

Research Article

Certainty and Disagreement in the Moral Domain: A Moral Dilemmas Perspective

Alejandro Rosas¹

1. Universidad Nacional de Colombia, Colombia

In *On Certainty*, Wittgenstein persuasively argued that certainties are like hinges, grounding moves in epistemic space. He provided instances from several fields of discourse, excepting morality. Pleasants proposed 'Killing innocent, non-threatening people is wrong' as a moral certainty. His critics pointed to some uncontroversial counterexamples, whose common element is, implicitly, a greater good that justifies harm. I therefore propose including the greater good as an equally foundational moral certainty that sometimes prevails over individual rights. Such prevalence is sometimes found in moral dilemmas, i.e., in conflicts between rights and the greater good. Their discussion has led to notice auxiliary rules active in their rational resolution. I describe three such rules and their rationales and present preliminary empirical evidence suggesting that rights, the greater good, and the auxiliary rules mediating in their conflicts are widely shared and used in moral judgment. I propose to see these certainties examples of overlapping and shared epistemic hinges potentially resolving deep disagreements. The point is not to contribute more subtleties to the conceptual discussion, but to submit the certainties involved in moral dilemmas as illustrating the notion of overlapping and shared hinge commitments.

Correspondence: papers@team.qeios.com — Qeios will forward to the authors

1. Introduction

Wittgenstein's epistemological reflections in *On Certainty* argued in favour of the existence of basic certainties. These certainties configure the common world we inhabit. While Wittgenstein's examples spanned empirical claims, mathematical truths, scientific propositions, semantic norms, and religious beliefs^[1], notably absent was any mention of moral certainties. However, it seems to follow from Wittgenstein's reflections that any discourse is only meaningful if it presupposes certainties functioning

as hinges. Without further discussion of this basic idea I will assume in the following that, however one wants to understand it in detail, it implies that some moral certainty or certainties are necessary to give meaning to moral discourse. These certainties belong to the basic conceptual scheme, the fundamental framework in which we make, justify, and argue about moral judgments^[2].

Pleasants^{[3][4]} argued that

(1) Killing (harming) an innocent, non-threatening person is wrong.

is a Wittgensteinian kind of certainty in the moral domain. In some sense, it constitutes this domain. However, numerous counterexamples were cited to challenge this claim^{[5][6][7][8][9][10][11][12]}. *Prima facie*, the list of exceptions to (1) is so long that this basic moral certainty apparently ‘suffers the death of a thousand qualifications’^[4].

Some exceptions advanced by the critics are not genuine counterexamples. They concern beings who are denied the status of innocents and/or of persons. Of this sort are Nazi Germany killings, slavery, and some arguments in favour of abortion. They depend on demoting or dehumanizing people, ‘re-categorized as dangerously threatening non-persons’^[4], or as a ‘kind of restricted human being, suited only for enslavement’^[4].

However, genuine exceptions to (1) are known, and not merely as minority opinions, implying, apparently, that (1) is not a basic certainty. Some actions, like euthanasia and abortion in carefully delimited cases, are morally justified publicly with beliefs that contradict (1). Those beliefs belong in a different category as the dehumanizing beliefs of a Nazi or a slaver. Two examples are uncontroversial. I give reasons for dismissing a third one.

- a. Killing of innocent civilians as a foreseen side effect of military operations in war is, under certain conditions, not consensually judged wrong. Philosophers have provided the rationale^[13] (Kamm 2012). Pleasants says they are not exceptions because the choice is directed at the lesser of two evils^[4]. But this implicitly concedes that killing an innocent person (or more) is sometimes not wrong, namely when it is the lesser of two evils^{[8][6]}.
- b. Some philosophers have argued that suicide^[14], or even killing an innocent person in medically assisted dying^[15], is not wrong. Many people are prepared to agree and act on this belief. We can say the same about doctors who ‘kill’ a person during surgery, which the patient, knowing the risks, agreed to undergo for a chance at prolonging her life. Understandably, proponents avoid using the

word ‘killing’, which arguably contains an evaluative hue evoking images or feelings attached to the word ‘murder’. We here use ‘killing’ only to express that an intentionally performed action or omission was causally active in bringing about a person’s death, even if the action aimed at her greater good, as is clearly the case in the aforementioned cases.

