Krister Bykvist[1] & Karsten Klint Jensen[2]

Having It Both Ways?
On the Prospects for a Cooperation-Friendly Harmonization of Individual and Collective Maximization in Moral Hi-Lo Cases[3]

This paper analyses moral Hi-Lo Cases, which were introduced by Donald Regan's Utilitarianism and Co-operation. Moral Hi-Lo cases are moral coordination problems where coordination equilibriums are ranked by strict betterness. We argue that moral Hi-Lo cases are not just abstract hypothetical cases, there are important real-life cases of this kind, e.g., some climate change cases; and that moral Hi-Lo cases are not just a challenge for utilitarians; they are challenge for all theories that can be represented by a maximizing teleological structure. Moral Hi-Lo cases pose the challenge for individually maximizing theories that they are not collectively maximizing. We show that the widespread solution to moral Hi-Lo cases of adding the option of taking a cooperative stance to the choice situation risks changing the topic. Moreover, in the changed situation, simply making available a cooperative attitude or act is not sufficient to harmonize individual and collective maximization. This suggests that the problem sticks deeper than exclusively act-orientedness, as Regan suggested. It is not sufficient for this harmonization to assume that it is possible to influence the other agent and make her cooperative, it is necessary to actually influence her, but even with this extra assumption about actual influence, taking a cooperative stance for the best outcome may not be mandatory, if the strategy as a whole involves costs, which is a realistic assumption.

[1] Institute for Futures Studies & Stockholm University, krister.bykvist@philosophy.su.se

[2] Institute for Futures Studies, karsten.klint-jensen@iffs.se

[3] Funding from Riksbankens Jubileumsfond (grant numberP22-0662) is gratefully acknowledged.

1. Introduction

Donald Regan (1980) presented what we shall call *moral Hi-Lo problems* in his groundbreaking book *Utilitarianism and Co-Operation*. His proto-type case, which he used throughout the book, looks like this: There are only two agents in the moral universe, Whiff and Poof. Each has a button which he can push or not. The possible outcomes are evaluated by numbers representing units of value for the overall state of the world. Neither agent can influence the other's choice.

Table 1. Regan's Whiff-and-Poof-case

		Poof	
		Push	Not-push
Whiff	Push	10	0
	Not-push	0	6

Reagan was inspired by a similar case, set up by Allan F. Gibbard (1965). As Gibbard saw things, such coordination problems pose a challenge to act utilitarianism, since it does not necessarily ensure the collectively best outcome. If Poof not-pushes, act utilitarianism requires Whiff to not-push as well. And the same holds for Whiff if Poof not-pushes. In other words, the act pattern (not-push, not-push) is individually maximizing, i.e., the best each agent could do on their own. But they could together bring about an outcome of value 10 by each pushing. Thus, the act pattern (push, push) is collectively maximizing, i.e., the best they could do together. The example thus shows that an individually maximizing act-pattern need not be collectively maximizing. Gibbard concludes that some form of institutional coordination is needed to achieve the collectively best outcome. However, Regan's aim is to demonstrate that coordination *can* be achieved by morally motivated agents.

The climax of Regan's analysis is the proof that no theory which fulfils a necessary condition for being exclusively act-oriented can be strongly collectively maximizing.[4] The necessary condition for a theory to be exclusively act-oriented in a Hi-Lo case is that the theory specifies, for each agent, some subset of the set of available acts, such that the agent satisfies the theory iff she does an act from the specified subset.

This led Regan himself to suggest that act utilitarianism should be supplemented by a somewhat complicated decision procedure, which he argues is able to ensure coordination for the collectively best outcome. No one else has followed him in that. But many have accepted the premise that a solution must involve going beyond an exclusively act-oriented theory, e.g. by adding a cooperative attitude or an act inviting

[4] We define the properties of individually and (strongly) collectively maximizing theories below.

to cooperation to the case and supplementing an act-consequentialist type of theory with a duty to take on the attitude and/or perform the invitation.

One of the aims of this paper is to make a rational reconstruction of this type of theory in order to explore the prospects of finding a cooperation-friendly harmonization of individual and collective maximization in moral Hi-Lo cases. More specifically, we shall show the following:

1. moral Hi-Lo cases are not just abstract hypothetical cases, there are important real-life cases of this kind, e.g., some climate change cases (sections 2 and 3);
2. moral Hi-Lo cases are not just a challenge for utilitarians; they are a challenge for all theories that can be represented by a teleological structure (sections 4 and 5);
3. adding the option of taking a cooperative stance to the Hi-Lo case in Table 1 risks changing the topic. Hence a solution to the changed Hi-Lo case may not be a solution to the original case (section 5);
4. in the changed case, simply making available a cooperative attitude or act is not sufficient to harmonize individual and collective maximization. This clearly suggests that the problem sticks deeper than exclusively act-orientedness, as Regan suggested (section 6);
5. in the changed case, it is not sufficient for this harmonization to assume that it is *possible* to influence the other agent and make her cooperative, it is necessary to *actually* influence her (section 6);
6. but even with this extra assumption about actual influence, taking a cooperative stance for the best outcome may not be mandatory, if the strategy as a whole involves costs, which is a realistic assumption (section 6).

Before we start arguing for these claims, we shall first give a more precise definition of a moral Hi-Lo case.

2. What is a moral Hi-Lo problem?

Moral Hi-Lo problems constitute a subclass of what can be called moral coordination problems. Regan does not provide a general definition of either of these; he mainly works from the generic case cited above. Let us first adapt from game theory the concept of a coordination equilibrium[5] to this context:

[5] Lewis (1969: 14). The concept is clearly modeled on the concept of a Nash equilibrium, which is defined in the

> A *coordination equilibrium* is a combination of acts in which the overall state of the world would not be better if any one agent alone acted otherwise.

