

WHAT IF I CAN NOT MAKE A DIFFERENCE (AND KNOW IT)¹

Felix Pinkert

— Accepted for publication in *Ethics* —

— pre-publication version, not for citation or circulation —

Abstract

Act Consequentialism faces difficulties in situations where several agents together produce less than optimal outcomes. Whenever a group of agents gratuitously fails to achieve best outcomes, then intuitively at least one of them must be morally at fault for this. However, in some cases, none of the agents individually can make a difference for the better by acting differently, and Act Consequentialism then gives the counterintuitive verdict that all involved agents act rightly. I discuss this problem in cases where agents know that they can not make a difference, and where Shelly Kagan's recent proposal to refer to agents' ignorance of this fact hence does not apply. Building on Donald Regan's and Michael Zimmerman's earlier work on this topic, I argue for supplementing Act Consequentialism with a requirement of Modal Robustness: Agents not only ought to produce best consequences in the actual world, but they also ought to be such that they would act optimally in counterfactual scenarios where others act differently. I argue further that the resulting Modally

¹I am grateful for the helpful comments on earlier versions of this paper from audiences at the Universities of Amsterdam, Munich, Oxford, Princeton, Stirling, and St Andrews. In particular, I am indebted to Ralf Bader, Marcia Baron, John Broome, Derek Ball, Mark Budolfson, Krister Bykvist, Cecile Fabre, Hilary Greaves, John Horty, Christian List, Brian McElwee, Paul McNamara, Tim Mulgan, Philip Pettit, Peter Singer, John Skorupski, Kai Spiekermann, Martin van Hees, and Bill Wringe for illuminating discussions on this topic.

Robust Act Consequentialism can be independently motivated by understanding it as a requirement of moral virtue, namely to reliably act rightly, and to do so for the right reasons. I then provide an inductive proof that Modally Robust Act Consequentialism identifies some moral fault whenever a group of agents gratuitously fails to bring about collectively optimal consequences. Finally, I discuss the limitations of my view, and contrast it with Donald Regan's Co-operative Utilitarianism.

We often collectively bring about bad outcomes. For example, factory farming is sustained by many people continuing to buy cheap supermarket meat, and anthropogenic climate change is brought about by the greenhouse gas emissions of millions of individuals. Intuitively, these bad outcomes are not just a matter of bad luck, but the result of some sort of moral fault. Yet in many of these situations, none of the individual agents could have made any difference for the better. Supermarkets only react to larger changes in demand than individual customers can bring about, and someone choosing to take the train instead of the plane may not make any morally relevant difference to the climate if a given threshold of greenhouse gas concentrations will be, or has already been, exceeded anyway. Consequentialism, at least of the simple Act Consequentialist type, here seems unable to find fault with anyone, as no one could have brought about better results by acting differently. Consequentialists hence face the challenge of somehow still finding fault with agents for the bad results which they together bring about.

In a recent paper, Shelly Kagan takes up Derek Parfit's discussion of such "no-difference cases". Kagan argues that standard Act Consequentialism can meet this challenge when agents do not know whether or not their contribution will make a difference.² He briefly argues that when agents *do* know that they can not make any difference for the better, Act Consequentialists should concede that everyone acts rightly even though all agents together could have produced better outcomes. Julia Nefsky, who strongly criticises most other parts of Kagan's argument in her follow-up paper, does not take issue with this claim, and likewise seems to assume that such cases of full knowledge pose no problem

²Shelly Kagan. "Do I Make a Difference?" In: *Philosophy and Public Affairs* 39.2 (2011), pp. 105–141, Derek Parfit. *Reasons and Persons*. Clarendon Press, 1984, ch. 3. Thresholds are only one of several possible reasons why individuals may be unable to make a difference in situations with multiple agents. The most important other such possible reasons are the seeming insignificance of small individual contributions (e.g. in pollution cases), and other agents' active frustration of one's attempts to improve outcomes (e.g. if other sellers will sell the weapons which one does not sell). While this paper does not discuss cases of seemingly insignificant contributions, frustration cases are covered by my proposed solution (see sections 3.1 and 7).

to Act Consequentialism.³

In this paper, I present a working example of agents who together bring about a collectively suboptimal outcome, but who each know that due to each other's actions and dispositions, they individually can not make any difference for the better (section 1). I then argue that contrary to Kagan's explicit and Nefsky's implicit position, Consequentialists should hold that some of the agents are morally at fault (section 2). I argue that this challenge can not be answered by Kagan's response, nor by any other common form of Consequentialism (section 3), and set aside an unnecessarily complicated and controversial solution (section 4). Building on the work of Donald Regan and Michael Zimmerman, I propose further that these cases demand a radically different Consequentialist approach, according to which agents are required not only to act optimally in the actual world, but also to be agents who *would* act optimally were others to act differently (section 5). I then show how this "Modally Robust Act Consequentialism" can be motivated by understanding it as Act Consequentialism plus a requirement of moral virtue (section 6), and subsequently provide an inductive proof to show that it is a satisfactory response for a wide range of cases with any number of agents (section 7). Finally, I discuss the limitations of my approach in cases of synchronous choice, and contrast it with Donald Regan's Co-operative Utilitarianism (section 8).

1 The Two Factories

Ann and Ben are owners of two factories which are located opposite each other on a river. Both agents opt for a production process which releases waste chemicals into the river, and thereby kill all fish in the river and destroy the livelihood of a fishing community downstream. The waste from one factory alone would suffice to kill all the fish, and adding the waste from the other factory does no additional damage whatsoever (say, the river flows into the sea where the waste is diluted below any harmful concentration).

³Julia Nefsky. "Consequentialism and the Problem of Collective Harm: A Reply to Kagan". In: *Philosophy and Public Affairs* 39.4 (2012), pp. 364–395.

If Ann or Ben were to unilaterally produce cleanly, this would make their production uncompetitive compared to the other factory, put them out of business, and destroy the livelihood of their employees. However, if they both were to produce cleanly, then this problem would not arise, and both factories would remain in business and the fishing community would flourish. Ann and Ben each employ 100 workers, the fishing community counts 100 people, and all that matters morally in this case are the livelihoods of the workers and fishermen. Further, the only available actions are either to pollute or to produce cleanly. In particular, Ann and Ben can not come together and suggest and discuss a common strategy.⁴ So far, the case is represented by figure 1.⁵

		Ben	
		ϕ_c	ϕ_p
Ann	produce cleanly (ϕ_c)	all workers and fishermen have their livelihood (300)	only Ben's workers have their livelihood (100)
	pollute (ϕ_p)	only Ann's workers have their livelihood (100)	only Ann and Ben's workers have their livelihood (200)

Figure 1: The Two Factories as strategic form game

The two entrepreneurs find themselves in a suboptimal equilibrium in a *coordination problem*: In the two (Nash-) equilibria where both agents pollute or produce cleanly,

⁴I here set this possibility aside because the structure of the problem case can be reproduced on the level of discussing a strategy: If we assume for each agent that the other would not cooperate in trying to find a common solution, and that trying to start a discussion has morally relevant cost, then the payoff structure precisely matches that of the actions of polluting or producing cleanly. Cf. section 8.2 for a discussion of situations where agents can communicate.

