Personality and Social
Psychology Bulletin
2014, Vol. 40(7) 923–930

DOI: 10.1177/0146167214530436
pspb.sagepub.com
SAGE

Efficient Kill–Save Ratios Ease Up the Cognitive Demands on Counterintuitive Moral Utilitarianism

Bastien Trémolière[1] and Jean-François Bonnefon[1,2]

Abstract

The dual-process model of moral judgment postulates that utilitarian responses to moral dilemmas (e.g., accepting to kill one to save five) are demanding of cognitive resources. Here we show that utilitarian responses can become effortless, even when they involve to kill someone, as long as the kill–save ratio is efficient (e.g., 1 is killed to save 500). In Experiment 1, participants responded to moral dilemmas featuring different kill–save ratios under high or low cognitive load. In Experiments 2 and 3, participants responded at their own pace or under time pressure. Efficient kill–save ratios promoted utilitarian responding and neutered the effect of load or time pressure. We discuss whether this effect is more easily explained by a parallel-activation model or by a default-interventionist model.

Keywords

moral cognition, utilitarianism, time pressure, cognitive load

Received December 13, 2013; revision accepted March 8, 2014

Is it morally permissible to kill one to save many? Different people react differently to this dilemma. Those who say yes are considered to give the utilitarian response, and the core claim of the dual-process approach to moral judgment (Greene, Nystrom, Engell, Darley, & Cohen, 2004; Greene, Sommerville, Nystrom, Darley, & Cohen, 2001) is that giving this response requires controlled, effortful mental processes. If the utilitarian response is controlled rather than automatic, manipulations that temporarily decrease cognitive resources (e.g., concurrent load or time pressure) should decrease its frequency. Although Greene, Morelli, Lowenberg, Nystrom, and Cohen (2008) failed to observe such an effect of cognitive load, Trémolière, De Neys, and Bonnefon (2012) did obtain the effect by using an extra-strong manipulation of load. In parallel, Suter and Hertwig (2011) observed that utilitarian responses were less frequent under time pressure, at least for a subset of their dilemmas, and Greene et al. (2008) observed that utilitarian responses took longer when participants were under cognitive load.

The claim that utilitarian responses to moral dilemma require controlled, effortful cognitive processing is a pillar of current research on moral thinking. Recently though, Kahane et al. (2012) argued that utilitarian responses to moral conflicts did not necessarily imply cognitive effort. They noted that the evidence for this claim had always been obtained with dilemmas for which the utilitarian response was highly counterintuitive, as it implied to kill or inflict severe harm. They then reported behavioral and neuroimaging data suggesting that giving the utilitarian response became intuitive when it implied a milder transgression, such as lying to prevent emotional or physical pain. The conclusion of the article was that the intuitive versus counterintuitive dichotomy was perhaps more relevant to moral reasoning than the utilitarian versus non-utilitarian dichotomy.

Our goal in this article is to nuance this conclusion, and to demonstrate that utilitarian responses can be delivered effortlessly, even when they imply a strong transgression such as killing. The basic idea underlying our experiments is that the moral acceptability to kill one to save *N* lives will increase with *N*, to the point that utilitarianism will become an undemanding response for large values of *N*. This is what we call the kill–save ratio effect: Efficient kill–save ratios (e.g., kill 1 to save 500) would encourage utilitarian responses, while making them so undemanding that they would survive manipulations of cognitive load or time pressure.

The novel idea here is that highly efficient ratios might generate *effortless* utilitarian judgments—not that they might encourage utilitarian judgments. Many studies demonstrated that people are sensitive to the ratio of positive versus harmful consequences of an action, when assessing the acceptability of this action. This is usually done by asking people how many saved lives (or trees, or species) would make it acceptable to sacrifice one life (or tree, or species). These studies are

[1]Université de Toulouse, France
[2]Centre National de la Recherche Scientifique, France

Corresponding Author:
Bastien Trémolière, Cognition, Langues, Langage, Ergonomie, Université de Toulouse, Le Mirail, Maison de la Recherche, 5, allées Antonio-Machado, 31058 Toulouse Cedex 9, France.
Email: bastien.tremoliere@univ-tlse2.fr

typically investigating the conditions under which harmful actions are considered as morally acceptable (e.g., Baron & Leshner, 2000; Bartels & Medin, 2007; Nichols & Mallon, 2006; Ritov & Baron, 1999). We expect that better kill–save ratios will increase the moral acceptability of sacrificing one to save many, and thus promote utilitarian thinking, but we are chiefly interested in the *cognitive cost* of utilitarian thinking as a function of kill–save ratio.

