



Rational Coordination Without Beliefs

Camilla Colombo¹ · Francesco Guala² 

Received: 3 February 2021 / Accepted: 7 November 2021 / Published online: 24 November 2021
© The Author(s), under exclusive licence to Springer Nature B.V. 2021

Abstract

Can rational agents coordinate in simultaneous interactions? According to standard game theory they cannot, even if there is a uniquely best way of doing so. To solve this problem we propose an argument in favor of ‘belief-less reasoning’, a mode of inference that leads to converge on the optimal solution ignoring the beliefs of the other players. We argue that belief-less reasoning is supported by a commonsensical Principle of Relevant Information that every theory of rational decision must satisfy. We show that this principle can be used to justify (some versions of) team reasoning, as well as other schemes of practical reasoning that do not involve sophisticated meta-representation.

1 Introduction

My partner and I have no way of concerting our choices. There must be some way, however, so let’s look for it.

Thomas Schelling

We coordinate all the time: when we queue at the bus stop, when we drive our cars, when we buy food at the grocer’s, or when we shake hands with a stranger. We do it so frequently and smoothly, in fact, that we hardly realize that coordination may be problematic. And yet, according to the best theories of action that the social sciences have conceived, coordination is puzzling. Rational choice and game theory are unable to account for coordination in situations that do not seem to bother ordinary humans at all.

✉ Francesco Guala
francesco.guala@unimi.it

Camilla Colombo
camillacolombo@gmail.com

¹ Scuola IMT Alti Studi Lucca, Piazza S. Francesco 19, 55100 Lucca, Italy

² Department of Philosophy and PhiLab, University of Milan, Via Festa del Perdono 7, 20136 Milano, Italy

The most striking case, which will be discussed in this paper, is a simultaneous coordination game with symmetric payoffs and a unique Pareto-optimal equilibrium called ‘Hi-lo’. Hi-lo is surprising because most people consider the choice of the optimal equilibrium absolutely obvious. Choosing any other profile of strategies seems silly, and in fact extremely high rates of successful coordination are observed in experiments with this game. And yet, the justification of this behavior is far from trivial: *any* strategy can be rationalized by some configuration of players’ beliefs, no set of beliefs is mandatory, and therefore no strategy can be ruled out as irrational a priori.

This surprising conclusion follows from the application of standard principles of rational choice, in particular the assumption that each player chooses her best response to the expected actions of the other players. This assumption has not gone unchallenged, to be sure: in his seminal analysis of tacit bargaining Thomas Schelling claimed that ‘the intellectual processes of choosing a strategy in pure conflict and choosing a strategy of coordination are of wholly different sorts’ (1960: 96), and that best-response reasoning is appropriate only in the former case. In *The Strategy of Conflict* however Schelling did not provide a full-fledged theory of coordination. Instead, he offered a number of examples centered around the notions of salience and focal point. Among other things, he argued that salience is *intentionally* used by players to coordinate:

What is necessary [for the players] is to coordinate predictions, to read the same message in the common situation, to identify the one course of action that their expectations of one another can converge on. They must ‘mutually recognize’ some unique signal that coordinates their expectations of each other (Schelling, 1960: 54).

As noticed by Sugden and Zamarròn (2006), expressions such as ‘finding the key’ (or the ‘clue’, or ‘solving the riddle’) recur frequently in *The Strategy of Conflict*, and not just metaphorically or for illustrative purposes. Coordinating players are portrayed by Schelling as goal-driven, intentional agents who try to find the solution of a puzzle. In this paper we will try to develop Schelling’s idea that people attain coordination by looking for a way to sidestep the circularity problem. The ‘key’ or ‘signal’ in games such as Hi-lo is obvious enough, for a single profile of strategies is better than any other profile for all the players involved. Because it is optimal, the choice of this equilibrium is also perceived as rational by most people who happen to play this game. We shall argue that this perception is warranted: convergence on the optimal equilibrium may be backed up by an inferential scheme called ‘belief-less reasoning’ that is rationally justifiable even though it departs significantly from best-response reasoning.

The core of our argument is that in simultaneous coordination games people have good reasons for being non-strategic. Thinking about people’s beliefs, even if it is done competently and systematically, does not lead anywhere in these games, while reasoning about the structure of preferences does. The justification of belief-less reasoning is provided by a Principle of Relevant Information that is implicitly

Fig. 1 A simple coordination game: Hi-lo. The preference ordering in the (b) matrix is $a > b > c$



presupposed by every theory of rational decision-making, and that prescribes to ignore other players’ beliefs in games such as Hi-lo.

The paper is organized as follows: Sect. 2 introduces Hi-lo and explains why it constitutes a puzzle for the standard theory based on best-response reasoning. Section 3 discusses the theory of team reasoning, and explains why it offers a partial solution to the puzzle. Belief-less reasoning is introduced in Sect. 4, and a rational justification based on the Principle of Relevant Information is offered in Sect. 5. In these sections we also show that some versions of team reasoning are structurally similar to belief-less reasoning, and we try to provide a unified approach to the solution of coordination problems. The rest of the paper is devoted to some objections against belief-less reasoning, focusing in particular on charges of irrationality (Sect. 6) and on its domain of application (Sect. 7). In the last section we summarize and conclude the argument.

