Overdemanding Consequentialism? An Experimental Approach

MARTIN BRUDER
University of Konstanz

ATTILA TANYI
University of Konstanz and University of Liverpool

According to act-consequentialism the right action is the one that produces the best results as judged from an impersonal perspective. Some claim that this requirement is unreasonably demanding and therefore consequentialism is unacceptable as a moral theory. The article breaks with dominant trends in discussing this so-called Overdemandingness Objection. Instead of focusing on theoretical responses, it empirically investigates whether there exists a widely shared intuition that consequentialist demands are unreasonable. This discussion takes the form of examining what people think about the normative significance of consequentialist requirements. In two experiments, the article finds that although people are sensitive to consequentialist requirements and, on average, find more extreme demands less reasonable, the level of disagreement with consequentialism falls short of qualifying as a widely shared intuition, even when demands are the highest. The article then ends with a general discussion of possible objections to its methods and its findings.

I. THE OVERDEMANDINGNESS OBJECTION

Act-consequentialism is the view that the right thing to do in any situation is the act with the best consequences as judged from an impersonal point of view. Act-consequentialism is thus a single-principle view; it is exhausted by what is called the optimizing principle of beneficence. Many disagree with this principle, however, emphasizing the significance of the personal in moral life. In this article our subject matter will be an objection that derives from this complaint and that can be employed against any form of act-consequentialism (henceforth 'consequentialism', unless qualifier needed).

We would like to thank audiences in Konstanz, The Hague, Nottingham, Bayreuth and Lucca, as well as all those who have commented on prior versions of the manuscript. We are also grateful to two anonymous reviewers of *Utilitas* for comments that have made us rethink many of the points in the article. The research reported in this article was funded by several grants from the Zukunftskolleg at the University of Konstanz. It was also supported by a grant to Attila Tanyi from the German Research Foundation (Grant number: TA 820/1-1).

© Cambridge University Press 2014
Utilitas Vol. 26, No. 3, September 2014
doi:10.1017/S0953820814000065

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

Let us call this charge the *Overdemandingness Objection* or OD.[1] OD is built upon two pillars: first, that consequentialism is unreasonably demanding; second, that an adequate morality cannot be unreasonably demanding. Consequentialism requires the agent to promote the good until further efforts would burden her as much as they would benefit others. However, the situation that determines what would be best overall is far from ideal: today's world involves, for example, significant levels of poverty, which charitable donations are insufficient to eradicate. At the same time, those institutions that might make things better are not effective, either domestically or internationally. Given that acting to alleviate poverty is likely to have, in sum, more positive consequences than pursuing individual goals and projects, it seems unavoidable that, if one fully accepts consequentialism, one must devote most of one's resources to humanitarian work or to the support of institutions that carry out this work.

At the same time, so advocates of OD assume, people have a firmly held intuition that this cannot be demanded of them, that they should not be required to forego all or most of their individual projects for moral causes. This is the second pillar of the objection. If this intuition indeed exists, it grounds a constraint on admissible moral theories, requiring them to avoid unreasonable demands. If they do not, the conclusion follows that these theories should not be allowed to guide people's conduct. OD is an attempt to articulate this constraint.

II. THE EMPIRICAL BACKGROUND OF OD

OD, as presented above, can be put more formally as the following argument:[2]

(1) Consequentialism makes demand D.
(2) Intuitively, demand D is unreasonable.
(3) Therefore, consequentialism makes unreasonable demands.
(4) No acceptable moral theory makes unreasonable demands.
(5) Therefore, consequentialism is an unacceptable moral theory.
(6) Therefore, we have reason to reject consequentialism.

[1] OD is one of those charges that are most clearly stated by those who oppose it. For an early statement see Henry Sidgwick, *The Methods of Ethics*, 7th edn. (London, 1907), p. 87; for a recent statement see Garrett Cullity, *The Moral Demand of Affluence* (Oxford, 2004), ch. 1. For further references see Brad Hooker, 'The Demandingness Objection', *The Problem of Moral Demandingness: New Philosophical Essays*, ed. T. Chappell (London, 2009), pp. 148–63, at p. 162 n. 4, and Alan Carter, 'Is Utilitarian Morality Necessarily too Demanding?', *The Problem of Moral Demandingness*, pp. 163–85, as well as the works to be cited in the next section.

[2] A similar presentation of OD can be found in Tim Mulgan, *The Demands of Consequentialism* (Oxford, 2001), p. 24 and in his *Understanding Utilitarianism* (Stocksfield, 2007), p. 99.

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

This more detailed structure suggests the possible ways of responding to OD. The dominant approach is to deny premise (1). This *strategy of denial* holds that consequentialism does not make the demands attributed to it, either because such attribution rests on false empirical beliefs or because consequentialism properly understood has a structure from which no such demands follow.[3]

Taking a different stance, the *strategy of extremism* does not deny that consequentialism makes the demands in question; what it denies is that these demands are indeed unreasonable. People might think so, but then they are wrong, because, as it is typically argued, their intuitions about these matters are unreliable. Thus, it is argued that the intuition premise (2) refers to, let us call it the OD-intuition, rests on lack of information, lack of clear thinking, lack of imaginative empathy or on some psychological 'failure', or that it tracks something entirely different from issues of excessive demands, or that intuitions in general are problematic because of their evolutionary background. Consequently, the inference to sub-conclusion (3) is warranted only if it is made from people's point of view. But if this is so, then the argument does not succeed since such unreliable intuitions offer no ground for rejecting consequentialism.[4]

Finally, the *strategy of moderation* points out that the argument implicitly assumes – already in its choice of the term 'unreasonable', instead of 'excessive' or 'extreme' – that consequentialism adheres to the thesis that moral reasons override non-moral reasons. But this need not be so: consequentialism can either make no claims about reasons, or even if it does, it does not have to hold that people have decisive reason to follow its precepts.[5] If either of these claims is true, then it is not the case that consequentialism makes unreasonable demands: again, people might assume so, but then they are wrong. Therefore

[3] The first approach is critically discussed in Mulgan, *The Demands of Consequentialism*. Influential examples of the second approach include Samuel Scheffler, *The Rejection of Consequentialism*, rev. edn. (Oxford, 1994) and Michael Slote, 'Satisficing Consequentialism', *Proceedings of the Aristotelian Society, Supplementary Volume*, 58 (1984), pp. 139–63.

[4] See e.g. Peter Singer, 'Famine, Affluence and Morality', *Philosophy and Public Affairs* 1 (1972), pp. 229–43; Shelly Kagan, *The Limits of Morality* (Oxford, 1989), and David Sobel, 'The Impotence of the Demandingness Objection', *Philosophers' Imprint* 7 (2007), pp. 1–17.

[5] The recent position of *scalar-consequentialism* opts for the first alternative. See Alastair Norcross, 'Reasons without Demands: Rethinking Rightness', *Contemporary Debates in Moral Theory*, ed. J. Dreier (Oxford, 2006), pp. 38–53. The second alternative is less controversial, and several consequentialists have held it. Note, though, that the position of *moral rationalism* holds that this is conceptually not possible; see Douglas W. Portmore, *Commonsense Consequentialism: Wherein Morality Meets Rationality* (Oxford, 2011) for this claim. See also Krister Bykvist, *Understanding Utilitarianism* (London, 2010), pp. 101–2 for a discussion of this strategy of moderation.

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

the inference to sub-conclusion (3) is unwarranted. Alternatively, 'reasonable' might just stand for 'excessive' or 'extreme', in which case, however, premise (4) does not hold. For consequentialism, although it makes excessive demands, does not claim overriding force for those demands. And, contrary to premise (4) so understood, a moral theory that makes excessive demands but does not claim to be overriding is acceptable. After all, why would it be objectionably demanding what one, all things considered, is not required to do?