We shall define a moral coordination problem thus:

> In a *moral coordination problem*, each of *n* agents chooses one act from a finite set of alternatives. Each outcome has an objective moral value. There are at least two coordination equilibria. Coordination equilibria are ranked by an 'at least as good as'-relation, and no non-equilibrium combinations is better than any equilibrium.

We get a *moral Hi Lo problem*[6] if none of the equilibriums are equally good; i.e. all equilibriums are ranked by strict betterness. Regan only considers simple two person cases with two available acts and we shall follow him in that. This simplifies the discussion considerably and can for the most part be done without any loss of generality.

It might be relevant to add to the ranking of outcomes an assessment of the value difference between the best equilibrium outcome and the second-best, and between the second-best equilibrium outcome and best non-cooperative outcome. Let us call it a *high-stake case*, when either or both of these differences are significant.

Regan's Whiff-Poof case is presented *as if* it were a game, more precisely a coordination problem. And clearly, the case shares with games the property that the outcome of an agent's choice is depending on the choices made by others. But the case cannot simply be identified with a standard game. The numbers represent an agent-neutral ranking of outcomes as overall states of the world. They do not necessarily represent the preferences of the agents, as they would do, if it was a game.

As Regan presents the case, there is uncertainty about the preferences of the agents, whereas in a standard game, full information about preferences is assumed. The agents in the situation may have both self-interested preferences and conflicting preferences, resulting in rankings of the outcomes which deviate from the objective moral ranking assumed in the case.

It is part of Regan's proposed decision procedure to provide shared information about the agents' preferences. Thereby the decision problem, initially under uncertainty, can be transformed into a moral coordination problem, which however only may be faced by the subgroup who shares an agent-neutral objective moral ranking and disregards those who are unwilling to cooperate.

Moral Hi-Lo cases show up in a more *indirect* way as well, for some standard games can be transformed to a moral Hi-Lo case. Consider a standard Hi-Lo game (the first

context of non-cooperative games, e.g. Luce & Raiffa (1957: 106).

[6] The name is of course inspired by Bacharach (2006).

number represents the preference intensities of agent A, the second those of B) (Bacharach 2006):[7]

Table 2. A standard Hi-Lo case

		B	
		Hi	Lo
A	Hi	5 / 5	0 / 0
	Lo	0 / 0	3 / 3

Bacharach calls this a *common interest* game. If (Hi, Hi) is better than (Lo, Lo), which is better than (Hi, Lo), we have a moral Hi-Lo problem. This would be so, if the preference intensities were simply summed.

It is important to note, however, that Hi-Lo cases are relevant to *non-utilitarian* moralities as well. The numbers need not represent the sum total of preference intensities or wellbeing; they can represent some other aggregation of all relevant values (not necessarily welfarist). Nor do we have to assume consequentialism for Hi-Lo cases to be of interest. As Portmore (2018) points out, the numbers can be seen as representing the moral value of the act-combination and its associated outcome. Even deontologists and virtue-ethicists can accept that things have impartial moral value. But, pace Zimmerman (1996) and Pinkert (2015), impartiality or agent-neutrality is not required either. Here is a moral Hi-Lo case where the numbers represent agent-relative moral value, degrees of moral reasons or choice-worthiness, which can depend on the agent's motivations or dispositions. The first number represents the agent-relative value for A of A's action (and its outcome, if that matters), the second number the agent-relative value for B of B's action (and its outcome, if that matters).

Table 3. A moral Hi-Lo case with agent-relative values

		B	
		Hi	Lo
A	Hi	-2 / 5	-8 / 0
	Lo	-8 / 0	-4 / 2

In this case, the available actions do not even have the same polarity for the agents. All actions are bad relative to A, and neutral or good relative to B. Still, if (Hi, Hi) is

[7] Throughout, when we for simplicity call acts 'Hi' and 'Lo', they need not represent the same act for the two agents. 'Hi' for each represents the act which combines to the best equilibrium, and 'Lo' for each agent represents which combines to the second-best equilibrium.

better than (Lo, Lo), which is better than both (Hi, Lo) and (Lo, Hi), we have a moral Hi-Lo problem.

Bacharach (2006) shows that some games of conflict can be transformed into common interest Hi-Lo games, which by the argument above then can be transformed into a moral Hi-Lo case. Stag Hunt is one example:

Table 4. Stag Hunt

		B	
		Stag	Rabbit
A	Stag	2 / 2	-1 / 1
	Rabbit	1 / -1	1 / 1

This becomes a moral Hi-Lo problem, if (Stag, Stag) is better than (Rabbit, Rabbit), which is better than both (Rabbit, Stag) and (Stag, Rabbit). This would be the case if the value is the sum of individual preference intensities

Other examples are some versions of the Prisoner's Dilemma, like the following one.

Table 5. Prisoner's Dilemma

		B	
		Cooperate	Defect
A	Cooperate	4 / 4	0 / 5
	Defect	5 / 0	3 / 3

This will become a moral Hi-Lo problem, if (Cooperate, Cooperate) is better than (Defect, Defect), which is better than both (Cooperate, Defect) and (Defect, Cooperate). This would be the case if value is the sum of preference intensities.

3. Why are moral Hi-Lo problems important?

That many important decisions involve coordination problems is widely accepted. Choices of great importance often require coordination to achieve the optimal results. Some of these problems can be seen as moral high stake Hi-Lo cases, either involving several agents or only two. There are numerous examples of such Hi-Lo cases, both from the individual and the political spheres. For example, many rescue cases have this form. They can be found whenever there are two rescue options, both of which requires cooperation to succeed, and where one is better than the other (cf. Colman et al. 2014: 36). Suppose that in the aftermath of an earthquake we find out that there is one person buried under a crumbled building and several people buried under

another. There is no time to rescue all people, because the buildings are far apart and the oxygen levels under the buildings are falling quickly. In order to successfully save any of the people, we need to both be there to move the heavy rubble. Our radios are not working so we cannot communicate. This case has the form of a moral high stake Hi-Lo case. We can either together save the larger group (Hi, Hi) or the lone person (Lo, Lo), but no one will be saved if we go to different buildings, (Hi, Lo) or (Lo, Hi).