2 The Problem

Consider the following situation: Ann and Bob are a young married couple with children. When they come home from work, they must pick up the kids at the nursery and shop for dinner. Let us suppose that Ann’s office is closer to the superstore, and Bob’s office is closer to the nursery. There are four possible combinations of actions: Ann picks up the kids and Bob goes shopping; Ann goes shopping and Bob picks up the kids; both go shopping and no one picks up the kids; both pick up the kids and no one goes shopping. The latter two solutions are obviously the least preferred ones. Of the two preferred solutions, however, one is clearly better, given Ann and Bob’s respective locations.

The possible combinations of Ann and Bob’s actions, as well as their outcomes, are represented in the matrix of Fig. 1a. As customary, we represent the actions or strategies of the players as the rows and columns of the matrix. The cells are the outcomes, and the numbers represent the preference orderings of the two players (the first number for the row player, the second one for column).

Now suppose for the sake of the example that Anne and Bob have not made any preliminary arrangement. As she is driving home, Ann realizes that she has left her mobile phone in her office, and she has no way of communicating with her partner. Should she go to the nursery school or to the superstore? Bob in the meantime is asking himself the same question.¹

¹ Although we are assuming the impossibility of communication, sending messages would not change the strategic nature of the problem in a significant way unless it is supported by social norms that modify

Ann and Bob's predicament is an instance of an interactive situation known in the literature as the Hi-lo game.² Hi-lo is a coordination game with two Nash equilibria in pure strategies, one of which (HH) Pareto-dominates every other outcome.³ Its abstract form is summarized in Fig. 1b. The labels of the actions are H (for High) and L (for Low), while the payoffs are ordered as follows: $a > b > c$. When they see Hi-lo for the first time, the overwhelming majority of people does not find it problematic at all. It seems obvious that Ann and Bob must coordinate by choosing H. In fact, it is more than obvious – it seems eminently *rational*, in a situation of this kind.

And yet, according to standard game theory, there is a problem of coordination. Both HH and LL are possible rational solutions, and there is no way to prove that a rational player should choose one instead of the other strategy. The reason is the following: even though HH is better for both players, LL is a Nash equilibrium, and hence a *rationalizable* outcome of the game.⁴ It can be rationalized, for instance, by assuming that each player believes that the other player will choose L – in which case L is the optimal action (the best response) for both.

Rationalizable is not the same as *rationalized*, of course. How could such beliefs be justified? Why should a player believe that the other will choose L (or H, for that matter)? Standard game theory has little to say about this. Two rational agents are trapped in a circularity problem: the identification of the optimal response requires the prior identification of the action (and hence the beliefs) of the other player. But the beliefs of the other player cannot be determined because she is in a symmetric position: her beliefs depend on the beliefs of the first player. Since both are aware of this mutual dependence, no belief and hence no choice can be ruled out as irrational, or justified as uniquely rational, a priori.

3 Solutions

Various theories have been proposed to tackle this issue. Some theorists have introduced extra criteria of equilibrium selection, such as Harsanyi and Selten's (1988) Payoff Dominance Principle, that prescribe choosing the outcome that is uniquely best for each player (if it exists). Such principles however are usually considered

Footnote 1 (continued)

the payoffs of the game (classic statements of this view can be found in Lewis, 1969 and Aumann, 1990). Communication may of course be used to create a focal point, but the interpretation of the message would constitute a higher-order coordination game, and the equilibrium selection problem would remain intact from a purely logical point of view.

² Hi-lo has been named differently by different authors. Early discussions can be found in Schelling (1960), Hodgson (1967) and Lewis (1969). The 'Hi-lo' label was introduced by Michael Bacharach, who is mostly responsible for the current revival of interest in this game (see in particular Bacharach 2006).

³ In a Nash equilibrium the strategy of each player is a best response to the strategy of every other player—or, equivalently, no player has a unilateral incentive to deviate. An outcome is Pareto optimal (or efficient) if and only if it is not (Pareto-)dominated by any other outcome, that is, if there is no other outcome that is better for some players and is not worse for any of them.

⁴ A strategy or set of strategies is rationalizable if it constitutes a best response to a rationally permissible belief set. See Spohn (1982), Bernheim (1984), Pearce (1984).

provisional and unwarranted.⁵ The problem of finding a justification is essentially the one we are tackling in this paper.

Others have conjectured that players' meta-representation capacities—their ability to represent higher-order beliefs—may be limited (Camerer et al., 2004; Lewis, 1969; Nagel, 1995; Stahl & Wilson, 1995),⁶ or that players make mistakes with a probability that is proportional to the payoffs of the game (Bach & Perea, 2014). Although they successfully explain convergence on the optimal equilibrium, none of these approaches provides a *rational justification* for this behaviour (for an overview see e.g. Gold & Colman, 2018).