However, we intend to keep the structure of consequentialism unchanged and we do not want to deny that the OD-intuition is reliable and can be used for supporting the conclusion of the argument. Nor do we want to question the assumption that consequentialism accepts the overridingness of moral reasons over non-moral reasons or that the assumption is needed to make OD work. All of these approaches have their – often much discussed – problems; it would therefore bring fresh impetus to the discussion if one approached the problematic from a different angle. What would this approach be? We propose to question premise (2) by using empirical research. After all, what the premise does is to assert the existence of an intuition, i.e. something that people intuit to be the case. Hence, it appears natural to investigate whether people do indeed have this intuition. For if they do not, premise (2) can be rejected and the argument of OD collapses.

There are two complications that must be dealt with before we can move on. The first is that what empirical research can establish is whether the OD-intuition is *widely shared*: how many people endorse premise (2) of the argument of OD? In other words, this kind of research has a statistical nature: it operates with percentages. Now, it cannot be reasonably expected that this research could establish that *no one* has the relevant intuition or that, conversely, *everyone* has it, which is what, strictly speaking, is needed to establish or deny the truth of premise (2). Is this a problem?

We do not think so: percentages matter. On the one hand, intuitive objections like OD appeal to those who have the intuitions in question. Hence the more people have the intuition, the more persuasive the given objection is. On the other hand, for many, intuitions are not merely devices of persuasion, but have justificatory power: they have the potential to establish the validity of a conclusion because, say, they serve as evidence for the truth of this conclusion. In the absence of any reasonable expectation that we can find uniformity – rejection or acceptance – in people's intuitions, percentages come to play a crucial role. To put it simply: the more people have the intuition in question, the more likely it is that they are getting things right and the validity of the objection can be established. Therefore, what we will do is to use a statistical criterion in assessing the validity of premise (2). For the

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

intuition to be able to support a persuasive argument and for it to serve as reliable evidence, it is not enough if only few people have it. It should be shared by a strong majority. What counts as a 'strong majority' is difficult to define. However, other research on moral intuitions has often found the agreement of around 90% of the respondents.[6] Yet, because we cannot establish either the uniform existence or the total absence of the relevant intuition, we will tone down our conclusion: we will not claim that we refute the idea that there exists an OD-intuition as premise (2) claims, but only that our research throws doubt on this claim.

The second complication concerns the intuitive background of OD. It is clear that the OD-intuition, given our acceptance of the role reasons play in the argument, is about what one is – all things considered – required to do; in short: what one has decisive reason to do. But is this the only intuition that has a bearing on the failure or success of OD? We do not think so. We believe that the *moral* intuition of what one is morally required to do should also be considered in assessing the above argument. This is because one's judgement of reasonableness, i.e. what one has decisive reason to do, and one's judgement of what one is morally required to do, can come apart.[7] People can deny what we called the overridingness thesis above and assert that moral reasons can be overridden by non-moral reasons. That is, people can hold that they have decisive reason to do X, that this is what it is reasonable to require, while maintaining that it is action Y that is required by morality. And, of course, the reverse can also happen. In the particular context, this means that people can affirm that a consequentialist requirement is reasonable, while taking it to be not demanded by morality; or they can affirm that a consequentialist requirement is unreasonable, while nevertheless taking it to be moral.

Now, if people can have conflicting intuitions about what they are morally required to do and what they have decisive reason to do, the moral intuition could be claimed to have a bearing on the validity of OD. This is because OD is used to reject consequentialism as a moral theory because it is unreasonably demanding: but what happens if people, although they intuit consequentialism to be unreasonably demanding, still think it gives rise to legitimate moral requirements? In this case, it seems, we have a moral intuition conflicting with

[6] E.g. Joshua Knobe, 'Intentional Action in Folk Psychology: An Experimental Investigation', *Philosophical Psychology* 16 (2003), pp. 309–24.

[7] We mentioned the opposite of this in our discussion of the strategy of moderation above, i.e. that people can think that whatever they are morally required to do, they also have the most reason to do. But of course there is no reason to think that this is the only way people can think of morality and its relation to reasons for actions.

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

consequentialism's alleged claim to the overridingness of its reasons (which, recall, we assumed to be the case so as to avoid endorsing the strategy of moderation). But if the intuition referred to in premise (2) is taken seriously, although it too conflicts with consequentialism's alleged claim to overridingness, then why not take seriously *this* moral intuition? This is important, because if the intuition is taken seriously, then it has a clear bearing on OD, since it can happen that although one intuition, that concerning the reasonableness of its demands, speaks against consequentialism, another, that concerning the morality of its demands, does not. And since OD is supposed to lead us to the rejection of consequentialism as a *moral* theory, this is significant. For it suggests that OD can be false, *even if* the intuitive unreasonableness of consequentialist demands is confirmed.[8] Consequently, although our research has its focus on the validity of premise (2), i.e. on how widespread the OD-intuition is, we will, whenever the above contrast occurs, take it into account in judging the validity of OD itself.

After these preliminaries, we can now turn to the substantive part of our article. This is how we will proceed. In what follows (sections III and IV) we will first present the findings of the experimental research we have conducted on the validity of premise (2). On the basis of these findings we will claim that there is reason to doubt the truth of the premise. We will then move on to a more general discussion of our argument, also concerning its methodology and the interpretation of its results (section V). We will end the article with a summary and some concluding remarks (section VI).

III. EXPERIMENT 1: A SCENARIO STUDY

We have carried out two studies to investigate the validity of premise (2): the first is a scenario study, which is frequently used in experimental research in philosophy and moral psychology; the second is a study based on an experimental game, the paradigm of which is adapted from behavioural economics. In each case, we describe the procedure (why the study was designed the way it was and how it was conducted), our predictions (the hypotheses we set out to test), and the results of our experiments, including a brief discussion (with a more general discussion, concerning possible problems with our approach, following in section V).

[8] How about the case when the demands of consequentialism are found intuitively reasonable but not demanded by morality? Although this case is interesting – not least because one is curious why these demands are contrary to morality – it does not have relevance for us. For though the moral intuition now speaks *against* consequentialism, it does so on grounds other than its alleged overdemandingness. And this has no bearing on OD and hence is not something we need to be concerned with in this article.

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

Procedure

We designed an online survey in which participants were confronted with fictitious scenarios that involved a decision between either following consequentialist requirements or acting in their own interest. We induced consequentialist demands by varying how attractive it was for participants to act in their own interest (narrowly enough conceived so that it excludes impersonal moral interests). If acting in one's own interest is very attractive, this results in the fact that the alternative action required by consequentialism places higher demands on an agent than if the non-consequentialist option is unattractive.

We distinguished two different kinds of demands corresponding to different kinds of attractiveness.[9] First, acting in one's self-interest may either be *objectively* attractive or not. By this we mean that doing something that is associated with a high reward (e.g. in terms of monetary value) is generally preferred over doing something that is associated with a lesser reward. Consequentialist requirements that compete with actions associated with a high reward should therefore be perceived as more objectively demanding than actions associated with a low reward. For example, most would agree that helping someone in need of money is more – objectively – demanding if the amount required to help effectively is high rather than low. However, even identical rewards may mean different things to different people and be differentially attractive to them, in the sense that people may feel differently about them. That is, second, alternative actions may differ in their *subjective* attractiveness; in contrast to the objective attractiveness of alternative actions, their subjective attractiveness has a certain – albeit not uniform – phenomenology. Again, consequentialist requirements that compete with actions of high subjective attractiveness should be perceived as more subjectively demanding than those competing with actions of low subjective attractiveness. For example, helping someone in desperate need of a thousand dollars will be less – subjectively – demanding if one does not care much about money than if one cares deeply.