Other examples come from the climate change context. The most straightforward and common cases are choices between climate actions that require extensive infrastructure to work and where each party's contribution to the infrastructure is crucial. One pair of coordinated climate options (Hi, Hi) might be better than another (Lo, Lo), but (Hi, Lo) and (Lo, Hi) would be worst, because unilaterally going for one option would provide an insufficient infrastructure. To take a mitigation case, 'Hi' can be electrification of aviation and 'Lo' expansion of fast trains. Both parties going for electrification of aviation is better than both going for expansion of fast trains, because of the time benefits of flight travel. But unilateral choice would not provide sufficient infrastructure for either aviation or expansion of fast trains. Instead, the unilateral choice would only incur futile costs. Similar cases can be constructed in which the choice is between different energy systems, for example, hydro and wind versus nuclear.

Another case is an adaption case, where the choice is between two adaption strategies against flooding: building seawalls and relocating the population. Both parties going for building seawalls is better than both going for relocation, because no one is forced to move if seawalls are built. However, the mixed options are worse because then not enough infrastructure will be put in for the seawalls to be effective and people will still have to relocate. To make things more concrete, assume that two neighbouring nations are each threatened by sea-level rise, where they share a salient geographic border. Without both nations building seawalls there would be flooding. If only one nation builds a seawall, the flood waters will just be pushed toward the part of their shared geographic boundary that is uncovered, and the same damage will occur, with the water just taking a different (slower) route. The country that built the wall will eat the entire cost of building the wall, which will then be useless.

As pointed out above, moral Hi-Lo cases may show up in a more indirect way as well, when a decision we face have the form of a Prisoner's Dilemma that can be transformed to a moral Hi-Lo case. One famous illustration is the so-called Polluter's Dilemma, which has the structure of a Prisoner's Dilemma.[8] Suppose A and B (who can be nations or individuals) have two options: not pollute (cooperate) and pollute

[8] For a thorough discussion of when a pollution case is best seen as a case of a Prisoner's Dilemma, see Pellikaan and van der Veen (2002).

(defect). The Pareto-optimal outcome is one in which both do not pollute. But it is not an equilibrium, since each agent would be better off unilaterally defecting and polluting. This assumes that the pollution produced by each agent is not significant enough to outweigh the benefits of not having to pay for making production pollution-free.[9] To get a moral Hi-Lo case it is enough to assume that both cooperating and not polluting is better than both polluting, which is better than one polluting and the other not polluting. In some cases, e.g. climate cases, the agents in the decision problem may not stand to benefit themselves but all benefits go to a third party (future generations). The problem would then not have the structure of a traditional Prisoner's Dilemma, but it could still be transformed into a moral Hi-Lo problem, if the agent-neutral ranking is this: (Hi, Hi) is better than (Lo, Lo), which is better than both (Hi, Lo), (Lo, Hi).

4. Assumptions and definitions

In the following we shall clarify the framework for our discussion, which we largely adapt from Regan (1980) who has framed the subsequent discussion. We are dealing with *objective moral theories*, according to which rightness depends on the facts, not the agent's beliefs or evidence about the facts. In contrast to game theory and decision theory, information and subjective probabilities are not normatively relevant.[10]

It is true that many authors find it necessary to add more information to the agents about the cases. As pointed out above, Regan's own strategy is to transform the moral Hi-Lo case into a coordination game where the agents have information about its structure. However, since our analysis is focusing on objective duties, we do not need such assumptions.

Second, we start from an individual duty perspective. The question is what each individual ought to do. As we shall see, some authors want to introduce collective duties, and our starting point does not exclude this possibility, since one could argue that individual duties can be derived from collective ones.

Third, in contrast to Regan who is concerned with act utilitarianism only, we are dealing more broadly with moral theories that are teleological in the following weak sense: theories that can be given a maximizing teleological representation, where rightness of an action is determined by the highest-ranked outcome (at least if other

[9] It is controversial whether this applies to national agents in climate change, since one could argue that here the pollutions are significant, at least for big emitters such as the US and China. For overviews of the relevance of Prisoner's Dilemma to climate change, see Chander (2018) and Magli and Manfredi (2022).

[10] Regan allows for objective probabilities as well, but as pointed out by Rabinowicz (1989), this creates more problems than benefits. Since probabilities play no role in his argument anyway, we ignore this possibility.

things are equal). This outcome can be an outcome of a group-action of which the individual action is a part.

Finally, we assume that each available alternative for an individual agent has a determinate morally relevant outcome given the pattern of behavior of the other agents.

We shall also list some properties of moral theories (denoted T) relevant for our arguments.

- T is *universally satisfied* in some pattern of actions iff all agents do what T requires of them in this pattern.
- T is *individually maximizing* (IM) iff for any agent, in any choice situation, if the agent satisfies T in that situation, he produces by his act the 'best' consequences he can possibly produce in that situation. This is a generalization of Regan's PropAU.
- T is *collectively maximizing* (CM) iff for any pattern of actions in which T is universally satisfied, the class of all agents produce by their acts taken together the 'best' consequences that they can possibly produce by any pattern of behavior. This is Regan's PropCOP, which is one version of what is often called moral harmony.
- T is *strongly collectively maximizing* (SCM) iff for any situation involving choices by any number of agents, the agents who satisfy T in that situation produce by their acts taken together the 'best' consequences that they can possibly produce by any pattern of behavior, given the behavior of agents who do not satisfy T. This is Regan's property of T being Adaptable. As pointed out by Regan (1980: 107), 'T is SCC' entails 'T is CM'.