A third, less orthodox approach is to allow the players to deliberate about the choice of *outcomes*, instead of individual actions, departing significantly from the logic of best-reply reasoning. The best-known theory of this kind is *team reasoning*, an approach proposed independently by several scholars but developed in detail by Michael Bacharach (1999, 2006), Robert Sugden (1993, 2000, 2003, 2018) and Natalie Gold (Gold, 2012, 2018; Gold & Sugden, 2007).⁷ Team reasoning moves from the insight that the players may see Hi-lo as a collective task rather than a problem to be solved individually: that they 'think as a group' or a 'team' rather than as individual agents. This shift in the unit of agency causes a transformation of the option space, for the team is allowed to choose among entire profiles of strategies, under the assumption that the members are going to implement the team strategy. In the Hi-lo game, obviously, the interests of the team are best served by converging upon the Pareto-dominant equilibrium.

An advantage of this approach is that the choice of HH is sanctioned as *rational for the team*, once the game is perceived as a collective problem. The disadvantage is that this perspective calls for a separate argument or justification: is it rational to engage in group-think? Susan Hurley (1989, 2005) has argued that units of agency can be the objects of instrumentally rational choice. Coordination, in her account, involves a two-step procedure: first individuals decide to think as a team, and then they follow the logic of team reasoning by choosing H. The problem with this story is that it is not clear how the first step can be justified. Each individual may try to defend the decision of joining the team by appealing to the superior consequences of team reasoning. But such consequences are conditional on the other players' adoption of team reasoning, and one cannot be assured that they will do that. The problem, in other words, is that the decision to engage in team reasoning constitutes another coordination problem with multiple equilibria, which must be solved prior to Hi-lo.⁸

⁵ Harsanyi and Selten for example propose payoff dominance only as a temporary solution to the problem of equilibrium selection, until a more satisfactory one is found Harsanyi and Selten (1988: 357–9).

⁶ If some players have a 'brute propensity' to choose the obvious strategy (H, in this case), then any other player with slightly more sophisticated cognitive capacities has an incentive to converge on the Pareto-optimal outcome.

⁷ Precursors or relatives of team reasoning can be found in the writings of Hodgson (1967), Gauthier (1975), Hurley (1989), and Gilbert (1989) among others.

⁸ See the discussion in Browne (2018), who also criticizes Hurley's attempt to use mind-reading as a regress stopper.

Other theorists have claimed that the question is misconceived, for rationality *presupposes* a unit of agency. According to Bacharach (2006), for example, the switch from individual to team reasoning is a non-rational framing effect, so one must only be confident that the other players are sensitive to the psychological mechanisms that trigger the change of frame.⁹ In his recent writings Sugden (2011, 2015, 2018) follows a different route, highlighting the role played by existing practices (regularities of behavior) in sustaining reciprocal expectations of team reasoning. But Sugden is careful to say that a practice needs not be optimal—only satisficing or mutually beneficial with respect to some benchmark—and that there is no reason why the participants ought to endorse it. We interpret these remarks as indicating that he does not rule out LL as a possible outcome of team reasoning, and that the adoption of team reasoning is not rationally sanctioned. In conclusion, both Bacharach and Sugden claim that the meta-coordination problem is solved outside of the realm of rationality.¹⁰

In what follows we will try to argue that this conclusion is too hasty. Although the mechanisms highlighted by Bacharach and Sugden are plausible from a descriptive point of view, the normative status of team reasoning and analogous theories of coordination deserves further attention. Perhaps there is a way of justifying convergence on the optimal equilibrium of Hi-lo within the boundaries of practical rationality. In the next few sections we will try to develop Schelling's insight that people look for a way to break the circle of reasoning that prevents coordination in such games. The search for an alternative mode of inference, we shall argue, is rational, given the situation the agents are in.

4 Belief-Less Reasoning

The starting point for any rational solution to problems of coordination must be the recognition that reasoning about higher-order beliefs does not work. As noticed by Sugden (1993: 87) 'it is because players who think as a team do not need to form expectations about one another's actions that they can solve coordination problems'. Karpus and Radzvilas (2018) have recently sketched a process that may lead to the adoption of team reasoning, following a similar line of thought:

decision-makers, who first approach games from the point of view of best-response reasoning, may switch to considering which outcomes of games are mutually advantageous when best-response reasoning is unable to resolve their decision problems definitively. The decision-makers' subsequent endorsement of team reasoning to guide their actions can depend on their beliefs about its endorsement by others as well as the outcomes they can expect to attain from

⁹ See also Sugden (1993, 2000, 2003), Gold and Sugden (2007).