- c. Some people have a notable difficulty in comprehending and internalizing the moral justification of actions. They cannot conceive the thought: ‘This is morally wrong’^[5]. Though it is difficult to capture morality in a simple phrase, I go with Kant’s^[16] claim that morality implies conceiving and treating people not merely as means, but also as ends. People who consistently subordinate others entirely to their personal selfish goals are known to medical psychiatry as antisocial personalities. If it suits their needs, they will believe that killing an innocent person is not wrong, because ‘wrong’ means for them unfit to achieve personal goals, and many morally wrong actions promote personal goals. Since it is likely, in my opinion, that antisocial personalities do not endorse moral language for whatever reasons – physiological, environmental, or both – but rather an extreme form of ‘moral egoism’^[17], I will not consider this phenomenon as a genuine challenge to (1), which presupposes, in my understanding, that the speaker appreciates and endorses moral language.

However, I do think that (a) and (b) describe widespread beliefs that fly in the face of (1). Note that the intuitions that support these beliefs include the thought of a greater good, which also has the normative pull expected of moral certainties. In (a), the greater good is *interpersonal*: more lives are saved than lost, and in the permissible cases (e.g., just war), the lives lost are taken in self-defence. In (b), the greater good is *intra-personal*: all life must end, but it should not end in a prolonged agony predictably leading to death. Alternatively, a patient dies undergoing a risky surgery for a chance at prolonging her life, which promises to bring more joy or to allow the completion of a project. In sum, many people have had and still have beliefs implying that bringing about the death of an innocent person is permissible as long as the greater good is or was included in the direct intention. Even if we were not to find this view altogether compelling, most of us recognize the moral belief on which it is based. We present preliminary empirical evidence for this claim below in section 4.

2. A conciliatory solution: two certainties in conflict

As anticipated, the upshot of these challenges to (1) need not deny (1) as a moral certainty. On the contrary, we should accept a second moral certainty: the obligation to promote the greater good^[18]. This certainty can be formulated thus:

(2) The obligation to promote the greater good sometimes overrides (1).

More generally, (1) can be understood as the duty to uphold individual rights (among them, foremost, the right to life). Here, I part company with Pleasants^[4], who rejected that (1) implies the idea of rights, though the notion of right to life can be easily connected to (1). Indeed, the concept of a right is highly abstract and theoretical. “Humans (and plausibly animals) have a right to life” does not seem to qualify for the kind of visceral, arational certainties Wittgenstein meant to capture^[2]. Recall, however, how Mill grounded the idea of a right partly in “desires to retaliate” common to most animals^[18] and explained thereby the strong “feeling of repugnance” attached to the wrongfulness of injustice^[18]. In the case of the right to life, the connection to visceral feeling seems immediate, while in other rights (e.g., the right to inheritance), it is more dubious. The point now is that even a basic right like the right to life sometimes conflicts with promoting the greater good. If we accept that these two principles embody legitimate certainties, an important question is whether conflicts between them can be resolved.

Given that (a) and (b) describe widespread beliefs, let us assume, provisionally, that both certainties are fundamental and equal in strength. If their conflict led to contradictory propositions like ‘A is wrong and permissible’ or ‘Not-A is obligatory and A is permissible’^[19], then moral judgment would completely fail us in moral dilemmas. This is not what philosophers and social scientists have observed. People use rules or criteria to distinguish cases where (1) can be legitimately overridden from cases where it cannot.