The reason we talk about best in square quotes is that it is meant to also capture the case where the relevant consequences are *at least as good* as that of any other relevant alternative.

Note that the property of being IM is defined in an unqualified way, i.e., for *any* choice situation. We shall later discuss the prospect of theories being IM only in certain specific choice situations.

5. Why are moral Hi-Lo cases a problem for act utilitarianism? Generalizing the challenge and two strategies to address it

Regan presented his Whiff/Poof case as a problem for act utilitarianism. We have already hinted at what the problem is: In the Whiff/Poof case, given that the agents act independently of each other, there are two patterns of actions, where act utilitarianism is universally satisfied, (Push, Push) and (Not-push, Not-push). In (Not-push, Not-push), the agents together do not produce the best possible consequences they possibly can. Hence, act utilitarianism is not CM.

However, there is a simple generalization of the challenge at hand. Consider any theory which is individually maximizing applied to a two-person moral Hi-Lo problem. Any such theory is universally satisfied in the pattern (Lo, Lo), and therefore it is not collectively maximizing. Hence, no theory which is individually maximizing can be collectively maximizing. This result should not be surprising, since (Lo, Lo) is a coordination equilibrium and the concept of a coordination equilibrium is defined from an individually maximizing perspective.

There seem to be two possible strategies to address this challenge. One is to reach for a theory which is only CM and thus giving up the requirement that it should also be IM. Call this the *non-reframing approach*. For this approach, the standard two-person choice situation still has only two options, Hi or Lo. We could imagine both individualist and collectivist versions of this strategy. As an individualist example, consider a simple rule consequentialism, according to which you should always follow the best rule, i.e. the rule that would have the best consequences if everyone followed it. This theory would tell each agent to do Hi, since the best rule would require both to do Hi. But consider also a simple collectivist version, which would tell each person to do their part in the collective duty to do the best we can do together. Since the best we can do together is that both do Hi, each agent ought to do Hi.

The non-reframing approach is really a non-starter, however, since it would implausibly require each agent to do Hi, even if a *catastrophe* would ensue. In fact, we are uncertain if anyone seriously would want to defend it. There is one qualification, though, because the argument is based on high-stake situations. Imagine a low-stake situation, where the value of the third-best outcome is not very much worse than the value of the best. Here, there could be a hard choice whether or not we should accept that a collectively maximizing theory might involve some relatively minor costs or harms for individuals. It is defensible to accept such costs in low-stake situations. Nonetheless, a collectively maximizing theory can only be plausible if it rejects such

costs in high-stake situations, which is to accept that collective maximization does not apply universally.

Hence, the conclusion is that no theory should be CM. But given low-stake cases, perhaps no theory should be IM either. The best theory might combine individual and collective elements. It is clear, however, that a theory should agree with individual maximization in high-stake cases.

The other strategy we call the *reframing approach*. This approach accepts that IM theories cannot be CM in the standard moral Hi-Lo case. The aim is to reframe or transform the standard case to a situation where an IM theory *can* be CM. This reframing is done by altering the existing options of the original case or adding one or more options to them and thereby creating new compound alternatives. As we will see, these new compound alternatives need not be compound *actions* (strictly speaking); there may instead be combinations of attitudes and actions, or reasoning processes and actions. Though this approach confirms that the original challenge cannot be met, it might of course still be of interest to see if a theory can be both IM and CM in such reframed and perhaps more realistic Hi-Lo cases.

There are many different instances of the reframing approach, including the ones defended by Regan (1980), Zimmerman (1996), Portmore (2018), Pinkert (2015), Schwenkenbecher (2021), and Goodin (2012). Let us start with quoting a short summary of Regan's (1980: 135f.) approach, which appears to have been an inspiration for most re-framers (italics added):

"Each agent should

- *be willing* to take part in a joint attempt to produce the best consequences possible by co-ordinating his behaviour with the behaviour of other agents who are willing

- *consider* the other agents involved in the co-ordination problem he is making a decision about and determine which of those other agents are available to be co-operated with.

- *ascertain* how other agents who are not (for whatever reason) available to be co-operated with are behaving or disposed to behave

- *identify* the best possible pattern of behaviour for the group of co-operators […] given the behaviour (or dispositions to behave) of the non-co-operators

- *do* his part in the best pattern of behaviour just identified."

Regan presents his theory as a non-exclusively act-oriented IM theory which is also generally SCM.[11] It is designed as a decision procedure which determines which of the acts available in the original situation (Hi or Lo) is the right one to choose under the circumstances. But it is clear that he is re-framing the situation, since he adds that each agent should not just perform a certain action but also have a certain willingness, and moreover consider, ascertain, and identify aspects of the aspects of the choice situation.

The original options, i.e. simply choosing without going through (or fail or reject) the procedure, are not mentioned by Regan and presumably they are not available. Instead, the theory tells you to go through a procedure, consisting of a number of preparatory steps and concluding in the actual choice of an action.

Zimmermann is also re-framing. He presents a revised version of his preferred individually maximizing theory, according to which "(d)oing the best one can must be accompanied by the adoption of a certain *attitude*" (italics added).[12] This is a clear alteration of the options in the original Hi-Lo case.

Here are other examples of what re-framers propose:

- Be willing to cooperate with others who are willing (Sugden 2015)
- Form a disposition to cooperate (Portmore 2018; Pinkert 2015)
- Make a cooperative commitment (Goodin 2012)
- Take a cooperative attitude (Zimmerman 1996)
- Do we-reasoning (Bacharach 2006; Schwenkenbecher 2021)
- Say or signal that one is cooperative (Schwenkenbecher 2021)
- Identify cooperators (Regan 1980)

We suggest that these proposals can be summarized as adding a new option of 'taking a cooperative stance' (Cop). (Note that this extra option need not be an action, strictly speaking.) So, instead of just having the alternatives doing Hi or doing Lo, we now have the alternatives:

[11] Even though the procedure involves acquiring certain beliefs, Regan boldly claims that it is a completely objective theory. But this appears doubtful, since acting rightly according to the theory depends on having the correct beliefs.