Interestingly, the 12 scenarios used in Greene et al. (2008) featured vastly different kill–save ratios, and we were able to reanalyze these data. The three scenarios with the most inefficient kill–save ratios (1:1 to 1:5) elicited about 44% utilitarian responses, whereas the three scenarios with the most efficient kill–save ratios (1:100 to 1:1,000,000) elicited about 84% utilitarian responses. Across the 12 scenarios, the standardized effect size of cognitive load (h) showed a modest correlation (r = .24) with the kill–save ratio,[1] suggesting that, in line with our expectation, the effect of load (and accordingly, the cognitive cost of the utilitarian response) was smaller for the most efficient ratios. It is interesting in this respect to observe that the effect of load was in opposite directions for the scenarios featuring the three most inefficient ratios, and the three scenarios featuring the most efficient ratios (−0.04 and +0.12, respectively).

This re-analysis is promising, but it does not afford strong conclusions. To conduct an appropriate test of our hypothesis, we need to use a more controlled manipulation of the kill–save ratio, as well as stronger manipulations of cognitive load. Our first experiment will feature the modified dot matrix task introduced in Trémolière et al. (2012), and our second and third experiments will feature the time pressure manipulation of Suter and Hertwig (2011).

Our basic prediction is that efficient kill–save ratios will encourage utilitarian responses, and make them undemanding of cognitive resources. That is, concurrent load and time pressure should decrease utilitarian responding for inefficient kill–save ratio, but should not have any effect on utilitarian responding for efficient kill–save ratios. If we are correct, we will have to reconcile both the dual-process model of moral judgment and its reformulation by Kahane et al. (2012) to the fact that strong transgressions such as killing can be effortlessly accepted as moral.

Disclosure Statement

For all experiments, we report all measures, conditions, and data exclusions. The minimal sample size was set at 120 for all experiments. This target was typically exceeded in just a couple of days of online data collection, in which case all participants were kept in the sample.

Experiment 1

Method

The 213 participants (142 women, M age = 30.2, SD = 12.7) were recruited through a French online data collection platform. Participants were randomly assigned to one condition of a 2 × 2 × 2 mixed design, manipulating cognitive load (between-participant), kill–save ratio (within-participant), and scenario (within-participant).

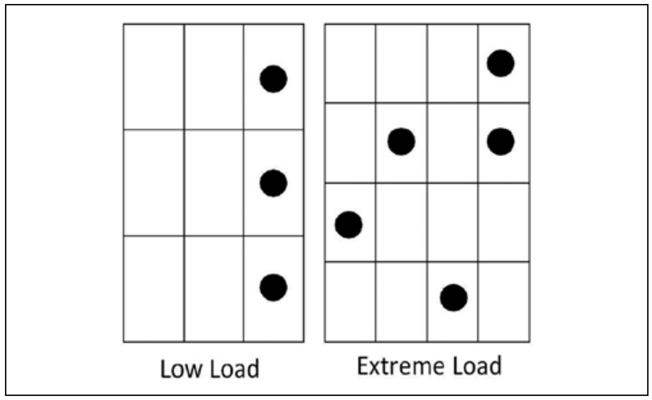

Figure 1. Example of dot matrices used in the two load conditions.

Each participant saw two moral dilemmas (Captive Soldier and Crying Baby, see Appendix A), which were interspeded by filler moral problems that did not feature a dilemma. The order in which these dilemmas appeared, as well as which one featured the 1:5 ratio and which one the 1:500 ratio, was randomly determined for each participant. For each dilemma, participants had to indicate whether they found it morally acceptable or inacceptable to perform the action.

To manipulate cognitive load, we used the dot memory task (Bethell-Fox & Shepard, 1988; De Neys, 2006; De Neys & Verschueren, 2006; Miyake, Friedman, Rettinger, Shah, & Hegarty, 2001), a classic spatial storage task already used by Trémolière et al. (2012) in the context of moral judgment. Before each scenario, participants were presented a matrix in which some cells were filled with dots.

Participants were instructed to memorize the location of the dots. After that, participants read and respond to the dilemma. Once done with the scenario, they had to reproduce the configuration of the dots in an empty matrix.

Participants in the low load condition saw very easy 3 × 3 matrices, similar to that presented in the left panel of Figure 1. These 3 × 3 matrices were presented for 850 ms. Participants in the extreme load conditions saw extremely difficult 4 × 4 matrices, similar to that in the right panel of Figure 1. To make this task feasible, these matrices were shown for 2 s. We then recorded the number of correctly located dots for each participant and each matrix. It is well established that the memorization of the low load pattern does not tap much cognitive resources (De Neys, 2006), while difficult matrices mobilize cognitive resources to a great extent (Bethell-Fox & Shepard, 1988; Miyake et al., 2001).