Aiming at a greater good is a necessary, but not a sufficient, condition for morally approving a suspension of rights. The greater good is surely part of the answer to the question about when and why (1) has counterexamples, but it is not the complete answer. Justifiably suspending rights needs to include other moral factors. A reasonable conjecture is that right-suspensions are legitimate when upholding them negatively impacts other right holders or when the party who would be harmed consents (e.g. to undergo a risky therapy) for their own greater good.

To derive the rules that justify suspending rights, philosophers imagined hypothetical scenarios where individual rights conflict with the greater good. In such thought experiments, philosophers observed their own judgment tendencies and proceed to extract plausible rationales. Since they were guided in each case by their own intuitions, and intuitions vary among philosophers and the folk alike (e.g. ^[20]), a prudent estimate of their universality requires probing suitably sized samples with surveys and experiments. This procedure tests whether their moral judgments find an echo in the population.

Both philosophers^{[21][22][23][24][25]} and social scientists^{[26][27][28][29][30]} (Mikhail 2007; and many others) have observed how judgments sometimes follow (2) and sometimes let (1) prevail. Although the mainstream scientific view^{[27][31]} appeals to intuitive vs. reflective modes of cognition to explain this variation, a rule-guided resolution could solve the conflict establishing which certainty should reasonably prevail. In section 3, I describe three rules postulated by philosophers. In section 4, I report the results of a study conducted with scenarios representing the three rules. This study shows two things. First, many participants are convinced that the rules sometimes fully justify the suspension of rights for the greater good. Second, though at least as many stop short of giving their full approval, their moral responses show sensitivity to the moral force of those reasons. Some disagreement therefore remains, but it is framed by rules whose moral relevance is recognized by all. In section 5, I note that it can be captured with the notions of deep^[32] and faultless^[33] disagreement, and agree with Pritchard^[34] that it is not necessarily intractable.

3. Three rules unveiled by philosophy and social psychology

Philosophers' thought experiments with hypothetical sacrificial moral dilemmas have inspired the stimuli used in actual social-science experiments. Formulating scenarios that vary in the ways and contexts in which harm is inflicted, philosophers and scientists disclosed shared intuitions tending to approve of the sacrifice in some cases and to disapprove in others^{[35][21][27][31][28][24][36][30][22]} (Christensen et al. 2014; Hauser et al. 2007; Mikhail 2007; and many others).¹ This pattern of judgments is not random. Researchers have detected features that significantly affect moral judgments. These features trigger the implicit rules that establish in which cases overriding (1) is morally legitimate provided the action aims at a greater good, even if the consent of the victim is absent.

3.1. Harm as a side-effect

One example is the centuries-old rule attributed to St. Thomas Aquinas, distinguishing between harm intended as a means and harm foreseen as a side effect, a rule often followed subconsciously^[26]. The first rule has been widely researched^{[21][28][24][22][23]} (Feltz & May 2017; Hauser et al. 2007; Mikhail 2002, 2007; Waldman & Dieterich 2007).

Folk intuitions condemn harming a person intentionally as a means, but tolerate harming her as a foreseen side effect, even though in both cases the ultimate intention is to save more^{[35][24]} (Hauser et al.

2007; Mikhail 2007; but see [28] for a different view²). Harming as a means is exemplified in *Footbridge*, where a bystander is pushed in front of a vehicle about to run over and kill five workers. Her body and person are thus used as a physical means to stop the vehicle and save the five. Harming as a side effect is represented in a similar scenario (*Switch*), except that, to save the five workers, the vehicle is diverted to the only available side-track, where, unfortunately, one worker will be killed.

Not all philosophers, however, agree that harming in *Switch* is more permissible than in *Footbridge*. Foot[21] imagined a similar scenario – *Fumes* – where a substance is fabricated that will save five patients, but will also release lethal fumes in the room of one patient. This patient will die as a side effect, but Foot believed the action to be nonetheless wrong. Her reason was that the obligation not to harm has precedence over the obligation to help. The action in *Fumes* violates this because the agent consents to harm one for the sake of helping more. Thomson[22] suggested that harming in *Switch* is more permissible than in *Footbridge* because diverting the vehicle is not in itself a violation of the victim's rights. However, she later claimed that both harms are wrong, even though the assault on the rights of the victim is, irrelevantly from a moral viewpoint, less drastic in *Switch* than in *Footbridge*[23].

