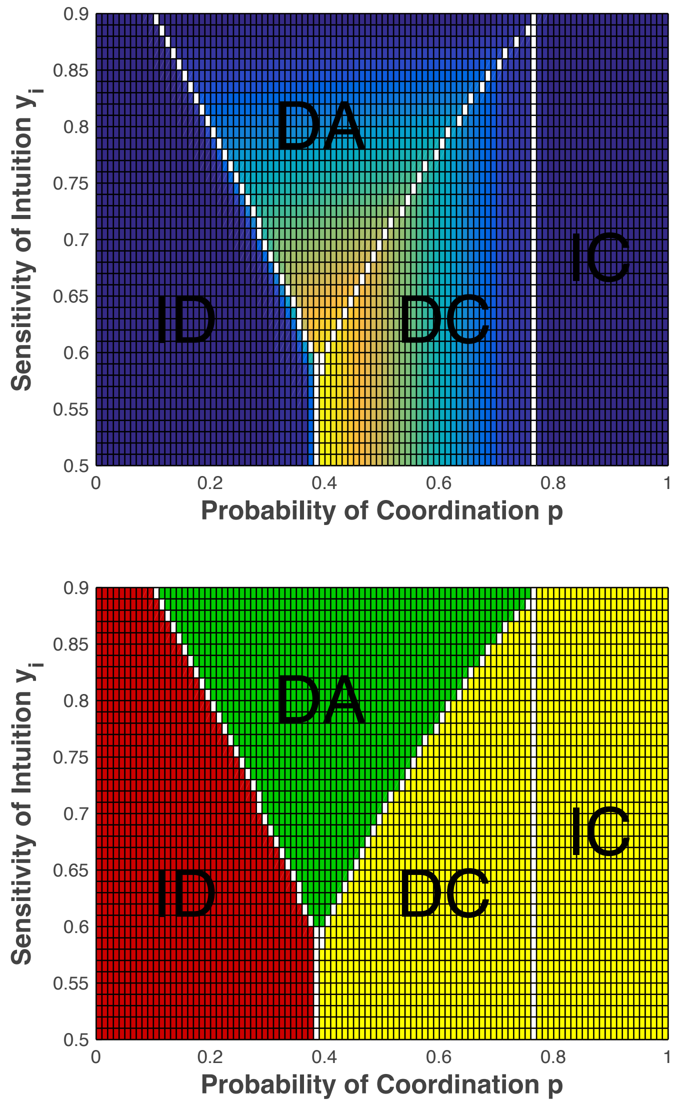

Figure S14: *Evolutionary calculations of the steady state distribution using* $N = 50$, $b = 4$, $c = 1$, $d = 1$, $w = 3$, *and* $y_d = .9$ *for various values of* p *and* y_i. *Shown are the average values of (a)* T, *where brighter colors indicate larger* T *values; and (b)* S_{iD} *and* S_{iC}, *both rounded to 0 or 1. In plot (b), red indicates a strategy that always defects when using intuition* ($S_{iD} = S_{iC} = 0$), *yellow a strategy that always cooperates when using intuition* ($S_{iD} = S_{iC} = 1$), *and green a strategy that intuitively discriminates* ($S_{iD} = 0$ *and* $S_{iC} = 1$). *White lines indicate approximate boundaries between the 4 equilibrium strategies ID, DA, DC, and IC, calculated based on the steady-state results.*