[9] Although it helped us to design our experiments, our main reason for making this distinction lay in our interest in the role of emotions not reported in this article. We hypothesized that whereas subjective demands should be strongly related to anticipated individual emotions in relation to the decision, objective demands should be more strongly related to anticipations concerning the emotions of others. Indeed, our analyses suggest that this is the case (research on this is ongoing and will be reported in another article). Beyond its role in relation to moral emotions and the design of our experiments, the distinction is, however, also helpful in identifying whether different kinds of demands are differentially susceptible to perceptions of overdemandingness and in examining whether such different kinds of demands interact with each other (i.e. do not merely have additive effects).

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

An exemplary scenario used in our studies may help to clarify further this reasoning and the design of the experiments. The beginning of one of the four scenarios – let us call it the *Africa* scenario – used in our first study reads:

> Imagine that you have completed your studies for becoming a civil engineer. You are now looking for a job. You are invited to go to an African country to build an orphanage for children who will suffer from lack of shelter if the orphanage is not built. There are no other applicants for the position and the project will be cancelled if you do not take up the offer. Just before taking the final decision, a good friend who is already an established civil engineer visits you.

The scenario continued to detail one of four possible job-related offers that the friend made to the protagonist. These four offers followed the differentiation between objective and subjective demands. The friend offered either,

- 'a job in his civil engineering office' (both objective and subjective demands are low), or
- 'a job in his civil engineering office. This is the position you have always dreamed of' (here, subjective demands are high, but objective demands are low), or
- 'a very lucrative job in his civil engineering office. This would mean that you could live in financial security' (here, objective demands are high, whereas subjective demands are low), or
- 'a very lucrative job in his civil engineering office. This is the position you have always dreamed of and it would mean that you could live in financial security' (here, both objective and subjective demands are high).

Participants were instructed that taking up the friend's offer would mean having to decline going to Africa. Given the significant level of suffering this would cause to the orphans, it seems plausible that taking up any alternative offer made by the friend would mean rejecting consequentialist requirements. In this study, we were interested in observing at what level of demands in the increasingly burdensome situation participants, putting themselves in the shoes of the agent in the scenario, would claim that consequentialism demands unreasonably much and is therefore no longer in line with what they have decisive reason to do.

The three further scenarios of this study revolved around:

- Being an investigator with the federal police whose only possibility of preventing a serial killer from murdering more people lies in paying an informant (the consequentialist requirement). This, however, would require selling her/his own car. [*Investigator* scenario]

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

- Inheriting a large sum of money that could be spent either to support a charity helping to stop the spread of deadly diseases in the developing world (the consequentialist requirement) or on buying a house. [*Inheritance* scenario]
- Being someone who is faced with either preventing a fire that could threaten people's lives (the consequentialist requirement) or going to a hearing that is critical to her/his own career. [*Fire* scenario]

Again, there were four versions of each of these scenarios following our distinction of objective and subjective demands. Thus, the highly demanding versions of the scenarios included the information that:

- The car held dear memories for the investigator (high subjective demand) and, in addition to the car sale, (s)he would have to take out a loan to pay off the informant that will indebt her/him for the next ten years (high objective demand). [*Investigator* scenario]
- The house was one that the agent who inherited the money had longed for since being a child (high subjective demand) and the property value is such that the agent will be financially secure in her/his old age (high objective demand). [*Inheritance* scenario]
- Missing the hearing would lead to the loss of the protagonist's professional licence (high objective demand), which was unbearable for her/him since the job was the most important aspect of her/his life (high subjective demand). [*Fire* scenario]

Each participant was randomly allocated to one of these scenarios and read all four versions in random order.

After each scenario, participants answered three questions that were central to our analysis:[10] first, they indicated what they would do in such a situation (e.g. in the *Africa* scenario they could 'go to Africa' or 'take up the job'). Second, they answered the question 'Overall, what is the thing to do?' (with the same answer options). Third, they indicated what they believed morality demanded them to do. Most important for our analysis is the second question, because it was designed to determine the overall assessment of reasons in this situation, hence the occurrence of the OD-intuition as it appears in premise (2) in the argument of OD. The analyses below therefore focus mostly on this question.

However, the other two questions are also important for our analysis, albeit for different reasons. The first question concerning expected actual behaviour allows for a more conservative analysis. If answers

[10] There were further questions that are not relevant to the present purpose. These focused, in particular, on anticipated emotional reactions to their decisions (see n. 9 for some details).

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

to the first and the second question differ, this may be because people anticipate internal (e.g. weakness of will) or external (e.g. coercion) processes interfering with the implementation of their intuition. However, in fictitious scenarios like ours, it sheds some doubt on the strength of an intuition if people anticipate not acting in accordance with their overall assessment of reasons. We therefore also conducted more restrictive analyses that required consistent responses to both questions. Thus, in this latter analysis, we only included people who believed they would actually act on their overall assessment of reasons and did not anticipate intervening obstacles. Whenever results for the two types of analyses diverged, we have included a footnote detailing this divergence.

As for the third question, concerning participants' moral assessment of the situation, we have used this question to see if the intuitive judgement of the unreasonableness of consequentialist demands is counteracted by an intuitive judgment of their morality. As explained earlier, this is relevant because OD is an argument against consequentialism as a moral theory. But if people intuit that there is nothing morally objectionable about the unreasonable demands of consequentialism (because what happens is only that non-moral reasons override the moral reasons), then the intuition, i.e. the answer to our second question, cannot be used to support OD. This leads again to a more restrictive analysis. If the answers to our second and third question differ in a particular way, namely, if consequentialist demands are found unreasonable but not demanded by morality, then these judgements of unreasonableness are not taken to count in favour of OD.[11] In our presentation, we will therefore always present the unrestricted analysis first, since this is what determines how widespread the OD-intuition is, and then the restrictive one, since this is what has the final word on the validity of OD itself.

One hundred and forty-three students and staff of a German university (84 female, 59 male) with a mean age of 24 years[12] participated in the study. Almost all were native speakers of German ($n = 134$); all others reported very good German-language skills.

Predictions

As argued above, proponents of OD should argue that there is a widespread intuition that moral agents do not have decisive reason to meet at least some consequentialist requirements. In our study, this

[11] For the reasons given in n. 8, the other way that answers to the questions can come apart – when consequentialist demands are deemed reasonable but not demanded by morality – does not change our analysis.

[12] $SD = 4$ years; range: 18 to 42 years.

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

should be reflected in substantial consensus amongst our participants that the overall assessment of reasons relevant to the decisions posed in the scenarios leads to a conclusion that is inconsistent with consequentialist requirements. This should be particularly true for those scenarios that involve consequentialist requirements posing high demands. In other words, as consequentialist demands increase, we should see an increasing number of participants who find more reason to follow the non-consequentialist than the consequentialist course of action (Hypothesis 1). This tendency of increasing dissent with consequentialism should then ultimately lead to an overwhelming majority of participants rejecting consequentialist requirements at least in the most demanding scenarios (Hypothesis 2).

Results and discussion

To test the predictions derived from OD, we conducted an analysis of variance.[13] Results showed that – across all four scenarios – increases in both objective and subjective demands led to higher levels of dissent with consequentialism.[14] That is to say, participants who were faced with higher demands (be they objective or subjective) saw less reason to comply with consequentialist requirements. The effects of objective demands and subjective demands on the level of dissent with consequentialism were additive rather than multiplicative.[15] Figure 1 details this pattern.