Results

Participants showed adequate performance in the Dot Memory Task. The percentage of correctly reproduced cells was 95% in the low load condition, and 81% in the extreme

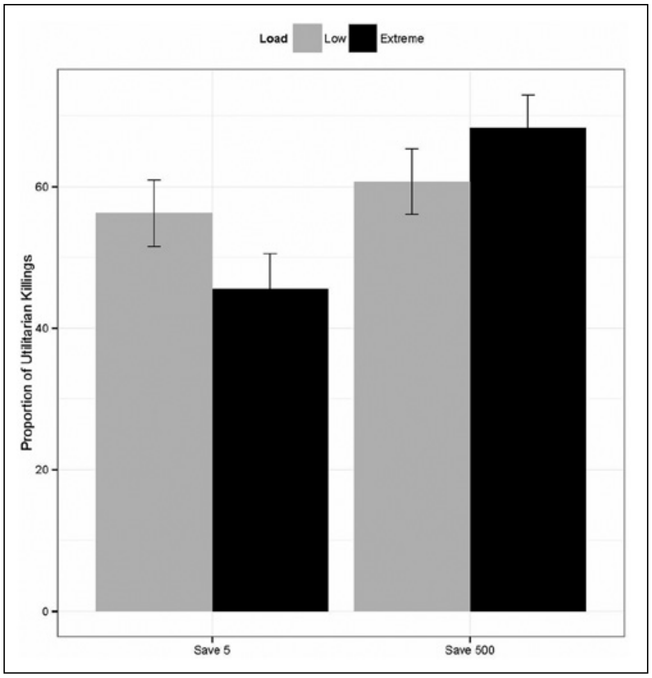

Figure 2. Percentage of utilitarian responses as a function of kill–save ratio and cognitive load (Experiment 1).
Note. Errors bars indicate standard error of the mean.

load condition, suggesting that participants gave high priority to the secondary task.

Figure 2 displays the percentage of utilitarian responses as a function of cognitive load and kill–save ratio. Visual inspection suggests that (a) utilitarian responses were more frequent when the kill–save ratio was 1:500 and (b) cognitive load impacted utilitarian responses only when the ratio was 1:5.

The fact that a better kill–save ratio encouraged utilitarian responses (McNemar test, χ^2 = 7.59, p = .006) was anticipated, but our hypothesis of interest is that of a different effect of load on 1:5 and 1:500 scenarios. To test this hypothesis, we constructed our main dependent variable as the difference between the decisions in the 1:500 and 1:5 scenarios. Specifically, we want to demonstrate that load impacts this difference variable. We coded the difference variable as 1 for participants who were utilitarian for the 1:500 ratio but not for the 1:5 ratio; 0 for participants showing the opposite pattern; and missing otherwise. We then analyzed this dichotomous variable by fitting a logistic regression model taking into account load, age, gender, scenario and performance on the Dot Memory task.

Results are displayed in Table 1. The effect of load is marginal, p = .093, and no other effect is detected. This marginal effect suggests (with all due caution) that load may have affected the 1:5 scenario more than the 1:500 scenario. A chi-square analysis did not detect an effect of load for any scenario, though χ^2 = 2.44, p = .12 for the 1:5 ratio, and χ^2 = 1.33, p = .26 for the 1:500 ratio. We postpone the interpretation of these results until after we report the results of our other experiments, starting with Experiment 2, which is a conceptual replication of Experiment 1, replacing concurrent load with time pressure.

Experiment 2

Method

As in Experiment 1, the 123 participants (83 women, *M* age = 23.3, *SD* = 4.9) were recruited through a French online data collection platform.[2] Participants were randomly assigned to one condition of a 2 × 2 mixed design, manipulating time pressure (between-participant), kill–save ratio (within-participant) and scenario (within-participant).

Each participant saw two dilemmas (again, Captive Soldier and Crying Baby, see Appendix A) and fillers. The order in which these dilemmas appeared, as well as which one featured the 1:5 ratio and which one the 1:500 ratio, was randomly determined for each participant.

Time pressure was manipulated as in Suter and Hertwig (2011). All participants had 35 s to read the scenario, after which the question appeared automatically. Participants in the time pressure group had 8 s to respond, whereas participants in the control group could respond at their own pace.