One way of expressing both Foot's and Thomson's[23] view is that killing a victim as a side effect, foreseen with complete certainty, treats the victim as a means although in a more subtle sense. One decides to help, knowing that one person will be harmed (killed) by the act of helping. Even if the action does not use a person and her body as a physical means, it can be described as consenting to harm one person to help more. This argument seems fatal for the view that sometimes the greater good should prevail. It implies that it is always wrong to save lives if this necessarily implies killing one (or more) and this is known to the agent. Basically, if helping involves harming, you are not allowed to help unless, perhaps, the harm does not exceed a low threshold. This, however, opens a door to justifying killing one as a means if the number of people saved is very large, a view labelled 'threshold deontology'[37][38].

Reading Foot[21] in this way is possible, but not necessary. More probably, she meant to put the accent on this: a positive numerical balance of lives saved minus lives lost is a weak reason to override the prohibition to kill innocents. Some other moral factor must concur if one is to consent to harming a person for the sake of helping more. This implies that Foot deferred to the authority of the greater good, as long as it was not just about numbers. And she said indeed: 'I have not been arguing for any general conclusion such as that we may never, whatever the balance of good and evil, bring injury to one for the sake of aid to others, even when this injury amounts to death.'[21]

This is probably a widespread attitude. McMahan^[39] argued that killing as a foreseen side effect is wrong when it violates the victim's rights, but also recognized that weighing harm against benefit merits careful consideration. Before pointing to other factors that might incline the balance towards the greater good, it is fair to mention Baumard and Sheshkin's (2015, p. 45) deontological-contractualist account of why people tend to approve in *Switch* but not in *Footbridge*. In *Switch*, they say, both the person on the sidetrack and the five people on the main track are on potential paths of the threat that would kill them. Since no party is actually safe from the threat, diverting the threat is allowed, whereas making a man who is *safe* from any potential path of the threat unsafe by pushing him, as in *Footbridge*, is forbidden. But even if one abstracts from the fact that the threat in *Switch* is usually depicted as directed by default at the five, this would not by itself justify diverting the threat, as one can easily see by imagining that the same number of people stand on both branches of the fork. Thus, it is reasonable to claim that saving more is also a necessary condition. Thus, their contractualist approach also suggests that permissibility intuitions target circumstances where upholding the right of one would unduly ignore the rights of the many. Two more cases, explained below, illustrate this last notion. These cases were formulated by Foot^[21] as she reflected on a scenario in which several workers are threatened with drowning by rapidly rising waters in an underground cave. The burly man leading the escape gets stuck at the cave's exit, blocking it for the remaining co-workers. Let us call this scenario *Flooding Cave*.

3.2. Inevitable death of the victim

In one version of *Flooding Cave*, the burly man leading the escape gets stuck at the exit with his head inside the cave. He will drown with all the others unless something is done, e.g., blowing the narrow rocky passage where the man is stuck. Generally speaking, when a person, call her the victim, is imminently going to die by a threat that will also kill several others, doing something to save the others is permissible even though it would kill the victim, who would die anyway if no action were taken. Her inevitable death seems to justify action. In Foot's words: 'it is reasonable that the action that will save someone should be done'^[21]. This very closely approximates a Pareto optimal solution, making some people better off without making anyone worse off^{[30][29]}. Despite, or even because of the sacredness of human life, cases of inevitable death force us to consider the other lives in danger who could be saved by accelerating the death of one.