[12] Zimmerman (1996: 263).

- doing Hi/Lo while or after[13] taking a cooperative stance (Cop&Hi, Cop&Lo)
- doing Hi/Lo while or after refusing to take a cooperative stance (¬Cop&Hi, ¬Cop&Lo)

We assume that taking a cooperative stance entails that one is *successfully cooperative* (in the pursuit of value). Of course, this is not always true. One can take a cooperative stance and fail to be cooperative, because of weakness of will, dishonesty, or lack of crucial information. But we want to make the best case for the re-framers. We will assume that in the Hi-Lo two-agent case, being successfully cooperative entails being such that

- one would do Hi, if the other agent were to do Hi.

Note that to make taking a cooperative stance available in the choice situation is to reject the assumption which is a defining characteristic of Regan's original problem, namely that the agents are *disconnected* in the sense that they choose independently of each other. For if one agent takes a cooperative stance, this is one way she can connect to the other agent in the sense that if the other agent did Hi, she would respond with Hi.[14]

Refusing to connect is being such that one chooses independently of the other agent. We shall in particular be concerned with the case where both agents do Lo, in which case not taking a cooperative stance entails that:

- one would do Lo, even if the other agent were to do Hi.

In the general case of any group of agents G, a member of G taking a cooperative stance towards G entails

- one would do Hi, if all other members of G were to do Hi.

For our argument we only need these plausible ways of (not) taking a cooperative stance. Furthermore, we shall conduct our argument concentrating on the standard two-person case. In this case, the reframed situation looks like this:

[13] For some re-framers, it is important that taking a cooperative stance (or not) may be at a time earlier than the choice between 'Hi' and 'Lo' so the other may be able to respond to the stance. We discuss cases where this issue becomes important below.

[14] It is also possible to attempt to connect by taking an *uncooperative* stance, i.e. fulfilling the condition 'One would do Lo if the other agent were to do Hi, and one would do Hi, if the other agent were to do Lo'. Bykvist (forthcoming) calls this 'the contrarian option'. It was first envisaged by Feldman (1986). It plays no role in our argument, and we shall therefore not include it in the model.

Table 6. The re-framed choice situation

			B			
			Cop		¬Cop	
			Hi	Lo	Hi	Lo
A	Cop	Hi	10	n.a.	10	0
		Lo	n.a.	6	0	6
	¬Cop	Hi	10	0	10	0
		Lo	0	6	0	6

N.a. (not applicable) signifies an impossible combination. A cannot fulfill the Cop-condition 'if B were to do Hi, then A would do Hi', if A does Cop&Lo when B does Cop&Hi, and vice versa.

This matrix makes it clear that re-framers are changing the topic, simply because the availability of new alternatives reflects that the original assumption of independent choices has been rejected. In other words, the reframed situation is a changed situation. Could it nevertheless be argued that these alternatives were available in the original situation, in other words, that the characterization of the original case ignored certain available options?

It is clear that the original case has been changed, if taking a cooperative stance is identified with a *physical* action, like saying or signaling that one is cooperative, identifying cooperators, or making others cooperative. These physical actions need not be available to an agent, not even implicitly. But perhaps the option of taking a cooperative stance can be seen as implicit in the original case, if they are identified with *mental* acts, such as doing we-reasoning. This is not clear, however, since on the most famous account of we-reasoning (Bacharach's 2006), such reasoning is not a voluntary option.

There is no need for us to decide exactly when the original case can be said to be changed, since our main aim is to examine under which conditions satisfaction of a theory being IM is compatible with it being CM.

6. The prospects for a cooperation-friendly harmonization of individual and collective maximization

6.1 The Nice Case

What Regan thought he could achieve with his decision procedure is a coordination game where both take a cooperative stance. This is true. See Table 7.

Table 7. The nice case

		B	
		Cop&Hi	Cop&Lo
A	Cop&Hi	10	n.a.
	Cop&Lo	n.a.	6

In this situation, if B does Cop&Hi, it is IM for A also to do Cop&Hi. In fact, there is no other choice available! The same holds for B. Any IM theory is universally satisfied in the pattern (Cop&Hi, Cop&Hi) and thus trivially satisfies being CM. But note this only holds, if taking a cooperative stance involves no costs, a topic we will come back to below.

6.2 The challenge from disconnection

Suppose that there is a mutual disconnection between the agents in the sense that choosing Cop would not make the other agent choose Cop. Suppose further that both A and B choose ¬Cop&Lo (marked in bold in table 8). Consider Table 8. In the yellow area, you find the nice case. The grey column shows the available outcomes from A's point of view, when B does ¬Cop&Lo, and A cannot do anything to make B choose Cop. As can be seen, even if A were to do Cop, the best she can do in that case is also to choose Lo (the value would be 6, which is marked by green), since her cooperative stance would have no effect on B. In other words, A can do nothing to move B into the nice case. Choosing ¬Cop&Lo would have the same value (again with value 6, marked green). The same holds for B, since the situation is exactly symmetrical. Hence, we get the result that any theory that is IM in this situation is universally satisfied in the pattern (¬Cop&Lo, ¬Cop&Lo). But this pattern is not CM. We are in effect back to the predicament of the original case where both agents do Lo and no agent can influence the other agent's actions. Furthermore, note that there is a tie for each agent between choosing (¬Cop&Lo) and (Cop&Lo) when it comes to IM, since each option would lead to an outcome with value 6.