¹⁰ Another strategy is to argue that group identification is justified by normative considerations (such as commitments—cf. e.g. Gilbert 1990, Gauthier 2013). Units of agency, again, cannot be rationally chosen.

the application of best-response considerations, and the first of these factors may depend on the second. (Karpus & Radzvilas, 2018: 25)

This process involves an inversion of standard game-theoretic reasoning, as highlighted at the end of the paragraph ('the first of these factors may depend on the second'). The key insight is that individual expectations may be derived from an evaluation of the outcomes that team reasoning may deliver compared to those delivered by standard best-response reasoning. Expectations thus are not an input of the reasoning process, as in standard game theory, but one of its outputs. We expect everyone to think as a team because we (and they) can see that this mode of reasoning succeeds where others fail.

Karpus and Radzvilas do not claim that team reasoning is rational, to be sure. To see that a purely pragmatic argument does not necessarily provide a rational justification, it is sufficient to notice that various irrational or boundedly rational decision processes are occasionally superior, in practical terms, to best-response reasoning. (One may simply help herself with arbitrary assumptions about the beliefs of the other players, for example, and it may work.) But the fact that they are successful does not make such schemes of inference less fallacious—after all, why should we expect rationality to always serve us well?

In this paper we would like to go one step beyond Karpus and Radzvilas' analysis, and show that the problem of belief circularity can be bypassed using a mode of reasoning that is rationally sanctioned. We shall call such a scheme of inference 'belief-less reasoning'.¹¹ We shall argue that some versions of team reasoning, as well as similar theories proposed in the literature, can be seen as particular cases of belief-less reasoning. The key idea of belief-less reasoning is that, instead of trying to predict the actions of the other players from their preferences *and beliefs*, the players try to identify what is objectively the best or most obvious way to coordinate, using only their preference rankings about the outcomes. The beliefs of the other players do not play a significant role in the inferential scheme, although in principle they may be derived from its conclusion.

Seen from the point of view of an individual player, a process of belief-less reasoning for the Hi-lo game may be reconstructed as follows¹²: (BR)

1. My goal is to maximize my payoff and your goal is to maximize your payoff.
2. The best way for you and me to achieve my goal and your goal is that I choose H and you choose H.
3. I will choose H and you will choose H.

The premises (1–2) include information about the preferences of the players, and the actions or strategies that must be implemented by each player in order to satisfy them. Although BR is formulated in 'individualistic' mode, it can be slightly modified to obtain 'collectivistic' inferential schemes such as team reasoning. A strong

¹¹ On schemes of practical reasoning, see Bacharach (2006), Gold and Sugden (2007).

¹² See Guala (2020).

version of team reasoning such as Bacharach's, for example, would look like this: (TR)

1. The team's goal is to maximize its payoffs.
2. The best way to achieve this goal is that we, as team members, choose HH.
3. I will choose H and you will choose H.

There are versions of team reasoning that do not put a strong emphasis on the transformation of the unit of agency, but merely require that the players recognize the existence of a 'mutually beneficial' outcome (e.g. Sugden, 2011, 2015, 2018). The idea of mutual benefit may be cashed out in different ways, but in Hi-lo, where HH is preferred by both players, convergence may be achieved as follows: (TR*)

1. My goal is to attain a mutually beneficial outcome and your goal is to attain a mutually beneficial outcome.
2. The best way for you and me to achieve my goal and your goal is that I choose H and you choose H.
3. I will choose H and you will choose H.

Although BR, TR, and TR* differ slightly with respect to their content, they share the same pattern of reasoning.¹³ This suggests that the essential feature of this reasoning mode is the way in which goals, means, and beliefs are arranged in the inferential scheme. Another theory of coordination that displays the same pattern is Adam Morton's 'solution thinking': when solving a problem of coordination,

One first thinks of an outcome which one can imagine the other person or persons both would want to achieve and would believe that one would try to achieve. One then thinks out a sequence of actions by all concerned that will lead to it. Lastly, one performs the actions that fall to one's account from this sequence [...] and expects the other(s) to do their corresponding actions. [...]
(Morton, 2003: 120)

Again, the reasoning starts from the identification of a goal and of a set of actions that leads to it. Since expectations figure only—and inessentially—among the conclusions of the inference, we take Morton's solution thinking to be a form of belief-less reasoning.

From now on we shall focus on BR as the core argumentative pattern that these approaches have in common. Clearly BR falls under the umbrella of instrumental (means-ends) rationality. The question we want to ask is whether it is a sound inferential scheme or not. In order to answer we must tighten the argument, make some hidden assumptions explicit, and examine their normative basis.

¹³ It has been argued, in fact, that 'team reasoning is simply a matter of using certain patterns of inference' (Pacherie, 2013: 1834). Since we are not interested in exegetical issues we remain neutral on what counts as 'genuine' team reasoning, and highlight the similarities between argumentative patterns that display collectivistic and individualistic premises to various degrees.

5 The Rationality of Belief-Less Reasoning

Belief-less reasoning is a form of instrumental reasoning that makes optimal use of the information the agents have about the situation. The scheme BR, in the previous section, offers a ‘thin’ version of belief-less reasoning based on two premises about players’ goals and the best means to achieve them. We deliberately stated the conclusion (3) in such a way that it could be read either as a straight prediction, or as a mix of prediction (‘you will choose H’) and intention (‘I will choose H’). We now have to see whether the intention is backed up by a good argument or not.