This pattern of results supports Hypothesis 1 in that our participants were indeed sensitive to the varying demand levels. If consequentialism posed higher demands, they saw less reason to follow its requirements. However, the findings are inconsistent with Hypothesis 2. On average across the four scenarios, even the most demanding situations had less than 50% of the participants who indicated that they would have more reason to follow the non-consequentialist as compared to the consequentialist course of action. This calls into doubt premise (2) of the argument of OD in so far as we found no widespread intuition concerning the unreasonable nature of some consequentialist demands.

Of course, it might have been that people reacted differently to each of the four scenarios and that OD-intuition was only apparent in a subset

[13] 2 (Objective Demands: high vs. low) × 2 (Subjective Demands: high vs. low) × 4 (Scenario: Africa vs. investigator vs. inheritance vs. fire) with the first two variables manipulated within subjects and the latter variable manipulated between subjects.

[14] Objective demands: $F(1, 136) = 13.88$, $p < .001$, $\eta^2_p = .093$; subjective demands: $F(1, 136) = 9.95$, $p = .002$, $\eta^2_p = .068$.

[15] There was no interaction between objective and subjective demands, $F(1, 136) = .07$, $p = .78$, $\eta^2_p = .001$. For the more conservative analysis, there was a tendency for such an interaction, $F(1, 78) = 3.37$, $p = .070$, $\eta^2_p = .041$, with a particularly pronounced effect of subjective demands when objective demands were low.

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

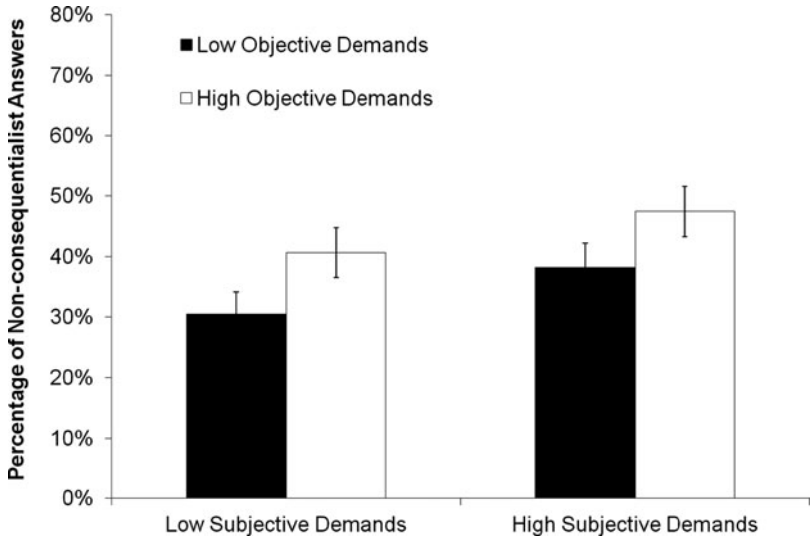


Figure 1. Percentage of non-consequentialist answers to the overall assessment of reasons relevant to the decision across all four scenarios for each level of objective and subjective demands.

of our scenarios. Indeed, the absolute level of responses varied across scenarios (see figure 2).[16] This is mainly relevant to Hypothesis 2: some scenarios were characterized by higher dissent with consequentialism than others. However, even in the most demanding version of the scenario with the strongest dissent (in which one is required to donate a substantial inheritance), only 63% of our participants believed that there was more reason to follow the non-consequentialist than the consequentialist course of action. In relation to Hypothesis 1, the overall patterns across the four differentially demanding versions of each scenario – despite their different absolute level – did not differ significantly.[17]

Let us turn now to the more restrictive analysis. The idea, recall, is that if the answers to questions (2) and (3) differ in a particular way, namely, consequentialist demands are found unreasonable but moral, the intuition concerning unreasonable demandingness does not support OD. In the present case, this was easily quantifiable. What we did here was to check, as the only constellation of answers that clearly supports OD, the percentage of those who found the consequentialist

[16] $F(3, 136) = 2.49$, $p = .016$, $\eta^2_p = .072$.

[17] All two-way interactions with the scenario manipulation: $F(3, 136) < 1.90$, $p > .139$, $\eta^2_p < .040$. Three-way interaction: $F(3, 136) = .43$, $p = .730$, $\eta^2_p = .009$.

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

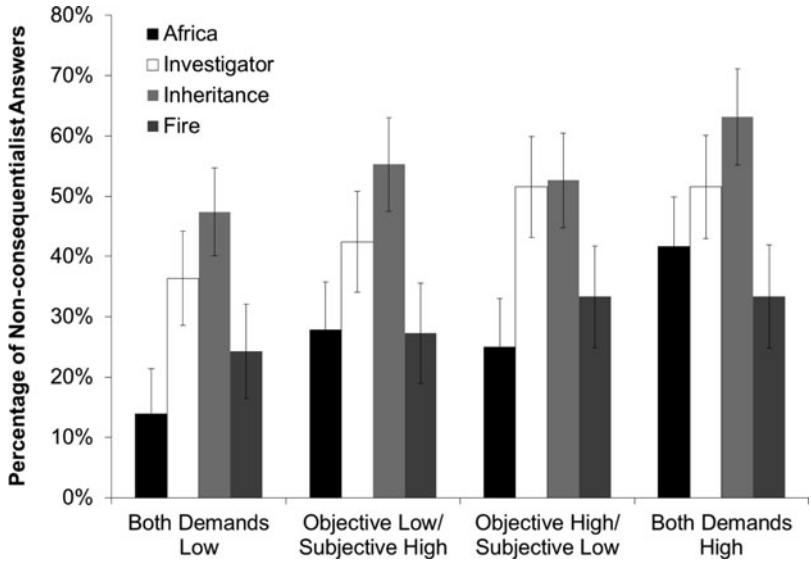


Figure 2. Percentage of non-consequentialist answers to the overall assessment of reasons relevant to the decision for each of the four scenarios for each level of objective and subjective demands.

requirement not only unreasonably demanding, but also not demanded by morality.

This restrictive analysis again revealed support for Hypothesis 1: increases in objective and (marginally significant) subjective demands led to higher levels of dissent with consequentialism.[18] More importantly, however, when we only took into account those patterns of replies that are fully consistent with OD, the absolute level of dissent with consequentialism dropped substantially from the unrestricted analysis. On average across all scenarios, even in the most demanding condition less than 28% of our respondents gave replies that fully conformed to what proponents of OD would have to expect (see figure 3). That is to say, even of those (in the most demanding condition) who thought that they did not have decisive reason to follow consequentialist demands (the most right-hand bar in figure 1), nearly half thought that the course of action demanded by consequentialism was what morality required them to do. A possible – and in our eyes probable – explanation of this finding is that, while subscribing

[18] Objective demands: $F(1, 136) = 10.91$, $p = .001$, $\eta^2_p = .074$; subjective demands: $F(1, 136) = 3.89$, $p = .051$, $\eta^2_p = .028$.