⁵Note that contrary to a typical strategic form game as used in game theory, we here only write one value per combination of strategies. This is because we are here only interested in how agents affect overall value, agent-neutrally considered. Thus the relevant “payoff” is the same for both agents.

neither Ann nor Ben can improve the outcome by *unilaterally* acting differently. But from this fact alone, it does not yet follow that Act Consequentialism can not find fault with Ann’s and Ben’s decisions to pollute, because it is not yet clear that they could not have improved matters by acting differently *simpliciter*. For example, if Ann were someone who *would* produce cleanly if Ben produced cleanly, then Ben *could* improve matters by producing cleanly, and Act Consequentialism would judge that Ben acts wrongly by polluting. The Two Factories becomes a challenge for Act Consequentialism only once we assume that Ann and Ben are both “uncooperative”, i.e. each would pollute *even if the other produced cleanly*.⁶ The additional stipulation of mutual uncooperativeness is needed because Act Consequentialism also takes into account how other agents would respond to the different actions available to an agent when it evaluates the consequences of these actions. And in The Two Factories, it is only if both agents are uncooperative that neither could have improved matters by acting differently.

Lastly, Ann and Ben are fully aware of this situation. They know all the relevant facts that determine the consequences of the available actions, i.e. the consequences of the different possible combinations of their actions and each other’s actions and uncooperativeness. Hence they know that they individually can not make a difference for the better by producing cleanly rather than polluting the river. These additional assumptions of known mutual uncooperativeness are shown by the extensive form game in figure 2.⁷

⁶For parallel examples, see Donald Regan. *Utilitarianism and Co-Operation*. Oxford University Press, 1980, p. 18; Christopher Woodard. *Reasons, Patterns, and Cooperation*. Routledge, 2008, p. 47; Michael J. Zimmermann. *The Concept of Moral Obligation*. Cambridge: Cambridge University Press, 1996, p. 257.

⁷As with the strategic form game, the “payoffs” represent overall value, agent-neutrally considered. Note further that the branches in the tree are meant to represent *counterfactual* relations and are not intended to represent any temporal order in which the agents act. However, as I discuss in section 8, certain morally desirable counterfactual relations presuppose that agents decide sequentially. I am indebted to Krister Bykvist for convincing me of the importance of tree representations.




Figure 2: The Two Factories as extensive form game

2 Consequentialism and collectively suboptimal outcomes

Cases like The Two Factories pose the following challenge to Act Consequentialism: Ann and Ben each individually could only have made matters worse by producing cleanly, as the other agent would then still have polluted the river, and the livelihoods of 100 workers in the cleanly producing factory would have been destroyed. Now Act Consequentialism holds that individuals morally ought to perform those actions which have the best expected consequences. Consequently, it judges that both Ann and Ben act rightly by polluting, even though they together could easily have brought about much better outcomes by both producing cleanly. Act Consequentialism thus lets them off the hook for together producing collectively suboptimal outcomes, and is hence at odds with the intuitively plausible claim:

On-the-hook: In any collection of agents who together gratuitously fail to bring about morally collectively optimal outcomes, some agent(s) must be morally at fault for the collectively suboptimal outcome.⁸

⁸In order to bypass the controversy about whether Consequentialism should be conceived of as maximising or merely satisficing, I stipulate that all discussed suboptimal outcomes are so bad and so strongly suboptimal that even those who conceive of Consequentialism as satisficing would concede that some fault needs to be found.

For a failure to bring about optimal outcomes to be “gratuitous” here means that the failure can not be explained by mitigating circumstances due to which we could not expect a given group to collectively act optimally. Typically, such circumstances consist in non-culpable misinformation or lack of information. Since Ann and Ben know all relevant facts, I assume that their failure to bring about optimal outcomes is gratuitous.⁹

On-the-hook is a widely shared assumption in the philosophical discussion of Consequentialism and no-difference cases. For example, Donald Regan’s Prop-COP is the mirror image of On-the-hook, Michael Zimmerman endorses the same claim, as does Derek Parfit in his discussion of collectively self-defeating moral principles. Shelly Kagan’s and Julia Nefsky’s recent discussions of no-difference cases are likewise motivated by an intuition along the lines of On-the-hook.¹⁰

On-the-hook is commonly defended along the following lines: The fundamental assumption of Consequentialism is that the right (i.e. the deontic status of actions) is solely determined by the good (i.e. by differences made to overall value in the world). In light of this connection, however, it should not be possible for each of us to act rightly while we together fail to produce optimal outcomes.¹¹ However, this motivation for On-the-hook is problematic, since it presupposes a Collective Consequentialist interpretation of the core Consequentialist position. According to this interpretation, the difference that many agents *together* make to value in the world bears on the moral status of individuals’ actions. Since this interpretation is denied by Act Consequentialists, basing On-the-hook on this interpretation already assumes that Act Consequentialism is mistaken. On pain of circularity, On-the-hook can then not pose a further challenge to Act Consequentialism.

⁹See the conclusion for cases with excusable collectively suboptimal outcomes.

¹⁰Nefsky, “Consequentialism and the Problem of Collective Harm: A Reply to Kagan”; Kagan, “Do I Make a Difference?”; Regan, *Utilitarianism and Co-Operation*, pp. 4–5; Parfit, *Reasons and Persons*, pp. 55–59; Zimmermann, *Moral Obligation*, ch. 9.

¹¹Parfit, *Reasons and Persons*, p. 54. See also Cf. Kagan, “Do I Make a Difference?”, p. 108; Regan, *Utilitarianism and Co-Operation*, pp. 18–21; Woodard, *Reasons, Patterns, and Cooperation*, pp. 45–47; and Zimmermann, *Moral Obligation*, 257f.

I contend that for this reason, On-the-hook should not be understood as a Consequentialist position at all. Instead, it should be understood as the contraposition of a second-order claim about morality in general, and hence as a desideratum for *any* moral principle. According to this claim, the relation between morality and overall value is such that if everyone always acted morally rightly, the world would be as good as it can be (as far as agents' influence is concerned). Whether or not individuals' moral obligations are in any way determined by consequences, they "conspire" to direct all agents together towards the best outcomes they collectively can bring about.¹² Conversely, if collectively optimal outcomes are not produced, at least one agent must be morally at fault.