Results

Figure 3 displays the percentage of utilitarian responses as a function of the kill–save ratio and time pressure. Visual inspection suggests that (a) utilitarian responses were more frequent when the kill–save ratio was 1:500 and (b) time pressure impacted utilitarian responses only when the ratio was 1:5.

As in Experiment 1, the fact that a better kill–save ratio encouraged utilitarian responses (McNemar test, χ^2 = 8.51, p = .003) was anticipated, but our hypothesis of interest is that of a different effect of time pressure on 1:5 and 1:500 scenarios. We thus applied exactly the same coding scheme and analysis strategy as in Experiment 1. We coded our main dependent variable as 1 for participants who were utilitarian for the 1:500 ratio but not for the 1:5 ratio; 0 for participants showing the opposite pattern; and missing otherwise. We then analyzed this dichotomous variable by fitting a logistic regression model taking into account time pressure, age, gender, and scenario.

The model (see Table 1) revealed a detectable effect of time pressure, p = .016, suggesting that time pressure had a stronger effect in the 1:5 condition than in the 1:500 condition. This result is in line with our expectation that time pressure is less disruptive of decisions about moral dilemmas featuring efficient kill–save ratios. A chi-square analysis detected an effect of time pressure in the 1:5 scenario, χ^2 = 3.67, p = .055, but not in the 1:500 scenario, χ^2 = 0.74, p = .39. Experiment 3 aimed at consolidating these findings by

Table 1. Parameter Estimates of Logistic Regression Models for Experiments 1, 2, and a Linear Model for Experiment 3.

	Experiment 1	Experiment 2	Experiment 3
Gender	−0.27 (0.47)	−0.64 (0.85)	42.08 (27.81)
Age	0.03 (0.02)	−0.44* (0.17)	−0.20 (0.85)
Load/time pressure	0.85 (0.51)	2.71* (1.12)	56.31* (25.83)
Scenario	0.62 (0.45)	−1.82 (0.98)	
Dot Memory Task performance	−0.002 (0.02)		
Constant	−1.12 (1.81)	13.00** (4.51)	165.18** (59.44)
Observations	96	47	203
R^2			.04
Adjusted R^2			.02
Log likelihood	−58.27	−18.47	
Akaike information criterion	128.54	46.93	
Residual *SE*			180.39 (*df* = 199)
F Statistic			2.51 (*df* = 3; 199)

*p < .05. **p < .01.

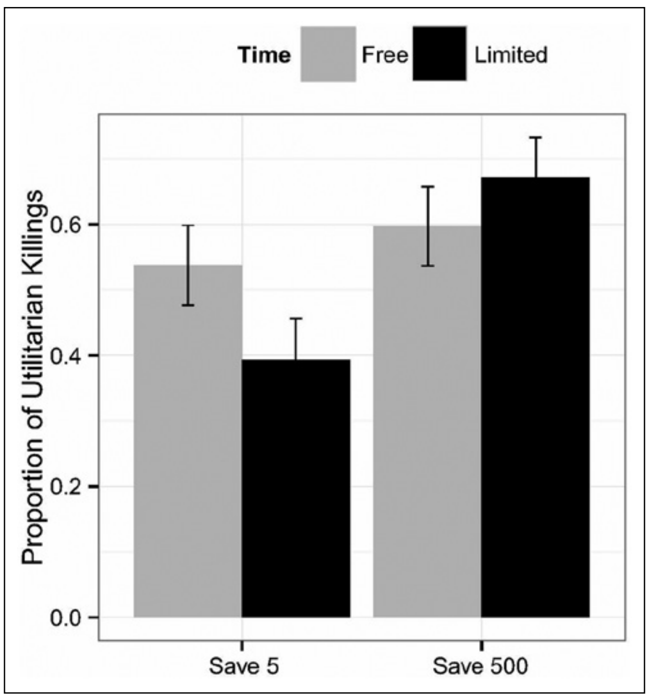

Figure 3. Percentage of utilitarian responses as a function of kill–save ratio and time pressure (Experiment 2).
Note. Error bars indicate standard error of the mean.

introducing a new set of four dilemmas, two new values for the kill–save ratio, and additional experimental controls.

Experiment 3

Method

The 234 participants (164 women and 70 men, *M* age = 34.5, *SD* = 15.1) were recruited through a French online data collection platform.[3] They were randomly assigned to one condition of a 4 × 4 × 2 mixed design, manipulating time pressure (between-participant), kill–save ratio (within-participant), and scenario (within-participant). The manipulation of time pressure was the same as in Experiment 2.