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

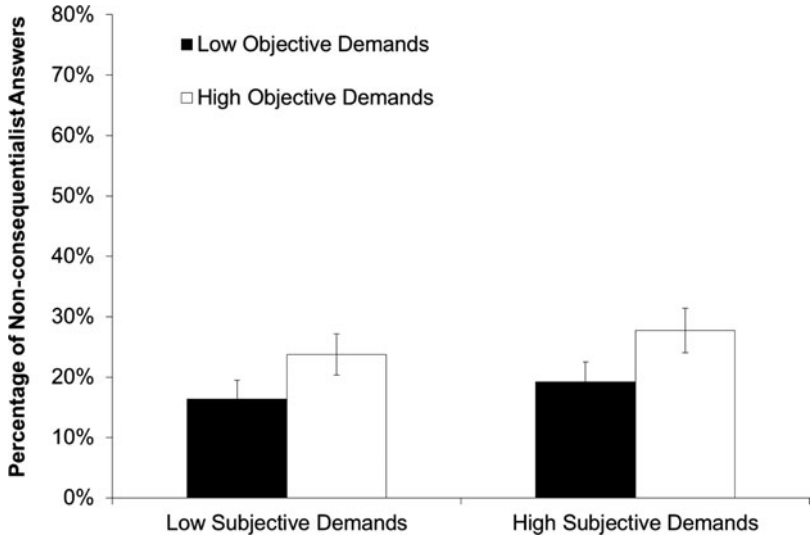


Figure 3. Percentage of non-consequentialist answers to the overall assessment of reasons relevant to the decision across all four scenarios for each level of objective and subjective demands. This restrictive analysis is limited to those who reported that (a) they had no decisive reason to follow the consequentialist demands and (b) the consequentialist course of action was not demanded by morality.

to a consequentialist morality, they do not consider consequentialist reasons to override non-moral reasons for action.

And again, to make sure that we did not miss important differences between the individual scenarios, we examined each scenario separately. In no condition of any one of the scenarios did answers consistent with OD (i.e. dissenting with consequentialism) as defined in this more restrictive analysis exceed 45% of the respondents. Also, the overall pattern of results shown in figure 3 did not differ significantly between the four scenarios.[19]

IV. EXPERIMENT 2: AN EXPERIMENTAL GAME

Critics of Experiment 1 may find intuitions in response to fictitious scenarios lacking (see e.g. a recent exchange of letters between David Carmel and Shaun Nichols in *Science*, 10 June 2011).[20] The goal of

[19] All two-way interactions with the scenario manipulation: $F(1, 136) < 1.08$, $p > .362$, $\eta^2_p < .023$. Three-way interaction: $F(3, 136) = 2.40$, $p = .070$, $\eta^2_p = .050$.

[20] David Carmel, 'Experimental Philosophy: Surveys Alone Won't Fly', *Science* 332 (2011), p. 1262; Shaun Nichols, 'Response', *Science* 332 (2011), pp. 1262–3.

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

the second study, therefore, was to create an actual decision situation with real-life monetary consequences. This was done to investigate the OD-intuition while avoiding the possible adverse effects of the 'as if'-character of imaginary scenarios.[21]

Procedure

We adapted a paradigm previously used in behavioural economics.[22] Participants divided a certain sum (€100) between themselves and the vaccination programme of the United Nations Children's Fund (UNICEF) benefiting children in developing countries. Participants knew that for one randomly determined participant, this choice would be implemented in real. Prior to the choice, we varied both objective and subjective demands of the consequentialist option, which, we assumed, was to donate as opposed to keep the money. To induce subjective demands, half of the participants first brainstormed on what they personally could do with €100, thereby rendering this money more affectively relevant (high subjective demands), whereas the other half reflected on what they could do with a piece of paper (low subjective demands). Specifically, in the high subjective demand condition participants were asked to imagine that they would win €100 in this study. They were then instructed to 'consider three possible ways to spend this money' and to write down these three possibilities.

The aim of the second task was to induce objective demands. Prior research in economics has shown that participants distributing money are sensitive to where this money comes from. In particular, in contrast to so-called 'windfall endowments', which are simply given to participants by the experimenter, endowments that have to be earned are less readily shared.[23] To avoid the possible problem that those higher in ability or motivation may earn higher endowments, we created a task that allowed us to make all participants believe that they managed to earn the full amount (i.e. the chance of distributing €100). Participants saw five pictures with a varying number of black dots on a white background and were asked to estimate the average number of dots. All of them were later told that they were sufficiently close to the correct number (which was 150) to participate in the study. Those in the low objective demand condition were simply informed that they

[21] Brian Parkinson and Antony S. R. Manstead, 'Making Sense of Emotion in Stories and Social Life', *Cognition & Emotion* 7 (1993), pp. 295–323.

[22] Fernando Aguiar, Pablo Brañas-Garza and Luis M. Miller, 'Moral Distance in Dictator Games', *Judgment and Decision Making* 3 (2008), pp. 344–54; Catherine C. Eckel and Philip J. Grossman, 'Altruism in Anonymous Dictator Games', *Games and Economic Behavior* 16 (1996), pp. 181–91.

[23] Todd L. Cherry and Jason F. Shogren, 'Self-interest, Sympathy and the Origin of Endowments', *Economics Letters* 101 (2008), pp. 69–72.

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

received the opportunity to distribute the money, without the need to do anything to gain this opportunity (i.e. the money came as a windfall endowment).

After completing these tasks, participants received some information on the UNICEF vaccination programme and then answered three questions that were central to our analysis.[24] First, they indicated how they would actually divide the money between themselves and the charity (under the supposition that they would be randomly selected to be the person for whom the decision would be implemented). They could assign any share of the €100 to either themselves or UNICEF. This question was designed to assess their actual behavioural responses. Second, they replied to the question 'Overall, what is the division to make?' referring to their overall assessment of reasons for the division. Third, they indicated which distribution they regarded to be morally required of them. As in Experiment 1, our main focus was on the second question because it most clearly speaks to premise (2) of the argument of OD. At the same time, answers to the first question help to establish the credibility of these results and, in particular, their relevance to actual behaviour, while answers to the third question provide for an importantly more restrictive analysis. Take someone who holds that what there is decisive reason to donate is not in line with consequentialism, while affirming that the consequentialist division would be the morally supported division to make. This person does not support OD.

We recruited 368 participants online through a university participant panel. Eleven either did not indicate their age or were underage and were therefore excluded. The remaining 357 participants were 18 to 57 years of age with a mean of 23.5 years.[25] Most participants were women ($n = 215$), all but three possessed at least a university entrance diploma, and the overwhelming majority were native speakers of German (93.9% of all participants).

Predictions

As in our previous study, we examined the predictions derived from the relevant literature on OD that (a) increasing (objective and subjective) demands would lead to higher levels of dissent with the consequentialist course of action (i.e. lower donations would be perceived as reasonable; Hypothesis 1) and that (b) the absolute level of dissent with the consequentialist course of action would be very high at least in the most demanding condition (i.e. the objection

[24] As in Experiment 1, further questions mainly concerned anticipated emotional reactions to their decisions and are not relevant to the present purpose.

[25] $SD = 3.9$ years.

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

to consequentialism should be based on a *widespread* intuition; Hypothesis 2).

Results and discussion

To test our hypotheses, we conducted an analysis of variance on each of the three questions we posed.[26] Results for the central second question concerning participants' overall assessment of reasons revealed that both increased subjective and objective demands led participants to keep a higher percentage of the endowment for themselves; that is, to reject consequentialist requirements concerning the support of the vaccination programme.[27] If participants had had the opportunity to elaborate on what they would do with the money if they won it (i.e. if they had brainstormed about the possible use of €100) and if they had the impression that they had 'earned' the possibility to keep the €100 for themselves (i.e. if they thought they had gained the money by correctly guessing the number of points presented to them), they, on average, saw less reason to donate. Consistent with the previous study, this pattern supports Hypothesis 1: people are sensitive to the demands made by consequentialism in their assessment of reasons for action. However, the detailed pattern of means shown in figure 4 again sheds doubt on the viability of Hypothesis 2. Even in the most demanding situation, participants saw reason to keep, on average, less than 60% of the money for themselves (i.e. they believed they had reason to donate more than 40%).