Table 8. The case of disconnection

			B			
			Cop		**¬Cop**	
			Hi	Lo	Hi	Lo
A	Cop	Hi	10	n.a.	10	0
		Lo	n.a.	6	0	6
	¬Cop	Hi	10	0	10	0
		Lo	0	6	0	6

Would things change if *one* of the agents took a cooperative stance? Many re-framers think that in that case the non-cooperative agent would no longer be an individual maximizer, since the cooperative agent is such that she would respond with a Hi to a Hi. The cooperative agent should then choose what is IM under the circumstances, which in the case of Table 8 would be Cop&Lo.[15]

But it need not be true that the non-cooperative agent is not IM if we consider a *diachronic* case in which the cooperative agent acts first. In this case, it would be IM for the cooperative agent to choose Cop&Lo, since the other agent is uncooperative and would choose Lo, no matter what she did. Now, once Lo has been chosen by the cooperative agent, the uncooperative agent faces a choice between Hi (with or without Cop), which would lead to an outcome of value 0, and Lo (with or without Cop), which would lead to an outcome of value 6. So, choosing ¬Cop&Lo would be an IM option for the uncooperative agent. So, if the cooperative agent chooses Cop&Lo and the uncooperative chooses ¬Cop&Lo, they are each choosing what is IM. But this act-pattern is not CM. The situation can be illustrated with the following tree diagram:

Figure 1. A diachronic case

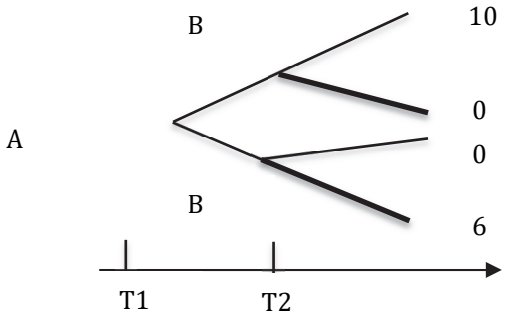

Suppose A is the cooperative agent, and B is the uncooperative agent. A bold line indicates what is actually chosen or would be chosen. Going up for A means choosing Cop&Hi; going down for A means choosing Cop&Lo. Going up for B means choosing Hi (with or without Cop); going down for B means choosing Lo (with or without Cop). (We have omitted the branches that involve A choosing ¬Cop, since they do not make a difference for the argument.) It is clear that if A chooses Cop&Lo at T1, this is an individually maximizing choice. Once this has been done, B would

[15] In a larger group, an individually maximizing theory would still be SCM when the subgroup of cooperative agents together chooses the best outcome, given the behavior of the non-cooperative agents. Regan (1980) considers this feature a major advantage of his theory.

be an individual maximizer if she were to choose ¬Cop&Lo at T2. But this act combination is not CM.

Note that how this situation differs from the *synchronic* one, where both agents act at the same time. In that case, since A is cooperative and would do Hi if B did Hi, B does not perform an IM action by doing Lo. If she had done Hi instead, A would have done Hi and the outcome of value 10 would have been achieved. This also holds for the diachronic case where the uncooperative B acts first. If B chooses Lo at T1, she does not act in an IM manner, since if she had done Hi instead at this time, A would have followed it up later at T2 with a Hi and the outcome with value 10 would have been achieved.

6.3 Adding the value of being cooperative

Many re-framers work from the intuition that refusing to cooperate (choosing ¬Cop) is wrong. They are willing to accept that if they face someone who would not cooperate and chooses Lo, the best you can do is to choose Lo. But they find it hard to "allow two wrongs to make a right" (Zimmermann 1996: 257), which seems to be the case for a theory which is universally satisfied by the pattern (¬Cop&Lo, ¬Cop&Lo). Hence, in that case, the theory should require Cop. It does not make the outcome better,[16] but it allows the theory not to violate the property of being SCM.

But note that a theory which merely stipulates that refusing to cooperate (choosing ¬Cop) is wrong jeopardizes the property of being IM. One might think that one could address this problem by assigning some extra final value v to an outcome if it results from Cop rather than ¬Cop. The matrix would then look like this.

Table 9. Added value of being cooperative

			B			
			Cop		¬Cop	
			Hi	Lo	Hi	Lo
A	Cop	Hi	10+2v	n.a.	10+v	0
		Lo	n.a.	6+2v	v	6+v
	¬Cop	Hi	10+v	v	10	0
		Lo	v	6+v	0	6

Note that by adding this value v we are breaking the tie for A: ¬Cop&Lo is no longer an IM choice; Cop&Lo would have a better outcome with value 6+v. Suppose that B is choosing ¬Cop&Lo and would do so, no matter what A did. Then choosing Cop&Lo is the only IM choice for A, as can be read from the grey column. But this

[16] As pointed out by Feldman (1986).

maneuver does not help in general. If A acts first, then B will later have a choice between doing Cop&Lo with value $6+2v$, doing ¬Cop&Hi with value v, and doing ¬Cop&Lo with value $6+v$. So, choosing Cop&Lo would be an IM choice for B. But the combination (Cop&Lo, Cop&Lo) is not CM. So, even if we have managed to break the tie for A, we still have not succeeded in establishing harmony between IM and CM in this case.

6.4 A stronger form of connection

We have seen that we cannot combine CM and IM in all cases where there is *one* agent who is not taking a cooperative stance. Let us now look at cases where one can do Cop and at the same time make sure that the other agents take a cooperative stance, perhaps by 'getting assurance' that the other person will cooperative, as in Sugden (2017), 'promote' cooperation as in Cripps (2013), or create a collective agent as in Collins (2019). Consider the situation in Table 10 from A's point of view, where 'Cop*' is defined as

- one would do Cop *and* at the same time make sure the other agent does Cop.