Each participant saw four dilemmas (different from the dilemmas used in Experiments 1 and 2, see Appendix B). The order in which these dilemmas appeared, as well as the kill–save ratio they each featured, was randomly determined for each participant. The four kill–save ratios were 1:5, 1:50, 1:500, and 1:5,000.

Immediately after one of the dilemmas (randomly determined for each participant), a surprise memory check asked how many people could be saved in case action was taken. This allowed us to check whether participants in the two conditions paid similar attention to these numbers, and to eliminate participants who did not pay attention.

Results

Because there was room for interpretation of our memory check question (see Appendix B for contents), we adopted a lenient criterion and accepted as correct the target number plus or minus 1 (i.e., in the 1:500 scenario, we accepted as correct the responses 499, 500, and 501). Only 12% participants failed this surprise memory check, and the proportion was sensibly similar in the two conditions (free time: 14%, time pressure: 11%). We eliminated these 30 participants from further analysis, leaving a final sample of 204 participants (143 women, *M* age = 33.9, *SD* = 15.0).

Figure 4 displays the percentage of utilitarian responses as a function of the kill–save ratio and time pressure. Visual inspection suggests that (a) utilitarian responses increased when the kill–save ratio increased and (b) the effect of time pressure decreased when the kill–save ratio increased.

Because each participant read four scenarios instead of two, our main dependent variable in Experiment 3 is the slope of the utilitarian response as a function of the log of the kill–save ratio. We then analyzed this variable by fitting a

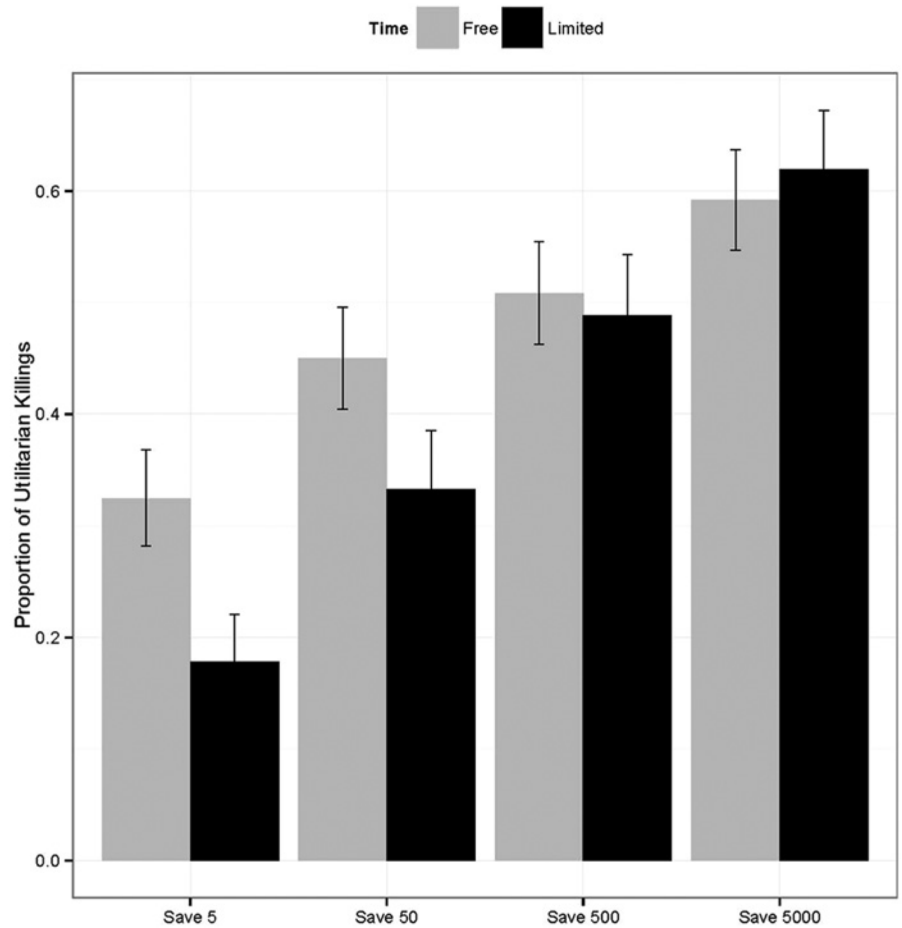

Figure 4. Percentage of utilitarian responses as a function of kill–save ratio and time pressure (Experiment 3).
Note. Error bars indicate standard error of the mean.

simple linear regression model featuring time pressure, age and gender. The model detected a significant effect of time pressure, $p = .03$ (see Table 1), suggesting that the slope of utilitarian responses was steeper for participants under time pressure. As is visually clear in Figure 4, participants under time pressure gave less utilitarian responses than control participants to scenarios featuring low kill–save ratios, but reached the same rates of utilitarian responses for the highest kill–save ratios.