From the point of view of standard game theory there is a gap in BR, between the identification of the optimal profile of actions and the choices made by the individual agents. For each player, it is optimal to choose H only if she believes that the other player is going to choose H, otherwise the right choice is L. The fact that the *profile* of actions HH is optimal does not constitute a sufficient reason for choosing H, because each agent can only implement part of that profile. We begin by asking whether rationality and common knowledge of rationality can provide the missing reason.

For analytical ease, we introduce two new premises, one identifying rationality with means-end reasoning (3*), and another one making common knowledge of rationality explicit (4*). (BR*)

- 1*. My goal is to maximize my payoff and your goal is to maximize your payoff.
- 2*. The best way for you and me to achieve my goal and your goal is that I choose H and you choose H.
- 3*. A rational player chooses the best means to achieve her goals.
- 4*. There is common knowledge between us that I choose the best means to achieve my goals and that you choose the best means to achieve your goals.
- 5*. Rationally, I must choose H and you must choose H.

Notice that the conclusion in BR* is now stated in prescriptive mode (5*), indicating the actions that a rational player must choose in a game such as Hi-lo. The first two premises are descriptive and, indeed, true. Premise 3* captures the instrumental notion of practical rationality that lies at the core of standard models of rational decision-making. And 4* says that each player is rational in this instrumental sense, she knows that the other one is rational, she knows that she knows, and so on.¹⁴ Which leaves us with the inferential step from premises to conclusion: does 5* really follow from 1*–4*?

According to standard decision theory, an agent is rational in an instrumental sense if she chooses the action that is most likely to lead to the achievement of her goals, *in light of her beliefs* about the choice situation. Or, in other words, standard

¹⁴ Whether common knowledge of rationality holds in real-life situations is an issue that must be settled empirically, of course. But, to repeat, we are primarily interested in normative issues here.

decision theory presupposes a subjective notion of rationality.¹⁵ Common knowledge of rationality seems to imply, a fortiori, that a rational player in interactive situations should consider the beliefs of the other decision-makers, for such beliefs provide (part of) their reasons to do what they must do. But no such beliefs are mentioned in the premises that a belief-less reasoner uses to identify the best means to achieve her goal (1*–2*). So either common knowledge of rationality fails (4* is false) or the conclusion (5*) does not really follow from the premises of the argument.

One cannot appeal to the standard conception of rationality, however, to settle the argument. The issue here is precisely whether such a conception is normatively adequate, so its authority cannot be invoked without begging the question. Moreover, there is an obvious reason why beliefs are not mentioned in the premises of BR*. If they were, as we have seen in the previous sections, players would get stuck with the circularity problem. Therefore, the players should better not focus on others' beliefs. Can they do it without violating common knowledge of rationality? We argue that they can.

Such a move is sanctioned by a commonsensical principle that is not hard to justify. We shall call it the *Principle of Relevant Information*:

(PRI) Whenever you make a decision, you must take into account all and only the information that is relevant for the problem that you are trying to solve. Everything else must be ignored.

The fact that the available information is potentially infinite makes this principle undisputable, which is probably why it is rarely stated explicitly in formal theories of decision-making. But the principle is always operative in the background, so to speak: every decision model presents a selective, abstract representation of the options (and of the properties of the options) available to the decision-makers, based on a subset of the available information. Every model, in other words, includes the information that supposedly matters.

We can use PRI to argue that the beliefs of the other players are irrelevant. This claim can be derived from a simple truth that has been stated earlier: reasoning about the beliefs of the other players does not lead anywhere in games such Hi-lo. Two perfectly rational players cannot solve the problem of coordination by climbing the ladder of meta-representation. Once they realize this fact, they must dismiss beliefs as irrelevant.

Schemes of belief-less reasoning such as BR and BR* allow each player to identify the best *objective* means to achieve their goals. From this identification the players can infer what they should do, without violating the common knowledge of rationality principle. Surely, it is rational to use only (and all) the information that is relevant for the task at hand. The common knowledge of rationality principle demands that the other players are portrayed in the same way—as ignoring the beliefs of others—when they are engaged in a task of this kind.

¹⁵ For example: if Jack believes (mistakenly) that the quickest way to reach the top of the Empire State Building is to take the elevator, but the elevator is broken and he gets stuck for two hours, he cannot be accused of irrationality—he did what he had to do, given the available information.

6 So Much Worse for Rationality?

The standard theory of rational choice is a powerful analytical tool that should not be given up lightly. The fact that its prescriptions occasionally conflict with untutored intuitions should not be surprising or disappointing. It suggests, on the contrary, that its implications are non-trivial and that the theory may help us to question unfounded presuppositions. It may help us to realize, for example, that individual rationality does not always serve us well. There may be cases in which it is better to be irrational, and Hi-lo could be one of them.