3.3. *The victim as an innocent threat*

The idea of an innocent threat, though not in these exact words, appears already in Foot^[21]. Nozick was perhaps the first to use this label (Nozick 1974: 34-5). Thomson's discussion^[40] lists some cases in which the victim could be an innocent threat to others. She argues for overriding the right to life of innocent threats based on the right of the threatened to self-defense. She fails, however, to mention one type of case that Foot^[21] did consider. This case is included below as case (iii). When more people are innocently threatened by one, the greater good gives, even to third parties, a moral reason to intervene in favor of the many. The cases illustrating an innocent threat are the following:

- i. A natural event (strong wind) throws the victim with killing force against another person, who has the right to defend herself, e.g., by 'disintegrating' the innocent victim, if she could.
- ii. The innocent victim is placed by others against her will in a situation where she constitutes a threat to others. The threatened have the right to self-defense even if that kills the victim.
- iii. The victim does something that, unintentionally and without her foresight, creates at the same time a threat to others. The others have the right to do something that would save them, even if that were to kill the innocent victim. This was illustrated by a second version of *Flooding Cave*^[21]. In this version, the burly man gets stuck at the cave's exit, with his head *outside* the cave. He will survive, but the others will drown, because he is blocking the exit.

More generally, a permissible action causally triggers, as a side effect, a threat of death to others. Precisely because the agent cannot anticipate this side effect, she is innocent of aggression. Nonetheless, the fact that he causes the threat to others increases the moral weight of the lives of those threatened, particularly if overriding his right to life is the only way to save those threatened.

This last case can be read as an allegory of the most frequent, and thus more relevant, structure of innocent threat cases. A recent real example, providing a less dramatic illustration of the basic rule, is the new attitude toward smoking. Smoking is a personal decision, which adults have the right to make. After medicine discovered the health hazard for passive smokers^[41], imposing this hazard on others became unjustifiable. Many countries have passed laws against smoking in public spaces shared with non-smokers. Consequently, the right to smoke has been partially overridden. Increasingly, people see this ruling as legitimate. One can also point to negative externalities of industrial economic activities exhibiting the same *innocent threat* structure.

4. Shared principles and auxiliary rules

Hypothetical moral dilemmas, because they pit the right to life of one against suspending that right to save more lives, reveal that people care both for upholding individual rights and for promoting the greater good. We are entitled to assert this because in cases like *Footbridge* most people disapprove of saving the many at the cost of one. However, in cases like *Switch* or *Cave*, responses favor the many. The question is how to explain this change of opinion, if the numbers of lives lost vs. saved are exactly the same in all three scenarios. This change reveals that a simple positive balance of lives saved vs. lost is not enough to override the right to life of one in order to save more. The scenario must also suggest that upholding the right of the one and ignoring the rights of the many would be unreasonable in some cases. This suggests that the greater good is nothing more, and nothing less, than the most extensive realization of rights compatible with respecting the claims each of the conflicting parties may legitimately advance in a given situation. The rules described above discriminate between cases where the rights of the one should prevail over the many and cases where the opposite is true. We need confirmation, however, that these rules are widely endorsed, because they belong in a unitary framework of moral certainty with the endorsement of rights and of the greater good.

We collected data to confirm the impact of the three rules. Three scenarios independently represented each of the three rules and were compared to *Footbridge*, where none of these rules apply: *Switch* represents harm as a side effect and no other rationale; *Lifeboat* represents harm when the victim will inevitably die and no other rationale; and *Cave* (the man stuck with his head outside the cave) represents harm to a victim who innocently causes a threat to others.

As can be seen in Figure 1, right panel, both participants with strong greater-good attitudes (continuous line) and those with strong deontological attitudes (dotted line) move up the approval rating scale in favor of the greater good when the three auxiliary rules apply each isolated from the other two.³ In both groups, confidence intervals (error bars) suggest that mean approval scores increase significantly compared to *Footbridge*, indicating that participants grant moral relevance to the circumstances of harm containing the triggers of the auxiliary rules (harm as side effect, inevitable death and innocent threat). However, not everyone moves up the scale by the same amount. Participants exhibiting strong deontological attitudes (dotted line) score below the mid-point (3.5 on a scale from 1 to 6), while those with stronger greater-good attitudes (continuous line) score above it. We provide an explanation in the next section. Nonetheless, moving up the scale when auxiliary rules apply signals a willingness to take

the greater good morally into account, for the reasons embedded in the rules as explained above. The data are available here: (https://osf.io/dte8b/overview?view_only=92233792cdb54dbfa227e7297dee8afb)