Thus understood, On-the-hook has considerable intuitive appeal, and a moral principle which can accommodate this intuition is, other things equal, strongly preferable to a moral principle that can not accommodate it. So while I follow Kagan's claim that Act Consequentialists should see little appeal in condemning actions that are known to make the world no worse (see section 3.2), I hold that they should investigate other ways to accommodate On-the-hook.

In the following, I argue for an extended form of Act Consequentialism which is superior to Act Consequentialism and all other Consequentialist principles because it meets the challenge of accommodating On-the-hook without introducing any new problems. Before turning to my own proposal, however, I first discuss and reject two prominent nonsolutions and an unnecessarily complicated and controversial solution to this challenge.

¹²For the related "principle of moral harmony", which holds that moral obligations conspire to direct agents to collectively maximize a given group's *welfare*, see Fred Feldman. *Doing the Best We Can: An Essay in Informal Deontic Logic*. Dordrecht, Boston: D. Reidel Pub. Co., 1986, chapter 7; and Fred Feldman. "The Principle of Moral Harmony". In: *The Journal of Philosophy* 77.3 (1980), pp. 166–179.

3 Nonsolutions

3.1 Subjective Act Consequentialism

Kagan offers a solution for those no-difference cases where agents do not know if their contribution will make a difference. Such situations typically arise when agents do not know the precise level of the threshold at which an additional contribution would make a difference, or they do not know how many other agents will contribute. According to Kagan, such cases call for a subjective Act Consequentialist assessment, where the deontic status of actions is determined not by actual but by expected consequences.¹³ Now while this solution is effective for the cases Kagan considers, it does not apply to cases like *The Two Factories*, where the involved agents know that they can not make a difference by acting differently. Here objective and subjective Act Consequentialism both judge that Ann and Ben each act rightly, and hence both moral principles fail to accommodate *On-the-hook*.

Kagan argues that this limitation of subjective Act Consequentialism is unproblematic because no-difference cases where agents know that they in fact can not make a difference are rare and unrealistic. I contend that this claim is mistaken, and that Kagan is misled by focusing on ordinary consumer choice and his unrealistic example of the Friends of Chicken Consumption organisation which informs customers when their purchase of a chicken will not make any difference to the number of factory-farmed chicken in the world.¹⁴ By contrast, realistic examples of no-difference cases with full knowledge are market situations where supply or demand of some good is known to be inelastic over a range of prices. Buyers and sellers then know that were they to stop participating, others would step in and compensate by taking over their share of supply or demand.

To give an example of price-inelastic demand, the number of weapons bought by a

¹³Kagan, “Do I Make a Difference?”, pp. 125–127. Essentially the same argument is also advocated by Peter Singer. “Utilitarianism and Vegetarianism”. In: *Philosophy and Public Affairs* 9.4 (1980), pp. 325–337, pp. 335–336.

¹⁴Kagan, “Do I Make a Difference?”, pp. 127–128.

rich warlord and distributed to child soldiers is arguably determined largely by strategic considerations and the number of available children, and is constant over a range of prices. Every individual weapons supplier can know that if they decide not to sell to the warlord, then another supplier will sell instead. With one supplier less in the market, the price of weapons may rise slightly, but this will not affect the total number of weapons bought nor the harm done to the child soldiers and to whoever gets into their way.¹⁵ The same considerations hold for dealers selling highly addictive drugs, and for corrupt officials selling favours for bribes. Examples of transactions affected by price-inelastic supply are the purchasing of natural resources from conflict regions, (other) stolen goods, and favours from corrupt officials. Unless Consequentialists are willing to bite the bullet and hold that agents on the price-*elastic* side of these collectively harmful and morally highly problematic transactions (i.e. those agents who can not affect the amount of goods traded) are at no moral fault, they must meet the challenge of accommodating On-the-hook in no-difference cases with full knowledge.

3.2 Collective Consequentialism

One way to accommodate On-the-hook is to move away from Act Consequentialism to Collective Consequentialist moral principles. In determining what *we individually* ought to do, Collective Consequentialist principles consider the consequences of what *we together* can do. Examples of such principles are Consequentialist Generalization (which holds that agents ought to do what would have best consequences if everyone did it), and simple compliance-based Rule Consequentialism (which holds that agents ought to act in conformance with the set of rules which is such that if everyone complied with it, best consequences would ensue). These simple Collective Consequentialist principles accommodate On-the-hook by judging that Ann and Ben both act wrongly by polluting.

This achievement, however, comes at the price of judging that Ann and Ben each ought

¹⁵The most likely response of the rich warlord to rising weapons prices is to reduce personal expenditure.

This might make the warlord do less environmental damage by luxurious living, but this damage clearly is not the main reason for condemning such weapons sales.

to unilaterally produce cleanly even though the other would then still pollute. Simple Collective Consequentialist principles then implausibly require Ann and Ben to perform actions which produce no good whatsoever and destroy the livelihoods of 100 employees. I contend that a principle that accommodates On-the-hook by giving such reckless verdicts is not to be preferred over Act Consequentialism which fails to accommodate On-the-hook, but at least does not require agents to act recklessly.

These reckless verdicts can be avoided by moving to more sophisticated Collective Consequentialist principles, e.g. by including caveats against recklessness or considering both the differences that we individually and that we together could make as reasons for action.¹⁶ This response, however, boils down to modifying the principles to imply that Ann and Ben are not required to unilaterally produce cleanly and hence act rightly by polluting. Collective Consequentialists can thus avoid reckless verdicts only at the cost of not accommodating On-the-hook. The appropriate Consequentialist response to this dilemma then seems to be to opt for the lesser evil, namely to not accommodate On-the-hook.

4 A solution with baggage: Collective Wrongness

The above dilemma occurs because so far, we have only considered Ann and Ben's individual actions as potential bearers of moral fault, and these actions can only be either permitted or prohibited. Consequently, the only way to avoid the dilemma is to allow for moral fault to attach to something other than Ann's and Ben's individual actions of polluting.¹⁷ A promising approach to extend the scope of moral fault is to let Act Consequentialism cover not only individual actions, but also conjunctions of such actions across agents. According to this view, first proposed by Betsy Postow and Derek Parfit, it

¹⁶Cf. e.g. Tim Mulgan. *The Demands of Consequentialism*. Oxford University Press, 2001, p. 60; Woodard, *Reasons, Patterns, and Cooperation*, 107ff.

¹⁷This is the essence of the negative conclusion of Donald Regan's argument that no exclusively act-oriented moral principle can be fully adaptive (*Utilitarianism and Co-Operation*, pp. 105–123).

is Ann and Ben *together* who act wrongly by both polluting, because together they could have brought about better outcomes.¹⁸ Because asserting collective wrongness is consistent with saying that Ann and Ben *individually* both act rightly, this position avoids the recklessness problem.