When looking at each kill–save ratio independently, chi-square analyses showed that time pressure impacted the 1:5 scenario, $\chi^2 = 5.44$, $p = .02$; had a marginal effect in the 1:50 scenario, $\chi^2 = 2.80$, $p = .09$; and no detectable effect in the 1:500 scenario, $\chi^2 = 0.08$, $p = .78$, or the 1:5,000 scenario, $\chi^2 = 0.15$, $p = .69$.

General Discussion

The dual-process model of moral judgment postulates that utilitarian responses to moral dilemma (e.g., accepting to kill one to save five) are demanding of cognitive resources. In line with this postulate, we observed that interference effects (time pressure and cognitive load) appeared to reduce the likelihood of a utilitarian response to classic, kill-1-to-save-5 dilemmas. These effects were on the weak side, though, lending support to the speculation that the cognitive demands of utilitarianism might be small (Greene et al., 2008; Trémolière et al., 2012). These demands, however, disappeared as the number of lives to be saved reached the hundreds or the thousands. Not only participants were more likely to be utilitarian when faced with these very efficient kill–save ratios, but these ratios protected the utilitarian response against the interference of concurrent cognitive load and time pressure.

A strict interpretation of the dual-process model would not predict this effect. No matter how many are saved, killing one remains the utilitarian response, and not killing one remains the non-utilitarian response. The suggestion of Kahane et al. (2012) to move away from the utilitarian versus non-utilitarian dichotomy, and to consider instead the intuitive versus counterintuitive dichotomy, does not help much to explain the finding: No matter how many are saved, the duty not to kill is as stringent, and its transgression as counterintuitive.

Recent syntheses on dual-process models (Evans & Stanovich, 2013; Kahneman, 2011) appear to favor default-interventionist architectures, in which default intuitive responses are generated first, followed (or not) by an effortful intervention of reflective processing. If we seek to reconcile our results with such an architecture, we are forced to assume that very efficient kill–save ratios allow utilitarian responses to be generated without the effortful intervention of the deliberative system. That is, we must assume that very efficient kill–save ratios generate *conflicting intuitions*, rather than a conflict between intuition and deliberation.

To understand how this might be the case, we need to consider one fundamental asymmetry between the utilitarian and the non-utilitarian responses in that the utilitarian response implies to actively kill someone, whereas the non-utilitarian response implies to passively let people be killed. It is well established that people are biased to find omissions more acceptable than actions (Cushman, Young, & Hauser, 2006; Ritov & Baron, 1999), and neuroimaging evidence suggests that this is an uncontrolled, automatic bias (Cushman, Murray, Gordon-McKeon, Wharton, & Greene, 2012). As a consequence, the emotional poignancy of killing one is greater than that of letting five die. But there must be a limit to the omission bias: Presumably, the emotional poignancy of letting 5,000 die is at least equal to that of killing 1. If this is correct, then a dilemma featuring a 1:5,000 kill–save ratio is no longer a tug-of-war between emotion and deliberation, but rather a case of emotion against emotion. Quite remarkably, this would mean that efficient kill–save ratios do not promote utilitarian responses because they increase global utility—but rather because they make it emotionally hard not to be utilitarian.

An alternative account of our findings would dispense with emotions but assume a different architecture for the two systems. In the parallel-activation architecture, the intuitive and deliberative components of the dual-process model (or even more neutrally, the arguments for each decision) are activated concurrently, and compete to have the final say. Although this architecture is not the most popular among reasoning or decision theorists, it recently gained traction in the moral domain (Baron, Gürçay, Moore, & Starcke, 2012; Koop, 2013).