Figure 1. Subjects (N=460) judged the sacrifice on a scale from 1 (definitely wrong) to 6 (definitely right). Left panel: *Footbridge* received the lowest mean approval (2.67), while *Innocent threat* (Cave) received the highest mean approval (3.66). Right Panel: When divided into low (N=265) and high (N=195), scores in both groups significantly increased from *Footbridge* to *Innocent threat*.

Alt Text: Graphical representation of the mean moral responses to *Footbridge* and three other scenarios where auxiliary rules apply separately. Error bars show the 95% confidence intervals, confirming that mean moral approval increased significantly in the three cases, compared to *Footbridge*. High scorers increased their approval with a steeper slope than low scorers.

5. Disagreement as difference in how people rank shared values

The resulting disagreement between low and high scorers satisfies the four characteristics of deep disagreements^[42], adapted to the moral case (Tejedor 2025). First, the disagreement is *genuine* because, on the same issue, one party can see the other as mistakenly approving or disapproving. Second, both parties offer *moral reasons* for their view: the prevalence of individual rights on the one hand, and the prevalence of the greater good on the other. Third, disagreement is *systematic* because it extends to a range of moral cases. For example, in discussing the morality of abortion, one party would condemn it based on the right to life of the fetus. At the same time, the other would allow it, alleging greater harm to the mother's overall life prospects. And fourth, the disagreement is *persistent*, anchored on individual differences in how people rank the values in question.

Tejedor (2025) argued that on a hinge approach to moral certainty, a paradox arises for deep disagreement, because if the hinges are different, parties would fail to recognize each other's reasons as moral, and thus no moral disagreement would exist. However, Tejedor's description of a plausible interactive disagreement over abortion depicted each party as advancing towards recognition of the other's fundamental moral reasons. Disagreement persists, however, likely because the parties believe in a different order of precedence between the values in conflict (the fetus' right to life, the mother's dignity and welfare). In the hypothetical scenarios we used to collect responses, both participant types increased their approval scores on the rating scale when the triggers of the auxiliary rules were present. This is evidence that even those with strong deontological commitments also care for the greater good and view some circumstances as meriting greater moral approval to the action saving the many, probably because their rights also matter in the situation. I thus submit that individual rights and the greater good are both recognized as basic moral concerns by (nearly) all parties,⁴ differing only in how they rank them. In this sense, both principles structure together a moral language game that most people adopt. No paradox thus arises, though we do experience a phenomenon for which the label 'deep disagreement' feels nonetheless appropriate. However, it does not seem right to treat these disagreements as intractable in principle. Rational resolution of deep disagreements has been defended by Pritchard^[34] and others and is illustrated here in a concrete example of how partially conflicting hinge commitments are shared by all. Moral dilemmas thus illustrate what Tejedor calls crisscrossing, shared hinge commitments, in which a moderate conflict still remains. The shared certainties include auxiliary rules offering the possibility of compromises as rational solutions to the conflict, avoiding incommensurability and strong relativism. Note that there is no fault in assigning different relative importance to the two principles, i.e., deep disagreement is also *faultless*, similar to what happens in disagreement about tastes^[33].

Surely, the consequences of disagreement in moral cases are dire, given the pressure for universality in moral judgment. But it makes a great difference whether the disagreement is grounded on differences in fundamental values, or whether the fundamental values are the same, and disagreement arises instead from assigning them a different rank order. According to our data, those recognizing only one of these principles are a small minority. Recognizing that fundamental values are shared, though differently weighted, should promote a willingness to reach compromises for the sake of coordination and mutual understanding.