Whether this view also satisfactorily accommodates On-the-hook, however, is questionable. In its rough formulation so far, Ann and Ben individually are let off the hook, as there is nothing for which we could reproach either of them on their own. This problematic responsibility gap can be amended by stipulating that the collective wrongness reflects on Ann and Ben individually. On such a view, Ann and Ben act rightly *qua individuals*, but are at fault *qua members* of the group they constitute, due to the wrongdoing of that group.¹⁹

This version of the collective wrongness response, however, overshoots the target. Consider a variant of The Two Factories where Ben is uncooperative, but Ann is cooperative. As far as Ann is concerned, the overall pollution is avoidable, as she would not pollute if Ben did not pollute either. However, Ben will pollute anyway, the fishermen still lose their livelihood, and Ann and Ben collectively act wrongly. It then seems implausible to hold that this fault reflects on Ann: After all, by being ready to bring about the collectively optimal outcome, she already does all that she can be expected to do. Furthermore, her fault *qua member* of the group of Ann and Ben together would arise solely from Ben's uncooperativeness. Ann's moral fault would then depend exclusively on facts about another agent, and Ann could not have avoided being at fault even with her best efforts and intentions.²⁰

¹⁸Derek Parfit. "What We Together Do". 1988. Unpublished manuscript, p. 7; Betsy Carol Postow.

"Generalized Act Utilitarianism". In: *Analysis* 37.2 (1977), pp. 49–52.

¹⁹See e.g. Christian List and Philip Pettit. *Group Agency*. Oxford, New York: Oxford University Press, 2011, p. 164, Marion Smiley. "Collective Responsibility". In: *Stanford Encyclopedia of Philosophy, The*. Ed. by Edward N. Zalta. Fall 2008. 2008, section 4.

²⁰In such cases of *unilateral* uncooperativeness, Act Consequentialism can actually find fault with Ben, without any reference to collective wrongdoing: Because Ann is cooperative and would produce

Consequently, for the collective wrongness response to be plausible, collective wrongness must reflect on individuals more selectively. The above discussion shows that we must discriminate between cooperative and uncooperative agents: Ben's uncooperativeness allows us to say that the collective wrongness reflects negatively on him, while the cooperative Ann must be let off the hook.²¹

In this improved form, the collective wrongness approach still relies on highly controversial ascriptions of collective wrongness and hence collective obligations. Even if we grant that *some* collectives of agents can be subject to moral obligations, it is highly doubtful that just any odd collection of agents can be subject to moral requirements. In order for the approach to fully accommodate On-the-hook, it would have to be shown that *all* cases of gratuitous failure to bring about collectively optimal outcomes are also cases where collective wrongness can correctly be attributed to the group in question.

I am doubtful that this challenge can be met. But rather than pressing this point further, in the following section I provide a solution that is superior either way. My proposal follows the improved collective wrongness approach by identifying agents' cooperativeness and uncooperativeness as facts that matter to the moral evaluation of these agents.

cleanly if Ben produced cleanly, Ben's action of polluting the river is suboptimal. So one could restrict the collective wrongdoing approach to account for cases of *mutual* uncooperativeness like The Two Factories where Act Consequentialism can not find fault with either Ann or Ben, and could leave the remaining cases to be tackled by standard Act Consequentialism. However, such a patchwork approach seems implausibly gerrymandered. Furthermore, because the collective wrongdoing approach goes amiss in cases of unilateral uncooperativeness, the worry arises that even in cases of mutual uncooperativeness where this approach gives the right verdicts, it does so for the wrong reasons.

The possibility of inescapable fault qua member of a collective that acts wrongly will not worry those who are at ease with "metaphysical guilt" that one bears solely in virtue of one's membership in a blameworthy collective (Smiley, "Collective Responsibility", section 4). This position, however, is rare, and at any rate not mine.

²¹Note that the same qualifications would have to be added to a collective wrongness approach that does not evaluate Ann and Ben's actions as a collective, but other facts about them collectively.

However, my approach condemns agents' uncooperativeness directly, as opposed to viewing uncooperativeness as a mere conduit through which collective wrongness reflects on individuals. My approach thus bypasses the controversies that come along with referring to collective wrongness, and has the further virtue of simplicity.

5 A better solution: Being at fault for what one would have done

Recall that it is Ben's uncooperativeness that makes Ann's act of pollution rightful, and *vice versa*. This is because by being agents who would pollute even if the other agent produced cleanly, Ann and Ben make it impossible for each other to achieve better outcomes by acting differently. A moral principle that condemns such uncooperativeness would then find fault in cases where Act Consequentialism can not do so. If Act Consequentialism is supplemented with such a principle, then it always finds fault when a group of agents fails to achieve collectively optimal outcomes – or so I will argue in the following.²²

Now Ben's uncooperativeness means that even if Ann *were* to produce cleanly, Ben *would not* produce cleanly as well (and vice versa). In this counterfactual situation, Ben would act suboptimally, as he then could produce better outcomes by producing cleanly as well.²³ Hence while Ben satisfies the requirements of Act Consequentialism in the actual world, he would violate them in the counterfactual situation where Ann produces

²²My proposed solution is in part inspired by Michael Zimmerman's "openness" requirement, which holds that agents ought to leave it up to others whether a collectively optimal outcome will be produced (Zimmermann, *Moral Obligation*, pp. 263 ff.). Properly specified, Zimmerman's and my own view are extensionally equivalent, but I contend that my view is more straightforward, more easily motivated, and more suitable for generalisation proofs.

²³As a terminological simplification, let the term "optimal" not only refer to best outcomes, but also to actions or combinations of actions that among all relevant alternatives bring about optimal outcomes (and analogously for "suboptimal").

cleanly (and vice versa for Ann).

This connection between uncooperativeness and counterfactual violations of the requirements of Act Consequentialism holds generally: If an agent is uncooperative and makes it impossible for others to bring about a collectively optimal outcome, then it must hold that were all others to play their part in an optimal collective response, the uncooperative agent would (actively or passively) frustrate this attempt to reach a collectively optimal outcome. By doing so, the uncooperative agent would act suboptimally, because her actions would bring about a less good outcome than the collectively optimal outcome. Hence the uncooperative agent would violate the requirements of Act Consequentialism.

Because of this connection between uncooperativeness and Act Consequentialism, we can capture the moral fault that is in play in The Two Factories by supplementing Act Consequentialism with a requirement of *modal robustness*, i.e. a *requirement to act according to the demands of Act Consequentialism not just in the actual world, but also in certain counterfactual scenarios*.²⁴ More precisely, the resulting requirement reads:

Modally Robust Act Consequentialism: An agent ought to act optimally in the actual world, and be such that for all possible combinations of the actions of other agents, if that combination were instantiated, she would act optimally in these circumstances.