However, in contrast to Study 1, this experiment offered some opportunity for weakness of will to reveal itself. In response to our first question – the actual division of money under the supposition that participants would be randomly selected – participants kept a significantly larger share of the endowment – on average approximately 8 percentage points more – to themselves, than what they indicated was reasonable to keep in response to question 2.[28] In the most demanding condition, this resulted in participants claiming 67.5% of the endowment for themselves (see figure 5 for the detailed pattern of means).

Finally, the more restrictive analysis is not as easily quantifiable as in the previous study. What we were interested in here, recall, were constellations of a particular kind in participants' answers to our second and third question: when consequentialist demands were found

[26] 2 (Objective Demands: high vs. low) × 2 (Subjective Demands: high vs. low) between subjects.

[27] Objective demands: $F(1, 353) = 5.33$, $p = .022$, $\eta^2_p = .015$; subjective demands: $F(1, 353) = 6.49$, $p = .011$, $\eta^2_p = .018$.

[28] $F(1, 353) = 30.58$, $p < .001$, $\eta^2_p = .080$

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

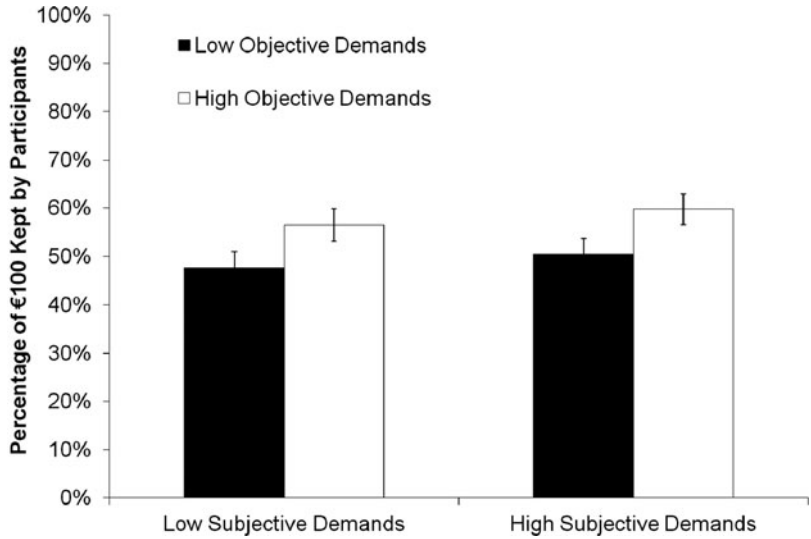


Figure 4. Average percentage of €100 that participants thought would be reasonable to keep for themselves (the remainder would be donated to UNICEF) for each level of objective and subjective demands.

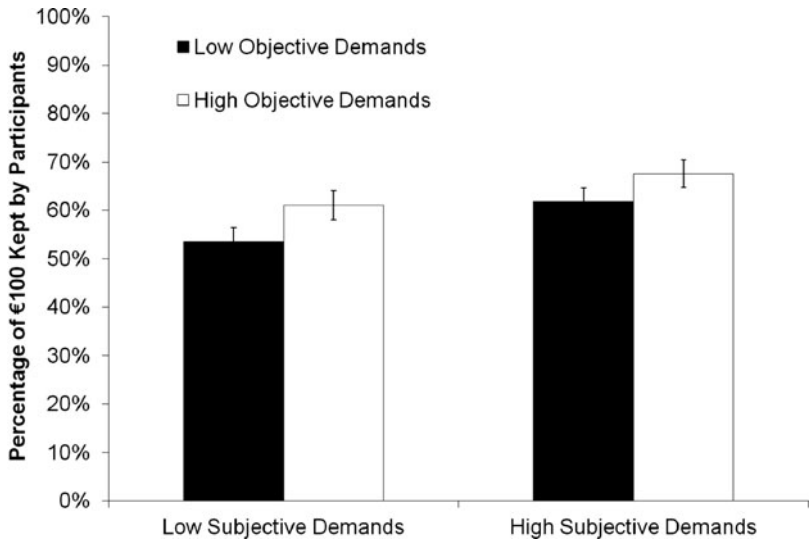


Figure 5. Average percentage of €100 that participants intended to keep for themselves (the remainder would be donated to UNICEF) for each level of objective and subjective demands.

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

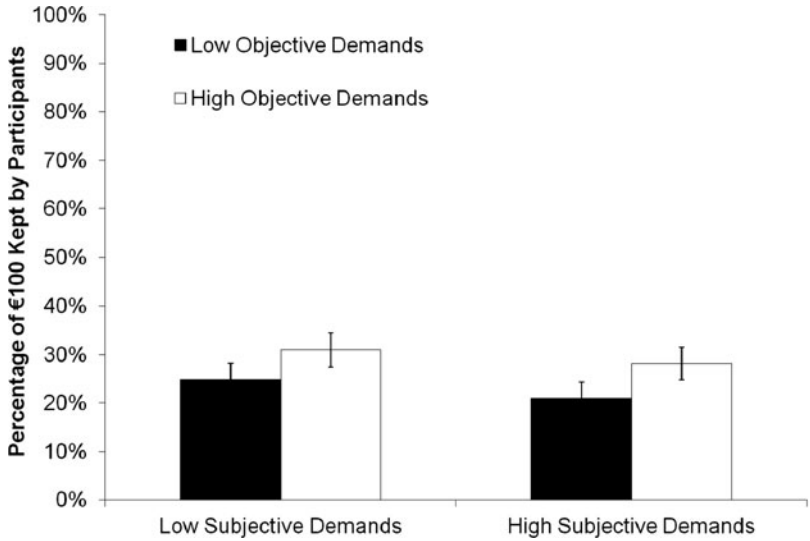


Figure 6. Average percentage of €100 that participants believed morality allowed them to keep for themselves (the remainder would be donated to UNICEF) for each level of objective and subjective demands.

unreasonable as well as not demanded by morality. We can provide the following data. Results showed that participants indeed believed that high levels of donations were morally required of them. As figure 6 demonstrates, they indicated that – assessing moral reasons only – they should keep only about one quarter of the endowment to themselves (26.1%). This was significantly less than both what they believed they, overall, had reason to keep and what they indicated that they would actually want to keep.[29]

The restrictive analysis suggests that what most participants experienced was a clash of moral and non-moral reasons in which the former were overridden. That is, most of our participants continued to endorse the morality of consequentialist requirements, contrary to their perceived unreasonableness. And this, as argued earlier, counts against the viability of OD, despite the more favourable – although, as pointed out above, not too favourable – judgements of reasonableness.

In sum, the second study corroborated the findings of Study 1 supporting Hypothesis 1. Again, there was evidence that with increasing consequentialist demands people see less reason to follow its requirements. However, strong doubts remained, even setting the more restrictive analysis aside, whether there exists a sufficiently

[29] Fs(1, 353) > 267.62, ps < .001, η^2_p > .431.

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

widespread intuition that can seriously challenge consequentialism on the basis of OD. The viability of Hypothesis 2 could again be doubted.