This line of reasoning is not farfetched: giving up standard principles of rationality calls for a strong justification. So we ought to explain more precisely what belief-less reasoning does and does not preserve of standard rationality. The core notion of rationality in choice theory is instrumental: as a rational agent I must choose the best means to achieve my goals. Under uncertainty, the principle is usually interpreted subjectively—I ought to choose what I *believe* the best means are, given the information that I have. For simplicity, we shall call this conception of rationality *preference-belief maximization*.

Instrumental rationality per se does not impose any restriction upon the information that one must use in order to identify the best means to achieve one's goal, so something like PRI is needed to screen useful from useless information. But if information about beliefs (and beliefs about beliefs) is irrelevant (if it does not help) then PRI prescribes that we ignore it in the deliberation process. Does this mean that preference-belief maximization must be abandoned? The PRI principle does not imply this: as a principle that screens information, it does not force us to abandon the way in which decision-making is standardly conceived. What it does imply is that we do not infer the beliefs of the other players from the information they have about our beliefs (and beliefs about beliefs)—because, once again, it is impossible to do so.

What information should each player consider then? The structure of payoffs provides relevant information—it tells both players which means are *objectively* the best ones to achieve their goals. We stress 'objectively' because this is a key feature of belief-less reasoning: the two individuals know the way the world is. As a matter of fact, the best way to achieve their respective goals is that they both choose H. Reasoning about the world thus delivers a unique prescription, which they ought to follow.

The process that leads to the identification of the best set of actions (and, a fortiori, of the related expectations) is a process of instrumental rationality. Whether common knowledge of rationality holds or not, then, depends on whether rationality can be identified with instrumental rationality simpliciter. Our defense of BR is based on the consideration that instrumental rationality and PRI are basic and unassailable principles of rational decision making. A rational agent facing a game with imperfect information such as Hi-lo must conclude that she cannot use information about the other player's beliefs to resolve her uncertainty. And she should conclude that, symmetrically, the other agent cannot do it either.

Fig. 2 Two mixed-motive games: **a** The Prisoner's Dilemma and **b** Chicken

	C	D
C	2, 2	0, 3
D	3, 0	1, 1

(a)

	D	S
D	0, 0	3, 1
S	1, 3	2, 2

(b)

We take this to imply that belief-less reasoning does not diverge from the core notion of instrumental rationality, because the latter does not imply that we use the other player's beliefs as a source of information to identify the best means to achieve our goals. Symmetrically, we can legitimately assume that the other player will not use our beliefs as inputs in the reasoning process, by common knowledge of rationality and PRI. Once the best means have been identified, we can—if we want—derive our expectations about the behavior of the other player and their expectations about ours. But neither of these pieces of information is necessary to solve the optimization problem that we face.

7 The Domain of Belief-Less Reasoning

A skeptical reader may wonder at this point whether belief-less reasoning is a general-purpose mode of reasoning. The answer, in a nutshell, is that it is aimed primarily at solving problems of equilibrium selection. So far, we have focused on the Hi-lo game because it constitutes one of the most blatant failures of standard game theory as a normative theory of interactive decision-making. If there is a puzzle that justifies a reconsideration of strategic thinking, then Hi-lo must be it: no other game elicits such a strong, univocal, and theoretically anomalous intuition.¹⁶ The intuition is strong because HH is both a Nash equilibrium and the unique Pareto-efficient outcome of the game.

Still, can belief-less reasoning be applied to other problematic games—such as games that do not have a unique optimal Nash equilibrium? In the prisoner's dilemma game, for example, the optimal outcome is not a Nash equilibrium (Fig. 2a). But since the equilibrium strategies (D, D) are strictly dominant, each player can solve the game without taking the beliefs of the other player into account. The Principle of Relevant Information is idle in this case—it does not prescribe to ignore anything that is not already ignored in the standard analysis of this game.¹⁷

Team reasoning, however, can be used to justify rational cooperation (e.g. Bacharach, 2006). So why do belief-less and team reasoners reach different conclusions in the prisoner's dilemma game? The answer is that team reasoning (in Bacharach's

¹⁶ Not even the prisoner's dilemma: dilemmas of cooperation and mixed-motive games more generally, as we explain below, elicit conflicting intuitions that pull in opposite direction.

¹⁷ It may be argued that some famous attempts to justify cooperation in prisoner's dilemma games are misapplications of belief-less reasoning (e.g. Davis, 1977). To make this point adequately, however, would require a separate paper.

‘strong’ version, at least) achieves this result by modifying the structure of the game. Once the unit of agency has been transformed, strategic reasoning becomes unnecessary because there is no ‘other player’ to begin with. The price to pay is that some prior non-rational psychological process, such as framing, must be presupposed: agency transformation itself cannot be justified rationally, as we have seen earlier.¹⁸

What about mixed-motive games with multiple equilibria? Fig. 2b represents Chicken, a game with two equilibria in pure strategies (DS and SD) and a mixed-strategy equilibrium with probabilities (1/2, 1/2). The main difference with respect to Hi-lo is that the two pure-strategy equilibria have asymmetric payoffs and neither Pareto-dominates the other, so no pair of strategies stands out as ‘obvious’ in the same way as HH does.¹⁹ The Principle of Relevant Information here does prescribe that the beliefs of the other player are ignored, and yet the coordination problem cannot be solved by belief-less reasoning, because the second premise of the BR* scheme is not satisfied: there is no ‘best way for you and me to achieve my goal and your goal’. As a consequence, again, coordination requires a modification in the structure of payoffs.