The simplest account of our results, within a parallel-activation model, is that increasingly large numbers increasingly activate the deliberative system (or increasingly support the utilitarian response), while the activation of the intuitive system (or the support for the deontic response) stays the same. Accordingly, the decision to be utilitarian would become increasingly frequent and increasingly easy.[4] One potential problem with this explanation, though, is that one may expect an overwhelming number of utilitarian responses for the higher kill–save ratios. However, the rate of utilitarian responses remains a meager 60% even for the massively efficient 1:5,000 ratio. This relatively low proportion of utilitarian responses, for extremely efficient ratios, has already been evidenced. In one study (Greene et al., 2008), 21% of participants rejected a utilitarian killing even though it would save millions of lives. In another (Nichols & Mallon, 2006), 24% of participants found that killing a person was wrong and not the thing to do even if it would save *billions* of lives. The fact that a substantial proportion of individuals can resist the utilitarian response no matter what the figures are, suggests the existence of very strong individual differences in the disposition to endorse utilitarian judgments (Baron et al., 2012).[5]

Clearly, contrasting the default-interventionist and the parallel-activation accounts will require to go beyond the data we report in this article. Critical predictions of the default-interventionist models can be tested with the help of measures of emotional arousal, and potential limitations of the parallel-activation model may be solved by the statistical modeling of individual differences. These extensions are outside the scope of our current contribution, but we are confident that they will be built on the empirical foundations we offered in the present article.

Appendix A

Dilemmas Used in Experiments 1 and 2

The original French version of the dilemmas are available from the authors. In each dilemma, *N* could be 5 or 500.

Captive Soldier. Sébastien is a military officer who was captured by the enemy with his *N* troopers. One of his men managed to escape and is now hiding. An enemy leader is looking for this man, and announces that he will kill the *N* troopers if he does not find the man within 2 hr. Sébastien knows where the man is hiding, and if he reveals where the man is, the man will be killed instead. The only way for Sébastien to save the *N* troopers is to reveal where the man is. Is it morally acceptable for Sébastien to reveal where the man is to save the *N* troopers? (Yes/No)

Crying Baby. Leo is a civilian during war. He, his children, and *N* other people are hidden in a cellar. If the enemy sees them, they will all be captured and killed. The youngest child is still a baby. Enemy soldiers are searching the house when the baby starts to cry. Leo puts his hand over the baby's mouth so that the noise does not attract the enemy soldiers' attention. The only possibility for Leo not to get caught with his children and the *N* other people is to leave his hand on the baby's mouth, which will deprive the baby of air for a few minutes and choke him to death. Is it morally acceptable for Léo to choke the baby to death to save the *N* other people? (Yes/No)

Appendix B

Dilemmas Used in Experiment 3

The original French version of the dilemmas is available from the authors. In each dilemma, *N* could be 5, 50, 500, or 5,000.

Bomb. Julien works in a company that employs *N* people. One morning, a masked and armed man comes in the building and warns that he planted explosives in the building, and that the countdown has started. He turns to Julien and offers a deal: If Julien kills a random colleague, the explosion will be canceled. The only way for Julien to save *N* employees is to kill one at random. Is it morally acceptable for Julien to kill one at random to save *N* others? (Yes/No)

Antibodies. Gilles is an army doctor. After a massive attack from the enemy, *N* soldiers have been urgently transported in the military medical facility. They all need rare antibodies. Another soldier is also within the facility, who was put in an artificial coma a few days earlier. His life is not in danger. This soldier could be harvested for the antibodies that would save all the others. The only way for Gilles to save the *N* soldiers is to harvest the antibodies of the other soldier, who will die in the process. Is it morally acceptable for Gilles to harvest this soldier for antibodies and kill him in the process, to save the *N* other soldiers? (Yes/No)

Hacker. Jean is an army general. He is based in a building in which *N* other people work. One day, the alarm unexpectedly warns people that a missile is about to be launched. A hacker announces by radio that he can direct this missile at any time on the building where Jean and the *N* other people are working. The hacker announces that if Jean accepts to direct the missile himself on the medical annex of the building, then he and the other *N* persons will be saved. However, one person is working in the medical annex, and will be killed. The only way for Jean to save the *N* person is to direct the missile on the medical annex and kill the person who is working there. Is it morally acceptable for Jean to direct the missile on the medical annex and kill the person who is working there, to save *N* others? (Yes/No)

Virus. Bertrand is in charge of a big pharmaceutical company. A terrible virus has spread in a part of the building where *N* people work. These people are now confined, and they will die if they do not get the proper antidote. Two chemical compounds were recently designed in the lab, one of which is the antidote, but Bertrand cannot tell which is which. There are two other men with Bertrand both under his responsibility. The only way for Bertrand to save the *N* employees is to inject each compound to one of these men, killing one in the process. Is it morally acceptable for Bertrand to kill one of these men to save the *N* other employees? (Yes/No)

Acknowledgment

We are grateful to Josh Greene for sharing the data of his 2008 Cognition paper and to the reviewers for their helpful comments.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