V. GENERAL DISCUSSION

Taken together, we suggest that our empirical studies illustrate two key points. First, in relation to Hypothesis 1, the more severe the requirements made by consequentialism, the less reason people think they have to follow these requirements. This is also reflected in their actual behaviour (Experiment 2). Second, despite this sensitivity to consequentialist demands, the absolute level of dissent with consequentialism is not such that one would talk of a widespread OD-intuition. We say this in the absence of normative criteria for what would constitute a widespread intuition. However, as mentioned earlier, other empirical research in moral philosophy has found the agreement of around 90% of respondents, whereas in our first study even the most severe version of the scenario (with the highest level of dissent with consequentialism) only found consensus among 63% of participants and in the second study, the donations were not more than 60%. And this only in the analyses solely focusing on judgements of reasonableness: when placing more severe requirements on respondents' patterns of answers (namely, that they considered the consequentialist course of action not only unreasonable but also not demanded by morality) the percentages are even lower, throwing further doubt on the validity of premise (2) of the argument of OD.

Hence, the implications for the viability of OD are mixed. On the one hand, OD seems to capture an important aspect of folk morality, namely, that, on average, moral demands are seen as less reasonable when they are high rather than low. On the other hand, the studies raise doubts about whether the OD-intuition is as uniformly shared as one might suspect on the basis of the discussions surrounding OD, which do not question this side of the objection. The article therefore contributes to these discussions by showing that while OD may reflect an important aspect of moral reasoning, those sympathizing with it will need to specify further the conditions under which the supposed intuition indeed is widely accepted and argue for the philosophical relevance of the situations in which these conditions are met. We will return to this point at the end of this discussion.

Of course, before accepting these conclusions, one may raise a number of critical points about this research: First, there is the issue of the status of intuitions in moral philosophy. How can intuitions constrain normative moral theories? Second, even if intuitions are relevant in moral philosophizing, how can we know about them and are the particular studies presented here suited to revealing them? Third,

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

given the often-observed variation in intuitions (more precisely, the variation in empirical observations that claim to correspond with intuitions), there is an issue of whether we need to specify whose intuitions should be assessed. Is a particular qualification needed to have the 'right' moral intuitions? Fourth, and finally, one may have more specific methodological qualms with the testing of our hypotheses. Even supposing that the prior three issues were resolved, do our studies lead to the conclusions that we draw from them?

Concerning the first three questions, this is not the place to discuss them in detail.[30] Intuitions, as noted earlier, can function as devices of persuasion and as evidence: it is in these two ways that they can constrain moral theories. The first such role we take to be uncontroversial (given the weak reading we have given of it) and we find it reasonable to assume that intuitions can serve as evidence. This view is certainly not undisputed, but this is not the place to argue for its correctness. [31] We also find it reasonable to assume that one can construct reliable methods of detecting intuitions. In fact, we think that we have gone to considerable lengths to investigate intuitions in a way that – given the state of knowledge about intuitions and the currently available procedures – offers us a good chance that we do at least as well as others in identifying intuitions. Similar considerations apply to the third difficulty mentioned. Since settling the particular debate goes beyond the scope of this article, we have acted on the assumption that there is at least as much argument speaking in favour of investigating folk intuitions as any other intuitions of a defined group of people.[32]

However, we do have to deal with the fourth worry at greater length since it concerns the particularities of our argument. Let us begin with an issue that connects directly to the matter discussed above. It could be argued that to test the existence of folk intuitions, the

[30] We say more about these questions in Martin Bruder and Attila Tanyi, 'How to Gauge Moral Intuitions? Prospects for a New Methodology', *Experimental Ethics*, ed. Christoph Lütge, Hannes Rusch and Matthias Uhl (forthcoming).

[31] In moral philosophy recent works include, on the sceptical side, Peter Singer, 'Ethics and Intuitions', *The Journal of Ethics* 9 (2005), pp. 331–52; on the anti-sceptical side, Folke Tersman, 'The Reliability of Moral Intuitions: A Challenge From Neuroscience', *Australasian Journal of Philosophy* 86 (2008), pp. 389–405; Matthew Liao, 'A Defence of Intuitions', *Philosophical Studies* 140 (2008), pp. 247–62. See also Matthew Tedesco ('Intuitions and the Demands of Consequentialism', *Utilitas* 23 (2011), pp. 94–104) who pursues this line further, in our particular context, by bringing in neuroscientific evidence.

[32] See Jonathan M. Weinberg, Chad Gonnerman, Cameron Buckner, and Joshua Alexander, 'Are Philosophers Expert Intuiters?', *Philosophical Psychology* 23 (2010), pp. 331–55; Eric Schwitzgebel and Fiery Cushman, 'Expertise in Moral Reasoning? Order Effects on Moral Judgment in Professional Philosophers and Non-Philosophers', *Mind & Language* 27 (2012), pp. 135–53. See also M. B. E. Smith, 'Does Humanity Share a Common Moral Faculty?', *The Journal of Moral Philosophy* 7 (2010), pp. 37–53.

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

participant sample should be representative of the general population and that ours is not. Indeed, this is a much-discussed problem across the empirical social sciences which, we believe, warrants more careful discussion also within experimental philosophy. And indeed, there is good reason to be cautious when drawing general conclusions from non-representative samples.[33] However, 'being cautious' does not translate into 'simply discarding' as soon as samples are non-representative. This would translate into the absurd claim that we cannot learn anything general about human nature (from the biological to the psychological) as long as we do not draw a *globally representative* sample. Thus, in short, this is not a black or white issue. Our task is to argue that our sample was appropriate for drawing the conclusions that we propose. We do believe that it was and have tried to ensure this by recruiting a sample that included students from all disciplines (in contrast to much research only running classroom experiments with students from one discipline). We did so because we were worried that any kind of moral education included in the curriculum could bias our results. Also, we want to argue that this work constitutes a significant improvement of the empirical basis of the overdemandingness discussion given that – to date – there were no empirical observations relevant to this controversy outside expert philosophers' introspection. Finally, we have tried to be mindful in our conclusions of the danger of generalizing too widely.

The challenge, let us add, is not only ours. Anyone worried about the representativeness of our samples should specify interesting predictions as to which groups will diverge from our findings and, ideally, why. If these predictions are promising and exciting enough to inspire others to invest time and money in new research that will qualify our findings, this is, of course, a positive outcome. Although we submit that it is likely that different groups within one culture and populations from different cultures may vary somewhat with respect to their consent or dissent with consequentialism, we, at this stage, do not see very strong reasons to think that the resulting patterns would lead to fundamentally different conclusions. The content of the scenarios would, of course, have to be adapted (to ensure equal demandingness in the different cultures); however, the pattern of results, we propose, should be largely unchanged. But again, that is an empirical claim that advocates of OD may contest and that would need to be the object of future research.

Another possible methodological criticism concerns our first experiment. One might ask: could it be that the scenarios are

[33] Joseph Henrich, Steven J. Heine and Ara Norenzayan, 'The Weirdest People in the World?', *Behavioral and Brain Sciences* 33 (2010), pp. 61–83.

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

unrealistic and so construed that people more easily give 'desirable' answers than in complex real-life settings where the influence of one's own actions on the world are less clear and all-or-nothing affairs? In response, let us first point out that our second study was specifically designed to address this potential shortcoming of the first experiment. Further, we also doubt that our scenarios are particularly unrealistic; they are certainly not more unrealistic than most similar philosophical thought experiments available in the literature. Thus, while our scenarios certainly do not perfectly reflect real-life situations, the question is how much better they could really be. In short, we hold that experimental philosophy must use certain simplifications and perhaps introduce some artificiality simply because there are many situations to which participants should not or even must not be exposed. Again, we have taken the potential artificiality of our scenarios into account in drawing our conclusions.