This should not be surprising: the standard way to solve Chicken problems – both practically and theoretically—is to introduce external mechanisms that facilitate coordination. Crossroads are mundane situations that mirror the payoff-structure of Chicken, for example. Most societies recognize that letting each driver think independently about the best means to achieve her preferred goal is *not* a practical way to avoid car accidents. The best and most common solution involves the introduction of rules and technological devices such as traffic lights. These devices effectively create new outcomes (‘correlated equilibria’, in the jargon of game theory)²⁰ that are Pareto-superior to mixed-strategy equilibria while preserving some of their attractive features – such as payoff symmetry and fairness. But the new equilibria are backed up by a system of incentives that effectively modifies the structure of the game: drivers’ compliance with the rules of traffic is regularly monitored, and transgressors are fined.

There are many other mixed-motive games, and we cannot examine them systematically here. We suspect, however, that the above remarks hold. The domain of belief-less reasoning corresponds to the class of coordination problems with an equilibrium that Pareto-dominates all the other outcomes of the game.²¹ When this condition is not satisfied, we can try to change the nature of the problem, introducing

¹⁸ This point can be generalized to other theories that have been proposed to explain cooperative behaviour: a pair of players may successfully cooperate if they care about equality or fairness for example (e.g. Bolton & Ockenfels 2000; Fehr & Schmidt 1999; Rabin 1993) or if they are sensitive to social norms (Bicchieri 2006). These approaches presuppose, as it is well known, the transformation of the prisoner’s dilemma into a Hi-lo type of game, and this transformation lies outside the realm of practical reasoning.

¹⁹ Although the mixed-strategy equilibrium is symmetric, it is not Pareto-dominant either (it is inferior to SS).

²⁰ See Aumann (1974), Vanderschraaf (1995).

²¹ If there is some other non-equilibrium outcome that is also Pareto optimal, then a belief-less reasoner would be unable to identify the profile of strategies that is ‘best for me and for you’, as stated in premise 2* of the BR* scheme (Sect. 5). This is essentially the same issue that prevents its application to the game of Chicken. We are grateful to a referee for pointing this out.

new incentives, or other-regarding preferences, or changing the unit of agency as in the stronger versions of team reasoning. This should not be a disappointing conclusion: recall that according to Schelling competition and cooperation elicit two different modes of reasoning. Since mixed-motive games involve a bit of both, neither mode of reasoning can be dismissed as inappropriate from the start. Hi-lo in contrast elicits a univocal and strong intuition precisely because the players have a straightforward reason to think cooperatively. When this condition does not hold, the players should not be expected to converge spontaneously on a ‘nice’ solution simply in virtue of being rational. It is sobering to recall that the original game of Chicken did not end well for Jim and Buzz (James Dean and Corey Allen) in ‘Rebel without a Cause’.

8 Concluding Remarks

Problems of coordination have proven to be particularly difficult to solve within the realm of rational choice. One possible approach is to go ‘boundedly rational’ – to argue that people coordinate because they are imperfect reasoners. This approach, which is attractive in a number of settings, seems excessively defeatist in those games that have an obvious ‘logical’ solution, like Hi-lo. In this paper we have outlined a proposal that is able to save the intuition that choosing payoff dominant equilibria in games such as Hi-lo is rational. Unlike some versions of team reasoning, belief-less reasoning is compatible with the idea that the players think as individuals and do not modify the payoffs of the game. Moreover, we have argued, belief-less reasoning is consistent with instrumental rationality.

The main cost of belief-less reasoning is that we have to forsake the idea that higher-order beliefs are a relevant piece of information in all situations of strategic interaction. But we have argued that this is a reasonable cost, for information must always be screened according to the Principle of Relevant Information. If Schelling is right, then, there may be two forms of rational inference, and each one may be appropriate in a specific domain. While ‘Machiavellian reasoning’ works best in competitive settings, belief-less reasoning appears to be most appropriate for certain coordination problems. In those situations that involve a mixture of competition and cooperation, the players are torn between these two ways of reasoning. An important task of institutional design is to nudge them in one direction or another by means of incentives and norms.

Acknowledgements Versions of this paper have been presented at the universities of Firenze, Caserta, Leeds, San Raffaele, and at a conference of the Italian Society of Economists. We are grateful to the members of these audiences for their feedback, and especially to Bob Sugden, Natalie Gold, Luis Bermudez, Ivan Moscati, and various anonymous referees for detailed written comments. This research was funded by the Department of Philosophy “Piero Martinetti” of the University of Milan under the Project “Departments of Excellence 2018-2022” awarded by the Ministry of Education, University and Research (MIUR).