Footnotes

¹ There has been some debate on whether the unrealistic character of the scenarios is a handicap. Some believe, plausibly, that their hypothetical character is advantageous, as it avoids political partisanship interfering with moral responses^[43]. Moreover, some of the dilemmas collected from actual historical reports^[44] are, in their morally relevant factors and structure, reminiscent of the second and third hypothetical dilemma-types described below. Several of their examples are cases where the victim to be sacrificed innocently threatens the life of another one, and where the people involved would inevitably die if nothing were done, as in conjoined siblings where saving at least one is an option.

² Greene et al.^[28] and Waldman & Dieterich (2007) found evidence that harming as a means is irrelevant for most participants, while a confounded variable, namely directly harming the victim with personal force, is highly relevant. Challenging their results is important, but addressing this here would require a long detour, distorting the focus of this paper. It suffices here to mention that other studies reported diverging results^[35] (e.g. Hauser et al. 2007; Mikhail 2007).

³ One hundred and sixty of the 460 subjects answered on a scale from 1 to 6 in session 1. In session 2 with different participants, 300 responded on a scale from 1 to 5, and their scores were converted to a scale from 1 to 6 by adjusting the intervals in this way: thus: 1=1, 2=2.25, 3=3.50, 4=4.75, and 5=6. We separated low vs. high using a fifth scenario in both sessions: *Loop* (not shown in line graph). *Loop* represents a scenario where the victim's body is used as a means to stop the threat and save the five, but the agent acts on the threat, as in *Switch*, rather than on the victim, as in *Footbridge*. The agent redirects the threat to a loop in the tracks. Scores from 1 to 3.5 in *Loop* composed the low group (N=265) and from 4 to 6, the high group (N=195).

⁴ In our data, only 18 out of 460 participants (3.9%) consistently gave the lowest score to all scenarios, thus apparently denying moral weight to the utilitarian principle in conflict with individual rights. Even less participants deny weight in such conflict to individual rights.

References

1. ^ΔKusch M (2021). "Disagreement, Certainties, Relativism." *Topoi*. 40:1097–1105. doi:[10.1007/s11245-018-9567-z](https://doi.org/10.1007/s11245-018-9567-z).

2. ^a ^bPritchard D (2025). "Axiological Hinge Commitments." *Synthese*. 205:49. doi:[10.1007/s11229-024-04898-0](https://doi.org/10.1007/s11229-024-04898-0).
3. ^ΔPleasants N (2009). "Wittgenstein and Basic Moral Certainty." *Philosophia*. 37(4):669–679.
4. ^a ^b ^c ^d ^e ^fPleasants N (2015). "If Killing Isn't Wrong, Then Nothing Is: A Naturalistic Defence of Basic Moral Certainty." *Ethical Perspect*. 22(1):197–215.
5. ^a ^bAriso J (2022). "Why the Wrongness of Killing Innocents Is Not a Universal Moral Certainty." *Philos Investig*. 45(1):58–79. doi:[10.1111/phin.12298](https://doi.org/10.1111/phin.12298).
6. ^a ^bAriso J (2022). "Morality in Disguise. A Response to Laves." *Philos Investig*. 45(1):92–97. doi:[10.1111/phin.12332](https://doi.org/10.1111/phin.12332).
7. ^ΔBrice R (2013). "Mistakes and Mental Disturbances: Pleasants, Wittgenstein, and Basic Moral Certainty." *Philosophia*. 41(2):477–487.
8. ^a ^bFairhurst J (2019). "Problems in Pleasants' Wittgensteinian Idea of Basic Moral Certainties." *Ethical Perspect*. 26(2):271–298.
9. ^ΔFairhurst J (2024). "Some Concerns About the Idea of Basic Moral Certainty: A Critical Response to Samuel Laves." *Philos Investig*. 47:119–136. doi:[10.1111/phin.12403](https://doi.org/10.1111/phin.12403).
10. ^ΔLaves S (2020). "Why Natural Moral Certainties Exist: A Response to Fairhurst." *Ethical Perspect*. 27(3):297–315.
11. ^ΔLaves S (2022). "Why the Wrongness of Killing Innocent, Non-threatening People Is a Universal Moral Certainty." *Philos Investig*. 45(1):77–90. doi:[10.1111/phin.12323](https://doi.org/10.1111/phin.12323).
12. ^ΔRummens S (2013). "On the Possibility of a Wittgensteinian Account of Moral Certainty." *Philos Forum*. 44(2):125–147.
13. ^ΔHurka T (2005). "Proportionality in the Morality of War." *Philos Public Aff*. 33(1):34–66. doi:[10.1111/j.1088-4963.2005.00024.x](https://doi.org/10.1111/j.1088-4963.2005.00024.x).
14. ^ΔHume D (1987). "Of Suicide." In: Miller EF, editor. *Essays, Moral, Political, and Literary*. Indianapolis: Liberty Fund.
15. ^ΔSinger P (1979). *Practical Ethics*. Cambridge: Cambridge University Press.
16. ^ΔKant I (1993). *Grounding the Metaphysics of Morals, with On a Supposed Right to Lie Because of Philanthropic Concerns*. Indianapolis/Cambridge: Hackett Publishing Company.
17. ^ΔSidgwick H (1907). *The Methods of Ethics*. 7th ed. London: Macmillan.