The last worry we would like to consider can be stated as follows. Since Hypothesis 1 is confirmed, that is, we established that increasing demands led to increasing dissent with consequentialism, we probably would also observe a confirmation of Hypothesis 2 if we considered sufficiently extreme situations. Hence, one may argue, these results do not really question the truth of premise (2) of the argument of OD. Our response to this objection is threefold. First, what really matters for judging the truth of the premise is whether the situations we used are such that they invite the intuition underlying OD as it appears in the premise. In the most demanding versions of our scenarios, we certainly claim this to be the case, and the objection does nothing to question this assumption. Second, it is a question what use alternative studies utilizing more extreme situations would have for understanding the viability of OD. Responses to such situations arguably make a difference when what is at stake are conceptual or metaphysical truths, but the strength of OD lies exactly in its assumed embeddedness in everyday situations. This of course also connects to the previous concern: were we to test participants' reactions in such extreme situations, the charge of unrealistic experimentation would certainly stick.

Our third response is less speculative. In a follow-up study we affirmed that results are largely parallel in cases that have been introduced in the philosophical literature to illustrate the gist of the overdemandingness objection.[34] These cases are interesting not only because they contain more extreme situations, but also because they have among them scenarios that are more ordinary, yet giving rise to

[34] For the examples see Bykvist, *Utilitarianism*, pp. 98–9, and Mulgan, *Understanding Utilitarianism*, pp. 95–6.

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

extreme demands, hence immune to our second response above.[35] These are so-called iterative demands, i.e. they are small but continuously repeated requirements that may easily add up to altogether extreme demands. The following six cases have been used in our experiment:

S1. On your desk is an envelope addressed to a reputable charity seeking donations to save the lives of victims of a natural disaster. In the envelope is a cheque totalling the amount of all the money you possess.

S2. You see an advertisement on television from a charitable organization calling for volunteers to spend the next thirty years working with destitute people in a very poor country.

S3. You must decide whether to become a teacher or a merchant banker. Although you find it very unfulfilling, you have a natural aptitude for banking. You calculate that if you become a banker and donate all your earnings to charity, this will produce more happiness for others than if you do anything else with your life.

S4. After you have saved a child from a burning building you are told that there is another child still left in the building. You could go back into the building and save the second child, but you know that this will cause you third-degree burns. However, if you do not go back, the child will die.

S5. You are wondering whether to spend a pound on chocolate for yourself or to give it to a certain charity. You know that this charity is unusually effective and that even a small contribution can help them save a child from some crippling and painful illness. If you decide to give every spare pound you have to the charity, a lot of children can be saved.

S6. At the end of a long, tiring day you sit down to relax and watch a TV show. However, you know that on your desk is a letter from a reputable charity asking for donations and volunteer contributions. Instead of spending your time relaxing, you could help the cause of the charity and join in voluntary work. But if you decide to work for the charity whenever you have time on your hands, you will never have some time for yourself again.

Participants were 265 students and staff of a German university (154 female, 109 male, 2 unidentified) with a mean age of 24 years.[36] Two-hundred and fifty-four had German as their native language, the remaining 11 self-reported *good* or *very good* German-language skills.

[35] An anonymous reviewer for this journal has also pointed out to us that iterative demands could be a problem for our article. The best discussion of iterative demands in the context of OD is in Cullity, *The Moral Demands of Affluence*.

[36] $SD = 4$ years. Range from 18 to 51.

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

Table 1. Percentage of respondents identifying the consequentialist/non-consequentialist course of action as the one that they would choose (Q1), the one that, overall, was the thing to do (Q2), and the one that was demanded by morality (Q3).

	S1	S2	S3	S4	S5	S6
Q1	4/96	6/94	28/72	70/30	59/41	10/90
Q2	21/79	42/58	36/64	**78/22**	**83/17**	38/62
Q3	41/59	60/40	44/56	92/8	90/10	66/34
Q2 + Q3	44/56	***64/36***	***52/48***	***93/7***	***92/8***	***67/33***

Note: Q2 + Q3 denotes the restricted analysis in which only those claiming that the non-consequentialist course of action was both not the thing to do and not demanded by morality counted as supporting OD (and thus as non-consequentialist).

In this study, no differentially demanding versions of each scenario existed. Thus, all participants read all scenarios in random order. After each scenario, they were asked, as in our first scenario study, the same three questions concerning what they would do (Q1), what, overall, was the thing to do (Q2), and what morality demanded them to do (Q3). Just as in that study, we also carried out a restrictive analysis, counting only those patterns of answers as supporting OD that identified the consequentialist demand as both unreasonable and not demanded by morality (Q2 + Q3 both consistent with OD). Table 1 summarizes the results of this study.

As these data show, on the unrestricted analysis (see numbers in bold), that is, taking only answers to question (2) into account, two scenarios – S4 and S5 – clearly do not support OD, and, crucially, one of them is iterative in nature (S5). All other scenarios apart from S1 also do not reveal dissent with consequentialism at a level higher than that of the most demanding scenarios in Study 1, thus offering some reassurance that our scenarios were not less demanding than those discussed in the relevant literature. Therefore, Scenario 1 is the only one that appears more extreme than the ones we used. On close inspection, however, it can be questioned whether indeed giving *all* the money one has to charity is a demand that consequentialism actually would make (in particular, if this endangers one's own survival). Furthermore, on the more restrictive analysis (see italicized numbers in bold), parallel to Study 1, less than half of all responses to questions 2 and 3 are fully consistent with OD (again with the exception of Scenario 1).

What follows from this, in the final analysis, is that there can be extreme situations in which the OD-intuition shows itself in a stronger

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press

light. However, it is clear from our investigations that these situations must be rather unusual. And this fact, coupled with our finding that in the most highly demanding situations, including textbook examples of unreasonable consequentialist demands, the OD-intuition does not occur widely across our sample population, puts the burden of proof on advocates of OD. To maintain a strong challenge, proponents of OD would need to put forth a strong argument why the focus should be on these extreme cases rather than other more ordinary or similarly extreme cases in which the OD-intuition, *contrary to what they claim*, is far from uniformly present. Otherwise, OD loses its force and consequentialists can turn to other, potentially more pressing challenges to their theory.[37]

VI. SUMMARY AND CONCLUDING REMARKS

In this article, we have discussed the Overdemandingness Objection (OD) to consequentialism. In contrast to the existing extensive literature on the subject, we have focused on what we take to be the empirical foundation of OD: that people have an intuition that consequentialist requirements are in certain situations unreasonably demanding. We have put this claim under empirical scrutiny by designing and carrying out two experiments (as well as a follow-up study). We have found that the claim is far from obviously true. In our experiments, although people were sensitive to consequentialist requirements and, on average, found more extreme demands less reasonable, the level of dissent with consequentialism fell short of qualifying as a widely shared intuition, even when demands were the highest. At the same time, as we admitted in our general discussion, there are several sceptical questions that need to be answered to make this research and experimental philosophy in general more convincing than it presently is. Primary among them, we believe, is the task to work out a reliable methodology for detecting intuitions. This is the task, however, for another article.[38]

martin.bruder@uni-konstanz.de
attila.tanyi@liverpool.ac.uk

[37] One recalls here Rawls's well-known remarks about the usefulness of counterexamples, in his *A Theory of Justice* (Cambridge, Mass., 1971), p. 52.

[38] See our 'How to Gauge Moral Intuitions?', as well as Martin Bruder and Attila Tanyi, 'Consequentialism and its Demands: A Representative Study', forthcoming in a special issue of the *Journal of Value Inquiry*, ed. Sabine Roeser and Joel Rickard. See also John Bengson, 'Experimental Attacks on Intuitions and Answers', *Philosophy and Phenomenological Research* 86 (2013), pp. 495–532.

https://doi.org/10.1017/S0953820814000065 Published online by Cambridge University Press