References

- Aumann, R.J. (1990). Nash equilibria are not self-enforcing. In: J.J. Gabszewicz, J.F. Richard, L.A. Wolsey (eds.) *Economic decision making: games econometrics and optimization: Contributions in honour of Jacques H. Drèze*. Business and Economics, pp. 201–206.
- Aumann, R. (1974). Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1, 67–96.
- Bach, C. W., & Perea, A. (2014). Utility proportional beliefs. *International Journal of Game Theory*, 43(4), 881–902.
- Bacharach, M. (1999). Interactive team reasoning: A contribution to the theory of co-operation. *Research in Economics*, 53, 117–147.
- Bacharach, M. (2006). *Beyond individual choice*. Princeton University Press.
- Bernheim, B. D. (1984). Rationalizable strategic behavior. *Econometrica*, 52, 1007–1028.
- Bicchieri, C. (2006). *The grammar of society*. Cambridge University Press.
- Bolton, G. E., & Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review*, 90, 166–193.
- Browne, K. (2018). Why should we team-reason? *Economics & Philosophy*, 34, 185–198.
- Camerer, C. F., Ho, T. H., & Chong, J. K. (2004). A cognitive hierarchy model of games. *Quarterly Journal of Economics*, 119, 861–898.
- Davis, L. H. (1977). Prisoners, paradox, and rationality. *American Philosophical Quarterly*, 14, 319–327.
- Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, 114, 817–868.
- Gauthier, D. (1975). Coordination. *Dialogue*, 14, 195–221.
- Gauthier, D. (2013). Twenty-five on. *Ethics*, 123, 601–624.
- Gilbert, M. (1989). *On social facts*. Routledge.
- Gold, N., & Colman, A. (2018). *Team reasoning and the rational choice of payoff-dominant outcomes in games*. Topoi, online first.
- Gold, N. (2012). Team reasoning, framing and cooperation. In S. Okasha & K. Binmore (Eds.), *Evolution and rationality* (pp. 185–212). Cambridge University Press.
- Gold, N. (2018). Team reasoning and spontaneous collective intentions. *Revue D'économie Politique*, 128, 333–353.
- Gold, N., & Sugden, R. (2007). Collective intentions and team agency. *Journal of Philosophy*, 104, 109–137.
- Guala, F. (2020). Solving the Hi-lo paradox: Equilibria, beliefs, and coordination. In: A. Fiebich (Ed.), *Minimal Cooperation and Shared Agency* (pp. 149–168). Springer.
- Harsanyi, J., & Selten, R. (1988). *A general theory of equilibrium selection in games*. MIT Press.
- Hodgson, D. H. (1967). *Consequences of utilitarianism*. Clarendon Press.
- Hurley, S. (1989). *Natural reasons*. Oxford University Press.
- Hurley, S. (2005). Social heuristics that make us smarter. *Philosophical Psychology*, 18, 585–612.
- Karpus, J., & Radzvilas, M. (2018). Team reasoning and a measure of mutual advantage in games. *Economics & Philosophy*, 34, 1–30.
- Lewis, D. K. (1969). *Convention: A philosophical study*. Harvard University Press.
- Morton, A. (2003). *The importance of being understood*. Routledge.
- Nagel, R. (1995). Unraveling in guessing games: An experimental study. *American Economic Review*, 85, 1313–1326.
- Pacherie, E. (2013). Intentional joint action: Shared intention lite. *Synthese*, 190, 1817–1839.
- Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52, 1029–1050.
- Rabin, M. (1993). Incorporating fairness into game theory and economics. *American Economic Review*, 83, 1281–1302.
- Schelling, T. (1960). *The strategy of conflict*. Harvard University Press.
- Spohn, W. (1982). How to make sense of game theory. In W. Stegmüller, W. Balzer, & W. Spohn (Eds.), *Philosophy of economics* (pp. 239–270). Springer.
- Stahl, D. O., & Wilson, P. W. (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10, 218–254.
- Sugden, R. (1993). Thinking as a team: Toward an explanation of nonselfish behavior. *Social Philosophy and Policy*, 10, 69–89.

- Sugden, R. (2000). Team preferences. *Economics & Philosophy*, 16, 174–204.
- Sugden, R. (2003). The logic of team reasoning. *Philosophical Explorations*, 6, 165–181.
- Sugden, R. (2011). Mutual advantage, conventions and team reasoning. *International Review of Economics*, 58, 9–20.
- Sugden, R. (2015). Team reasoning and intentional cooperation for mutual benefit. *Journal of Social Ontology*, 1, 143–166.
- Sugden, R. (2018). *The community of advantage*. Oxford University Press.
- Sugden, R., & Zamarròn, I. E. (2006). Finding the key: The riddle of focal points. *Journal of Economic Psychology*, 27, 609–621.
- Vanderschraaf, P. (1995). Convention as correlated equilibrium. *Erkenntnis*, 42, 65–87.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.