Modally Robust Act Consequentialism accommodates On-the-hook in The Two Factories: Ann's uncooperativeness means that in counterfactual situations where Ben produces cleanly, Ann would act suboptimally, and likewise for Ben. Furthermore, this violation of the requirements of Modally Robust Act Consequentialism explains the suboptimal outcome, because universal satisfaction of the requirements of Modally Robust

²⁴This formulation is adapted from Philip Pettit's 2011 Uheiro lectures on robustly demanding *values* or *goods*. Here, the idea is that values such as friendship do not just require certain behaviour in the actual world, but also in counterfactual scenarios. My proposal differs by applying the idea of modal robustness to moral requirements and hence moral *rightness*.

Act Consequentialism would have guaranteed collectively optimal outcomes, for the following reason: First, the modal robustness part of this principle requires Ann to be such that she would produce cleanly if Ben did, and *vice versa*. If both agents satisfy this requirement, then the Act Consequentialism part of the principle requires each to produce cleanly. If they also satisfy this further requirement, then they both produce cleanly, thereby producing collectively optimal outcomes. Conversely, if Ann and Ben do not produce collectively optimal outcomes, then this fact is explained by at least one of them violating the requirements of Modally Robust Act Consequentialism.

To strengthen the case for Modally Robust Act Consequentialism, in the following two sections I first independently motivate Modally Robust Act Consequentialism, and second show that it accommodates On-the-hook not just in The Two Factories, but in a wide range of cases. In this discussion, I call agents who satisfy the requirements of Modally Robust Act Consequentialism in a given situation “cooperative”, and agents who violate them “uncooperative”.

6 Motivating Modally Robust Act Consequentialism

So far, I have supplemented Act Consequentialism with a modal robustness clause and have argued that the resulting Modally Robust Act Consequentialism accommodates On-the-hook in The Two Factories. In this section, I provide an independent motivation for Modally Robust Act Consequentialism. I argue that agents who satisfy the requirements of Modally Robust Act Consequentialism thereby display an important aspect of morally good or *virtuous* character, viewed from a Consequentialist perspective.

6.1 Consequentialist virtue and Global Consequentialism

The position I advocate differs crucially from a common Consequentialist reading of moral virtue. This view evaluates dispositions and character traits from a Global Consequentialist perspective, and understands as virtues those dispositions which are optimal for an agent to have. As I show in the following, Modally Robust Act Consequentialism

makes more stringent requirements on agents' dispositions than Global Consequentialism. Consequently, Global Consequentialist virtue does not capture the requirements of Modally Robust Act Consequentialism, and, as I shall argue, is insufficient for fully accommodating On-the-hook in cases where Act Consequentialism on its own can not find fault.

To begin with, note that satisfying the requirements of Modally Robust Act Consequentialism does not always lead to optimal outcomes. For example, in *The Two Factories*, outcomes would not be improved if Ann was cooperative and thereby satisfied the requirements of Modally Robust Act Consequentialism. This is because due to the fact that Ben uncooperatively pollutes the river, a cooperative Ann has to pollute as well if she is to save the livelihoods of her employees. The only difference between a cooperative and uncooperative Ann is that the former would make better outcomes *available* to Ben. But since Ben would not make use of that new possibility, making superior outcomes available to him does not in fact make the world better. Hence contrary to Modally Robust Act Consequentialism, Global Consequentialism here does not require Ann and Ben to be cooperative rather than uncooperative, and also does not find fault with either of them.

The Global Consequentialist about virtue can respond by holding that dispositions need to be evaluated not in single choice situations, but over longer time periods. Virtues are then those dispositions of an agent which in the long run lead to best outcomes. However, on this view, Modally Robust Act Consequentialism and Global Consequentialist virtue still come apart: First, if agents can have very fine-grained dispositions, then there are many competing possible dispositions that have just as good consequences as satisfying Modally Robust Act Consequentialism. For example, suppose that Ann and Ben are always cooperative *except in The Two Factories*. Because unilaterally cooperativeness in *The Two Factories* does not make the world any better, these dispositions have just as good consequences as fully satisfying the demands of Modally Robust Act Consequentialism. Hence Global Consequentialism can not find any fault with Ann and Ben's dispositions, and thus comes apart from Modally Robust Act Consequentialism

and does not accommodate On-the-hook.

Second, whether satisfying the demands of Modally Robust Act Consequentialism produces better outcomes in the long run than alternative dispositions depends on the prevalence of cooperativeness in other agents. If everyone were always disposed to frustrate any attempts to bring about better outcomes, then a wide range of choice situations would be similar to The Two Factories. Unilateral cooperativeness would then not bring about better outcomes than a range of alternative dispositions, even if we do not allow for extremely fine-grained dispositions. Contrary to Modally Robust Act Consequentialism, Global Consequentialism would then not require agents to always be cooperative, and hence can not accommodate On-the-hook in such situations.

Lastly, Global Consequentialists can avoid these limitations by understanding virtues as those dispositions which it would be optimal for *everyone* to have. On this account, satisfying the requirements of Modally Robust Act Consequentialism is a virtue, because (as I will argue in the next section), *universal* satisfaction of the requirements of Modally Robust Act Consequentialism by all members of a collective guarantees collectively optimal outcomes. The downside of this view is that modally Robust *Act* Consequentialism would then be motivated along *Collective* Global Consequentialist lines, which is an unattractive *ad hoc* patchwork of two theories that is hard to motivate: On this view, the deontic status of an agent's actions is solely determined by the consequences of *that individual agent's* actions, while the moral quality of her character traits is solely determined by the consequences of *everyone* having such traits. It is not clear why consequences should matter in such different ways when we assess actions and when we assess character. So in the following I argue that we can provide a more straightforward and convincing way of understanding the requirements of Modally Robust Act Consequentialism as a matter of moral virtue.

6.2 Consequentialist virtue as responsiveness to value

The alternative picture of moral virtue that I propose starts from the independently plausible claim that a morally good character includes reliably acting morally and being

sensitive to those facts which are morally relevant. This claim is independent from Act Consequentialism, and any moral principle can be supplemented by a modal robustness requirement: If a given moral principle is the right moral principle, then being a morally good agent includes that one in fact satisfies the demands of that principle. Furthermore, a morally good agent satisfies these demands for the right reasons, i.e. because she appropriately and effectively takes into consideration those facts which the moral principle identifies as morally relevant reasons that determine her moral obligations. Her satisfaction of these demands is then not coincidental or morally lucky, and she would reliably satisfy these demands in various counterfactual scenarios. Hence modal robustness of acting rightly is arguably part of moral virtue, no matter what specific moral principle one adopts. A moral principle like Act Consequentialism simply adds an account of which facts are the morally relevant reasons.