Table 10. Making the other agent taking a cooperative stance

			B					
			Cop*		Cop		¬Cop	
			Hi	Lo	Hi	Lo	Hi	Lo
A	Cop*	Hi	10	n.a.	10	n.a.	n.a.	n.a.
		Lo	n.a.	6	n.a.	6	n.a.	n.a.
	Cop	Hi	10	n.a.	10	n.a.	10	0
		Lo	n.a.	6	n.a.	6	0	6
	¬Cop	Hi	n.a.	n.a.	10	0	10	0
		Lo	n.a.	n.a.	0	6	0	6

Suppose B does not take a cooperative stance, but chooses ¬Cop&Lo. If A were to take a strong cooperative stance (Cop*) towards B, she would make B take a cooperative stance towards her and thereby move her into upper right orange area doing Cop&Hi. The outcome would then be 10 (blue). Any other act by A would have a worse outcome (red). (It is possible that B would even want to do Cop*&Hi, thereby moving into the green area).

Suppose the same holds from B's point of view. Then any IM theory is universally satisfied in the patterns (Cop*&Hi, Cop&Hi), (Cop&Hi, Cop*&Hi) and (Cop*&Hi, Cop*&Hi), and thus also CM. Hence, we have found other cases than the 'nice case' (yellow area) where there is harmony between IM and CM. But it is important to note

that merely adding the option of taking a cooperative stance (Cop) does not help. Nor does it help to assume that one agent is taking the cooperative stance, as we pointed out in section 6.2. We also need to assume that the agents' Cop-actions are *mutually* connected.

This example can also be used to fulfil another aim of this paper: to show that cooperation is not necessarily involved in the act-patterns that are both IM and CM. For suppose both A and B choose ¬Cop&Hi. Any individually and collectively maximizing theory is universally satisfied in this pattern too (green). In order to make (Hi, Hi) combinations involving one or two Cop*-stances better solutions, these outcomes need to be assigned an extra final value v, perhaps because being cooperative itself has final value, as we discussed in section 6.3.

6.5 Taking a costly cooperative stance

So far, we have assumed that taking a cooperative stance is cost-free. But this is a very unrealistic assumption. There are two ways, in which taking a cooperative stance might involve costs to the outcome. One is that it may have unintended negative moral consequences, which must be subtracted from the value of the outcome.[17] This is something almost any kind of act risks, with some probability. We shall ignore this possibility in our argument, since it is not relevant in our framework which assumes that each available alternative for an individual agent has a determinate morally relevant outcome given the pattern of behavior of the other agents.

We find it more important that taking a cooperative stance, in most cases at least, appears *certain* to involve a personal cost in terms of spent time and energy for the agent who undertakes it.[18] If only this cost is positive, however small, it will have a dramatic effect on the evaluation of the outcomes in the cases where successful coordination is achieved. Consider Table 11:

[17] Regan himself (1980: 267ff.) introduced the possibility of a mad telepath, but he did not consider it a serious problem. However, most commentators have used this example against him.

[18] Regan (1980: 267ff.) also considers costs of this kind, but – as Zimmerman (1996: 260) says – he *deliberately* ignores them.

Table 11. The cooperative stance involves a cost

			B			
			Cop(*)		¬Cop	
			Hi	Lo	Hi	Lo
A	Cop(*)	Hi	10-2*c*	n.a.	10-*c*	-*c*
		Lo	n.a.	6-2*c*	-*c*	6-*c*
	¬Cop	Hi	10-*c*	-*c*	10	0
		Lo	-*c*	6-*c*	0	6

Remember there are two ways to achieve coordination: either if both agents take a cooperative stance (what we called the 'nice case' above) or if one agent takes the strong form of a cooperative stance (cop*). Now we assume that taking a cooperative stance (choosing Cop(*), i.e., either Cop or Cop*) involves a (morally relevant) personal cost of *c*. We contrast choosing Cop(*) with choosing ¬Cop, which serves as the baseline in terms of costs.

It is clear from Table 11 that no IM theory is universally satisfied in the pattern (Cop&Hi, Cop&Hi). It is also not the pattern where the agents together produce the best possible consequences they possibly can; this price goes to the pattern (¬Cop&Hi, ¬Cop&Hi) (blue). Ironically, a theory which were to be CM in this situation would have to recommend the agents *not* to take a cooperative stance.

Finally consider how a costly cooperative stance affects 'breaking the tie'. Remember that the tie came up if one agent were to choose ¬Cop. Then the other could equally well choose ¬Cop as Cop. This tie could be broken by assigning some final value to choosing Cop. Then (see *Table 12*), if one agent were to choose ¬Cop&Lo, the IM choice by the other would be Cop&Lo, and similarly, if one agent were to choose ¬Cop&Hi, the IM choice by the other would be Cop&Hi.

Table 12. Breaking the tie

			B			
			Cop		¬Cop	
			Hi	Lo	Hi	Lo
A	Cop	Hi	10+2*v*	n.a.	10+*v*	*v*
		Lo	n.a.	6+2*v*	*v*	6+*v*
	¬Cop	Hi	10+*v*	*v*	10	0
		Lo	*v*	6+*v*	0	6

But now assume that Cop involves a (morally relevant) personal cost *c*. This means that every *v* should be replaced with *v-c* (Table 13). There are three cases to consider.

Suppose first that $v>c$. Then the situation in *Table 12* is not affected, the tie is still being broken (green combinations in Table 13).

Table 13. Breaking the tie in spite of costs

			B			
			Cop		¬Cop	
			Hi	Lo	Hi	Lo
A	Cop	Hi	10+2*v*-2*c*	n.a.	10+*v*-*c*	*v*-*c*
		Lo	n.a.	6+2*v*-2*c*	*v*-*c*	6+*v*-*c*
	¬Cop	Hi	10+*v*-*c*	*v*-*c*	10	0
		Lo	*v*-*c*	6+*v*-*c*	0	6

Suppose next that $v=c$. Then the tie is back (Table 9 above), because v and c cancel each other out.