Notes

1. Where exact ratios were not specified, we used minimal interpretations, for example, 200 lives for "hundreds of lives."
2. A total of 15 other individuals started the study but did not complete it, including 11 in the time pressure condition.
3. A total of 93 other individuals started the study but did not complete it, including 70 in the time pressure condition.
4. A twist on this account is to consider that cognitive load prompts people to ignore numerical information, unless the numbers are huge and therefore very salient. That is, small numbers (5, 50) would not enter the moral judgment under load, but large numbers (500, 5,000) would. Once numbers enter the moral judgment, they support the utilitarian decision. As a consequence, people would not only become more utilitarian for large numbers, but they would be less sensitive to load for these same large numbers. We note however that when given a surprise memory check for the numbers involved in the dilemma (Experiment 3), our participants performed the same under time pressure and in control conditions.
5. Another relevant consideration here is that the deontological versus utilitarian distinction is typically conflated with omissions versus commissions, present article included. Perhaps utilitarian responses could gather larger support if they corresponded to omissions, rather than commissions. It would be accordingly useful to introduce the orthogonal manipulation of which response requires action within the experimental designs used in the current article.

References

Baron, J., Gürçay, B., Moore, A. B., & Starcke, K. (2012). Use of a Rasch model to predict response times to utilitarian moral dilemmas. *Synthese*, *189*, 107-117.

Baron, J., & Leshner, S. (2000). How serious are expressions of protected values? *Journal of Experimental Psychology: Applied*, *6*, 183-194.

Bartels, D. M., & Medin, D. L. (2007). Are morally motivated decision makers insensitive to the consequences of their choices? *Psychological Science*, *18*, 24-28.

Bethell-Fox, C. E., & Shepard, R. N. (1988). Mental rotation: Effects of stimulus complexity and familiarity. *Journal of Experimental Psychology: Human Perception and Performance*, *14*, 12-23.

Cushman, F., Murray, D., Gordon-McKeon, S., Wharton, S., & Greene, J. (2012). Judgment before principle: Engagement of the frontoparietal control network in condemning harms of omission. *Social Cognitive and Affective Neuroscience*, *7*, 888-895.

Cushman, F., Young, L., & Hauser, M. (2006). The role of conscious reasoning and intuition in moral judgment: Testing three principles of harm. *Psychological Science*, *17*, 1082-1089.

De Neys, W. (2006). Dual processing in reasoning: Two systems but one reasoner. *Psychological Science*, *17*, 428-433.

De Neys, W., & Verschueren, N. (2006). Working memory capacity and a notorious brain teaser: The case of the Monty Hall dilemma. *Experimental Psychology*, *53*, 123-131.

Evans, J. S. B. T., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science*, *8*, 223-241.

Greene, J. D., Morelli, S. A., Lowenberg, K., Nystrom, L. E., & Cohen, J. D. (2008). Cognitive load selectively interferes with utilitarian moral judgment. *Cognition*, *107*, 1144-1154.

Greene, J. D., Nystrom, L. E., Engell, A. D., Darley, J. M., & Cohen, J. D. (2004). The neural bases of cognitive conflict and control in moral judgment. *Neuron*, *44*, 389-400.

Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI investigation of emotional engagement in moral judgment. *Science*, *293*, 2105-2108.

Kahane, G., Wiech, K., Shackel, N., Farias, M., Savulescu, J., & Tracey, I. (2012). The neural basis of intuitive and counterintuitive moral judgment. *Social Cognitive and Affective Neuroscience*, *7*, 393-402.

Kahneman, D. (2011). *Thinking, fast and slow*. New York, NY: Farrar, Strauss and Giroux.

Koop, G. J. (2013). An assessment of the temporal dynamics of moral decisions. *Judgment and Decision Making*, *8*, 527-539.

Miyake, A., Friedman, N. P., Rettinger, D. A., Shah, P., & Hegarty, M. (2001). How are visuospatial working memory, executive functioning, and spatial abilities related? A latent-variable analysis. *Journal of Experimental Psychology: General*, *130*, 621-640.

Nichols, S., & Mallon, R. (2006). Moral dilemmas and moral rules. *Cognition*, *100*, 530-542.

Ritov, I., & Baron, J. (1999). Protected values and omission bias. *Organizational Behavior and Human Decision Processes*, *79*, 79-84.

Suter, R. S., & Hertwig, R. (2011). Time and moral judgment. *Cognition*, *119*, 454-458.

Trémolière, B., De Neys, W., & Bonnefon, J. F. (2012). Mortality salience and morality: Mortality salience makes people less utilitarian. *Cognition*, *124*, 379-384.