We can then view Modally Robust Act Consequentialism as Act Consequentialism plus a requirement of moral virtue to be appropriately and effectively responsive to the morally relevant facts. Just as people with psychopathic disorders who are insusceptible to typically deontological reasons are defective moral agents from a deontological standpoint, so uncooperative agents who do not appropriately and effectively care about promoting the good are defective moral agents from a Consequentialist perspective.

Modally Robust Act Consequentialism understood in these terms has two attractive features. First, the moral defect manifested by uncooperative agents is distinct from morally wrong action: Just as a person with a psychopathic disorder might get through life without violating any deontological requirements, so uncooperative agents might be morally lucky and in fact always act optimally and according to Act Consequentialism rightly. Thus a requirement to *be* a morally good agent in this respect stands next to the initial Act Consequentialist moral principle that concerns right *action*.²⁵ It is this twofold moral assessment which enables Modally Robust Act Consequentialism to accommodate On-the-hook while avoiding reckless verdicts. Second, Modally Robust Act Consequentialism does not consider any more facts as morally relevant than Act Consequentialism

²⁵I am indebted to Brian McElwee for stressing this point.

does, as it holds that the only morally relevant reasons are the consequences of individuals' actions. Modally Robust Act Consequentialism interpreted as a requirement of moral virtue only adds the second-order requirement for agents to appropriately and effectively care about these very same reasons, and thus has the virtue of theoretical simplicity.

Modally Robust Act Consequentialism, motivated by appeal to moral virtue, then responds to the challenge of accommodating On-the-hook in The Two Factories as follows: Under mutual uncooperativeness, Ann and Ben each act rightly by polluting the river, because only in this way can they remain competitive and maintain the livelihoods of their workers. However, their uncooperativeness shows that there is something wrong with them as moral agents: They do not satisfy the demands of Act Consequentialism reliably or modally robustly, and this shows that they do not appropriately and effectively care about the livelihoods of the workers and fishermen. That is, they either do not care about them at all, or they are weak-willed and their care would not always translate into optimal action. Because according to Act Consequentialism, the extent to which actions promote the good is the supreme (and only) moral consideration, and since a world with more people having livelihoods is better than a world with fewer people having livelihoods, Ann and Ben thus *each individually* show a morally problematic character trait.

7 Generalization about Modally Robust Act Consequentialism and On-the-hook

As argued above, Modally Robust Act Consequentialism accommodates On-the-hook in the specific case of The Two Factories. However, there is a vast range of possible cases where several agents each have arbitrarily many options for acting, and where they together gratuitously achieve only collectively suboptimal outcomes. In many of these cases, no individual could make any difference for the better by acting differently. If Modally Robust Act Consequentialism is to be a real improvement over Act Consequentialism in terms of accommodating On-the-hook, it must always find some moral fault in such no-difference cases. In this section of the paper, I prove that under certain epistemic

assumptions, Modally Robust Act Consequentialism meets this requirement.

My proof proceeds by induction: I first show that Modally Robust Act Consequentialism always finds fault when two agents gratuitously bring about collectively suboptimal outcomes in situations with any number of available actions and optimal combinations of actions. Second, I show that if Modally Robust Act Consequentialism always finds fault in situations where some number n of agents gratuitously bring about collectively suboptimal outcomes, then it also does so in all situations with $n + 1$ agents. By induction, it then follows that the same claim holds in all situations with $n \geq 2$ agents.

7.1 Induction Start: Modally Robust Act Consequentialism in situations with two agents

First, I establish the Induction Start, i.e. the claim that Modally Robust Act Consequentialism always finds fault in any situation where two agents gratuitously achieve only collectively suboptimal outcomes, irrespective of the magnitude and distribution of outcomes, the number of optimal collective responses, and the number of actions that are available to each agent.

So consider a case where two agents together gratuitously achieve only a collectively suboptimal outcome. Consider further all the different combinations of actions of the two agents which would have brought about an optimal outcome. Call all actions of a given agent which feature in at least one such optimal collective response her “potentially cooperative” actions. By contrast, “uncooperative” actions are actions which do not feature in any optimal collective response.

If there is only one optimal collective response, then each agent has only one potentially cooperative action available to her. Since only suboptimal outcomes are produced, it must hold that at least one of the agents performs an uncooperative action, and must therefore have had a potentially cooperative action available to her which she does not perform. If there is more than one optimal collective response, then at least one of the agents has more than one potentially cooperative action available to her. In both cases, there is at least one agent who has had a potentially cooperative action available to her that she

does not in fact perform, either because she performs an uncooperative action or because she performs another potentially cooperative action.

Now consider an agent a who could have performed such a (different) potentially cooperative action ϕ_c rather than the action ϕ_a that she actually performs, and ask whether she acted optimally by performing ϕ_a :

1. If she does not act optimally, then according to Act Consequentialism, she acts wrongly. *A fortiori*, Modally Robust Act Consequentialism finds fault with her.
2. If she does act optimally, then it must hold that had she performed the potentially cooperative action ϕ_c , outcomes would not have been better than the actual collectively suboptimal outcome. So had she performed ϕ_c , this would not have led to an optimal collective response. This is only possible if the other agent b would then not have performed a potentially cooperative action that would have completed an optimal collective response. Hence agent b would not have acted optimally in this counterfactual scenario. It follows that Modally Robust Act Consequentialism finds fault with agent b .

Hence either way, Modally Robust Act Consequentialism finds fault with at least one agent if a group of two agents gratuitously fails to produce collectively optimal outcomes.

7.2 Induction Step from n to $n+1$ agents

The argument for Induction Step requires that the assumption that agents know all relevant facts is made more precise, as follows:

Modally robust knowledge: The involved agents know all relevant facts.

If others acted differently, they would still have such knowledge, and this knowledge would likewise be modally robust.

The modal robustness requirement on agents' knowledge states that the agents can not affect the quality of each others' epistemic positions, i.e. deception or misleading each other are not among the available actions. This requirement needs to be recursive, i.e.

the modal robustness of an agent’s epistemic position must itself be unaffected by others’ actions, in order for step three of the below argument to be valid. Lastly, lest the assumption seems overly strong, for the purposes of my argument, “know” only needs to mean that agents have true beliefs and have such a high level of evidence that it does not make a difference to the verdicts of subjective Act Consequentialism whether the true belief is treated as fact or only as true with a very high probability.