Finally, suppose that $c>v$. The colors in Table 14 are changed from Table 13 to match this situation.

Table 14. Attempting to break the tie involves a net cost

			B			
			Cop		¬Cop	
			Hi	Lo	Hi	Lo
A	Cop	Hi	10+2*v*-2*c*	n.a.	10+*v*-*c*	*v*-*c*
		Lo	n.a.	6+2*v*-2*c*	*v*-*c*	6+*v*-*c*
	¬Cop	Hi	10+*v*-*c*	*v*-*c*	10	0
		Lo	*v*-*c*	6+*v*-*c*	0	6

Suppose the other agent were to choose ¬Cop&Lo. Then the IM choice by the other would be to likewise choose ¬Cop&Lo. Any IM theory is universally satisfied in this combination; but it is not CM. Hence, we are back in the original problem. A costly Cop does nothing to overcome it. On the contrary, in (¬Cop&Hi, ¬Cop&Hi) any IM theory is universally satisfied, and this combination is also CM. Any unilateral choice from a cooperative stance will make things worse.

7. Concluding Remarks

The analysis shows that the challenge raised by the original moral Hi-Lo problems stems from the fact that the agents choose independently of each other and are unable to influence each other. In this case, any IM theory is universally satisfied if each agent chooses Lo, but no such theory can be CM.

An attempt to give up the requirement that a theory should be IM and go for CM theory instead would lead to a catastrophe in high-stake Hi-Lo cases, and is thus not a viable option. However, it may be an option in low-stake cases, if the costs are not considered prohibitive.

It seems a better prospect to break the independence by attempting to connect cooperatively with the other agent in some way. However, this amounts to reframing of the choice situation and is thus changing the topic, at least if the reframing involves adding physical actions.

In the 'nice case', where each agent takes the cooperative stance, the IM option is also CM. For this reason, it seems natural to suggest that there should be a duty to take a cooperative stance. But if taking a cooperative stance involves a cost, the pattern where both do 'Hi' is no longer neither IM nor CM.

However, suppose that only one of the agents take a cooperative stance. If the other does Lo under these circumstances, the best the cooperative agent can do is also to do Lo. But this is not CM. To take a cooperative stance makes no difference in this case. The theory could be adjusted so as to assign taking a cooperative stance some final value in itself, such that it becomes the uniquely best answer for the agent. But the theory is still not CM. Hence, no IM theory, even if it is not exclusively act-oriented by allowing for unilaterally taking a cooperative stance, can be CM in all cases.

Suppose finally, by assuming that it is possible for an agent to successfully influence the other. The pattern (Hi, Hi) resulting from both taking a cooperative stance would be both IM and CM. But again, this only holds if taking a cooperative stance involves no cost. Collective maximization is satisfied only in the pattern (Hi, Hi) resulting from *not* taking a cooperative stance, in which individual maximization is also universally satisfied.

Going for a SCM theory rather than a merely CM one is to accept in cases of disconnection there is reason to make the best of the situation with those who are connected. This may be all right.

The lesson we draw is that a moral solution to a moral Hi-Lo problem is effective only if all agents become convinced and motivated through calls for coordinated action. But an attempt to convince others is likely to involve costs, whereby the coordinated action is no longer guaranteed to be collectively maximizing. Hence, it is not easy to have to both ways even if we reframe the original Hi-Lo case.

References

Bacharach, Michael (2006). Beyond Individual Choice: Teams and Frames in Game Theory. Princeton University Press.

Bykvist, Krister (forthcoming). High stake coordination problems – Do we need to reach beyond individual duties to solve them? In Säde Hormio & Bill Wringe (Eds.), Collective Responsibility: Perspectives on Political Philosophy from Social Ontology. Springer.

Chander, P. (2018) Game Theory and Climate Change. Columbia University Press.

Collins, S. (2019). Collectives' Duties and Collectivization Duties. Australian Journal of Philosophy, 91(2), 231-248.

Cripps, Elizabeth (2013). Climate Change and the Moral Agent: Individual Duties in an Interdependent World. Oxford University Press.

Feldman, Fred (1986). Doing the Best We Can: An Essay in Informal Deontic Logic. D. Reidel Publishing Company.

Gibbard, Allan F. (1965). Rule-utilitarianism: Merely an illusory alternative? Australasian Journal of Philosophy, 43 (2), 211-220.

Goodin, Robert (2012). Excused by the willingness of others? Analysis 71(1), 18-24.

Magli, A. C. & Manfredi, P. (2022). Coordination games vs prisoner's dilemma in sustainability games: A critique of recent contributions and a discussion of policy implications. Ecological Economics 192, 1-7.

Lewis, David Kellogg (1969). Convention: A Philosophical Study. Cambridge, MA, USA: Wiley-Blackwell.

Luce, Robert Duncan & Raiffa, Howard (1957). Games and Decisions: Introduction and Critical Survey. New York: Wiley.

Pellikaan H. & R. J. van der Veer (2002). Environmental Dilemmas and Policy Design. Cambridge: Cambridge University Press.

Pinkert, Felix (2015). What If I Cannot Make a Difference (and Know It). Ethics 125 (4), 971-998.

Portmore, Douglas W. (2018). Maximalism and Moral Harmony. Philosophy and Phenomenological Research (2), 318-341.

Rabinowicz, Wlodzimierz (1989). Act-utilitarian prisoner's dilemmas. Theoria 55, 1-43.

Regan, Donald (1980). Utilitarianism and co-operation. New York: Oxford University Press.

Sugden, Robert (2015). Team Reasoning and Intentional Cooperation for Mutual Benefit. Journal of Social Ontology 1 (1), 143–166.

Schwenkenbecher, A. (2021). Getting our act together. Routledge.

Zimmerman, Michael J. (1996). The Concept of Moral Obligation. New York: Cambridge University Press.