To establish the Induction Step, grant, for the sake of argument, the following Induction Assumption: If a group of n agents gratuitously produces collectively suboptimal outcomes, then Modally Robust Act Consequentialism finds fault with at least one of the group’s members. Now consider a situation where $n + 1$ agents fall short of producing optimal outcomes. By the above considerations, there is then at least one agent a who has a potentially cooperative action ϕ_c available to her but who in fact performs another (potentially cooperative or uncooperative) action ϕ_a . Now we can again ask whether a acts optimally by performing ϕ_a rather than ϕ_c :

1. If a ’s performing ϕ_a is not optimal, then according to Act Consequentialism, a acts wrongly. *A fortiori*, Modally Robust Act Consequentialism finds fault with a .
2. If a ’s performing ϕ_a is optimal, then it must hold that had a performed ϕ_c , the remaining n agents would not together have completed an optimal collective response that includes a performing ϕ_c .
3. Now in this counterfactual situation, a doing ϕ_c is held fixed (as it is the antecedent of the above counterfactual). So we can reduce this counterfactual situation to a choice situation of the remaining n agents, while treating a doing ϕ_c as an external parameter. So had a done ϕ_c , the remaining n agents would have been in a situation where they together failed to bring about collectively optimal outcomes. Moreover, due to modally robust knowledge, in this situation, they could not have been excused due to ignorance of the actions of other agents, and their failure would thus have been gratuitous. Hence by the induction assumption, Modally Robust Act Consequentialism *would* have found fault with at least one of these remaining

agents in the counterfactual situation where a does ϕ_c .

4. To see that Modally Robust Act Consequentialism also *in fact* finds fault with these agents, consider an agent b that Modally Robust Act Consequentialism would have found fault with if a had done ϕ_c . Modally Robust Act Consequentialism would have found fault with b for one of two reasons. First, b would have acted suboptimally in the counterfactual situation where a does ϕ_c . In this case, b would have acted wrongly according to Act Consequentialism, and in the actual world, b violates the Modal Robustness requirement of Modally Robust Act Consequentialism. Second, b would have acted optimally but would have been such that in the counterfactual situation where a does ϕ_c , it holds that there is then at least one further counterfactual combination of actions of the other $n - 1$ agents to which b would not respond optimally. Modally Robust Act Consequentialism *would* then find fault with b on grounds of violating the Modal Robustness requirement. Now take one such counterfactual combination of actions, and add a doing ϕ_c to it. This gives us another possible combination of actions such that it holds *in the actual world* that were this combination instantiated, b would not respond optimally. Hence Modally Robust Act Consequentialism *in fact* finds fault with the would-be culprit.
5. So if a 's performing ϕ_a is optimal, then Modally Robust Act Consequentialism finds fault with at least one of the remaining n agents.
6. Thus whether or not a 's performing ϕ_a is optimal, Modally Robust Act Consequentialism finds fault with at least one agent.

The above argument gives us the Induction Step: If Modally Robust Act Consequentialism identifies some moral fault in all situations where n agents gratuitously fail to bring about optimal outcomes (the Induction Assumption), then it also finds fault in all situations where $n + 1$ agents gratuitously fail to produce collectively optimal outcomes. Now since Modally Robust Act Consequentialism finds fault in all situations where $n = 2$ agents gratuitously fail to produce collectively optimal outcomes (the Induction Start), by induction it also finds fault in all situations where any number $n \geq 2$ agents gratu-

itously brings about only collectively suboptimal outcomes. Hence given modally robust knowledge, Modally Robust Act Consequentialism accommodates On-the-hook in any no-difference case with $n \geq 2$ agents.

Note that the above proof makes no reference to details of the strategic choice situation of agents, and hence does not assume that the situation is a coordination problem like The Two Factories. It thus applies to all cases where several agents together gratuitously bring about collectively suboptimal outcomes and have modally robust knowledge about each others' actions. These cases then include situations of active frustration, as in the market examples discussed in section 3.1.

8 Limitations and rivals

Modally Robust Act Consequentialism is limited as a response to the challenge of accommodating On-the-hook because the above generalisation proof relies on the assumption of modally robust knowledge. In this final section, I discuss cases where this assumption does not hold, and contrast Modally Robust Act Consequentialism with a prominent rival view which aims to do better.

8.1 Knowledge and synchronous choice

The assumption of modally robust knowledge most notably breaks down in cases of synchronous choice, like the following specification of The Two Factories:

The Two Factories*: Ann and Ben synchronously decide to pollute the river. When Ann makes her decision, she correctly infers from Ben's past conduct and expressions of intent that Ben will pollute the river.

If contrary to his past conduct and expressed intentions, Ben were not to pollute the river, then there is no way for Ann to know about this fact when she makes her choice. So all that Ann can do is to decide whether or not to pollute, depending on her beliefs about Ben's actions, but independently from what Ben in fact does. She thus can not both be such that she would act optimally were he to pollute, and were he not to pollute.

So in *The Two Factories**, Modally Robust Act Consequentialism asks the impossible and hence can not apply. Since modally robust knowledge can never hold in situations of synchronous choice, my discussion of Modally Robust Act Consequentialism so far has implicitly assumed that agents act in sequence.²⁶

The most promising way to avoid the problem of requiring impossible dispositions is for Modally Robust Act Consequentialism to require agents to be such that for all counterfactual situations, they would act optimally *unless being such would come at the cost of acting optimally in the actual world*. In *The Two Factories**, the only way for Ann to be such that she would act optimally in the counterfactual situation where Ben produces cleanly is for her to produce cleanly in the actual world. But given that Ben in fact does pollute the river, this action of Ann's would be strongly suboptimal in the actual world. When morally good agents who care about the good are confronted with such a choice between either acting optimally in the actual world, or acting optimally

²⁶Once we explicitly consider sequential choice, the following complication arises for my proof by induction: Those agents who get to choose earlier can determine not only the options faced by agents who get to choose later, but might also influence the order in which others need to make their choices. In essence, my proof is unaffected by this possibility, as it makes no reference to the specific order in which agents choose, and does not assume that this order is fixed. Instead, the proof only implicitly assumes sequential choice due to its reliance on modally robust knowledge. This assumption is compatible with a variable ordering of agents' choices, and as long as this assumption is granted, the proof is unaffected.

However, cases with variable ordering of choices may well introduce significant cognitive challenges for agents to obtain modally robust knowledge, so that the assumption of such beliefs may be unrealistic or, given the cost of obtaining such beliefs, unreasonable. These situations then provide an intermediate class of cases between situations where modally robust knowledge is plausible, and scenarios where it is impossible. I can here only express the conjecture that the future approach to cases of synchronous choice which I outline in the conclusion can be adapted to these intermediate cases as well. I am indebted to an anonymous referee for *Ethics* for pointing me to the possibility of variable orderings.

in counterfactual scenarios, they would act optimally in the actual world. So Modally Robust Act Consequentialism, understood as a requirement of moral virtue, must be modified to give preference to actual optimal action over counterfactual optimal action.²⁷ *However, such a modified Modally Robust Act Consequentialism* can not find any fault with Ann and Ben on the grounds that they would not act optimally in counterfactual situations where the other factory owner produces cleanly.* Hence this view does not accommodate On-the-hook in The Two Factories*.

8.2 Co-operative Utilitarianism as rival view

This limitation of Modally Robust Act Consequentialism makes Donald Regan's Co-operative Utilitarianism a particularly salient rival view, as it is explicitly developed for cases of synchronous choice. According to this view, agents ought to be cooperators who are willing to do their part in whatever the best response of all cooperators turns out to be. They further ought to follow a sophisticated procedure to identify the other cooperators and the best response available to all cooperators, and to then play their part in this response.²⁸ Since the ambition of Co-operative Utilitarianism is to accommodate On-the-hook in synchronous choice cases, it promises to be superior to Modally Robust Act Consequentialism in these cases. However, as I argue in the following, this promise remains unfulfilled, as Co-operative Utilitarianism tacitly assumes a sort of sequential choice situation.

For Co-operative Utilitarianism to be applicable to cases like The Two Factories*, it must be possible for Ann and Ben to satisfy its requirements. Now identifying the other cooperators is not a purely mental activity that agents can do all on their own before making their decisions. Instead, it consists in approaching other agents with proposals for collective strategies, and in asking whether they would be willing to participate. So if we apply Co-operative Utilitarianism in The Two Factories* and other cases of

²⁷Cf. Zimmermann, *Moral Obligation*, pp. 266–268 for analogous caveats for Zimmerman's openness requirement.

²⁸Regan, *Utilitarianism and Co-Operation*, ch. 8–10, especially pp. 148 and 153.

synchronous choice, we implicitly assume that there is a prior choice situation in which agents can communicate their willingness and assess each others' willingness to cooperate. If such a prior choice situation is not given, Co-operative Utilitarianism must be modified to not require an impossible identification of other cooperators. And just like Modally Robust Act Consequentialism*, such a modified Co-operative Utilitarianism* is no longer guaranteed to accommodate On-the-hook.

Furthermore, I contend that Co-operative Utilitarianism also does not have an edge over Modally Robust Act Consequentialism* in situations of synchronous choice where prior communication is possible. This is because such prior communication allows agents not only to carry out the procedures required by Co-operative Utilitarianism, but also to directly coordinate their actions.²⁹ Applied to The Two Factories*, agents then have a prior “coordination choice” to make, where they decide what to do with their ability to communicate, followed by the original synchronous “pollution choice”, which concerns their actions with regard to the river.

Now if Ann and Ben first face a coordination choice, then Modally Robust Act Consequentialism* can also accommodate On-the-hook: Suppose that Ann and Ben satisfy the requirements of Modally Robust Act Consequentialism* in both choice situations. Now if in the pollution choice, they were confident that the other agent was going to produce cleanly, then they would both produce cleanly. Consequently, the best that Ann and Ben can do in the coordination choice is to create that confidence in each other, which they can do by agreeing to produce cleanly. Since the coordination choice takes the form of a possible conversation, it allows agents to act in sequence, so that the caveats of Modally Robust Act Consequentialism* do not apply. We can then refer to the above generalisation proof and conclude that since Ann and Ben satisfy the requirements of Modally Robust Act Consequentialism*, they will do the best they together can do in the coordination choice. Ann and Ben will hence agree to produce cleanly, and then follow through with that agreement in the pollution choice and thereby produce collectively optimal outcomes. Conversely, if they do not bring about collectively optimal outcomes

²⁹Again I am indebted to Krister Bykvist for this observation.

even though a coordination choice was available, then at least one of them violates the requirements of Modally Robust Act Consequentialism*. Hence under the assumption that Ann and Ben face a coordination choice, Modally Robust Act Consequentialism* and Co-operative Utilitarianism both accommodate On-the-hook.

My conjecture is that this result generalizes: Modally Robust Act Consequentialism* and Co-operative Utilitarianism* can accommodate On-the-hook in the same cases, and so the limitations of the former do not give support to the latter. Given the potential for extensional equivalence of the two views, the main advantage of Modally Robust Act Consequentialism over Co-operative Utilitarianism is the conceptual simplicity of the former, and the possibility of motivating it as a requirement of moral virtue.

9 Conclusion and outlook

I have shown how Act Consequentialists can find fault with some agent in all cases where a group of agents who have modally robust knowledge of all the relevant facts gratuitously brings about collectively suboptimal outcomes, even if the agents individually can not make any difference for the better due the uncooperativeness of others. Act Consequentialists, I have argued, can supplement their view of morally right action with a basic requirement of moral virtue, which holds that agents ought to be such that they act rightly not only in the actual world, but also in counterfactual scenarios where others act differently. The resulting Modally Robust Act Consequentialism correctly finds fault with agents who fail to be morally virtuous in this respect and who thereby make collectively optimal outcomes inaccessible to others.

The main limitation of my proposal is its reliance on the assumption of modally robust knowledge, which breaks down in cases of synchronous choice. I have sketched how despite initial appearances, the competing view of Co-operative Utilitarianism fares no better than Modally Robust Act Consequentialism in such situations: I have conjectured that both views can accommodate On-the-hook in all and only those cases of synchronous choice that are preceded by coordination choices, but not in cases of isolated synchronous

choice. A rigorous generalisation of this claim, however, needs to wait for a future paper.

Provided that such a generalisation is possible, the main limitation of Modally Robust Act Consequentialism with regard to accommodating On-the-hook is that it does not apply to isolated synchronous choice cases where no prior coordination is possible. I contend that these cases can be accommodated by making independent use of the above motivation of Modally Robust Act Consequentialism via moral virtue. The idea is to approach these cases by means of epistemic game theory:³⁰ For any choice situation with a set payoff structure, this approach first identifies those sets of pairs of agents and probability assignments on others' actions for which universal satisfaction of Act Consequentialism is compatible with producing collectively suboptimal outcomes. For example, in *The Two Factories** without a prior coordination choice, Act Consequentialism permits Ann and Ben to both pollute if they attach a probability of less than $\frac{1}{3}$ to the other agent producing cleanly. For each of these "epistemic profiles" of a possible situation, we then ask whether it is explained by a moral shortcoming of some agent which can be captured by the understanding of moral virtue advocated above (e.g. if Ben is known to be more likely than not to be uncooperative, because he does not care about the good, or if Ann negligently did not acquire enough information about Ben), or by mitigating circumstances which make collectively suboptimal outcomes non-gratuitous (e.g. if Ann mistakenly but non-culpably believes that Ben is uncooperative). If it can be shown that all problematic epistemic profiles can be explained by one of these factors, then Act Consequentialists can accommodate On-the-hook: The failure to bring about collectively optimal outcomes is then due to some moral shortcoming that Act Consequentialists can account for, or it is not gratuitous and no fault needs to be found. Developing this argument, however, is a task for yet another paper.

³⁰I am indebted to Christian List for helping me specify the approach I am after.