18. ^a ^b ^cMill JS (1871). *Utilitarianism*. 4th ed. London: Longmans, Green, Reader, and Dyer. Reprinted in *Essays on Ethics, Religion and Society*. The University of Toronto Press, 1969.
19. ^aTrigg R (1971). "Moral Conflict." *Mind*. **80**(317):41–55.
20. ^aMachery E, Mallon R, Nichols S, Stich S (2013). "If Folk Intuitions Vary, Then What?" *Philos Phenomenol Res*. **86**(3):618–635.
21. ^a ^b ^c ^d ^e ^f ^g ^h ⁱ ^j ^kFoot P (1967). "The Problem of Abortion and the Doctrine of the Double Effect." *Oxf Rev*. **5**:1–5.
22. ^a ^b ^c ^dThomson J (1985). "The Trolley Problem." *Yale Law J*. **94**(6):1395–1415.
23. ^a ^b ^c ^dThomson J (2008). "Turning the Trolley." *Philos Public Aff*. **36**(4):201–212.
24. ^a ^b ^c ^dKamm F (1996). *Morality, Mortality: Volume 2: Rights, Duties, and Status*. New York: Oxford University Press USA.
25. ^aParfit D (2011). *On What Matters*. Oxford/New York: Oxford University Press.
26. ^a ^bCushman F, Young L, Hauser M (2006). "The Role of Conscious Reasoning and Intuition in Moral Judgment: Testing Three Principles of Harm." *Psychol Sci*. **17**(12):1082–1089. doi:[10.1111/j.1467-9280.2006.01834.x](https://doi.org/10.1111/j.1467-9280.2006.01834.x).
27. ^a ^b ^cGreene J, Sommerville B, Nystrom L, Darley J, Cohen J (2001). "An fMRI Investigation of Emotional Engagement in Moral Judgment." *Science*. **293**:2105–2108. doi:[10.1126/science.1062872](https://doi.org/10.1126/science.1062872).
28. ^a ^b ^c ^d ^eGreene J, Cushman F, Stewart L, Lowenberg K, Nystrom L, Cohen J (2009). "Pushing Moral Buttons: The Interaction Between Personal Force and Intention in Moral Judgment." *Cognition*. **111**(3):364–371.
29. ^a ^bHuebner B, Hauser M, Pettit P (2011). "How the Source, Inevitability, and Means of Bringing About Harm Interact in Folk–Moral Judgments." *Mind Lang*. **26**:210–233. doi:[10.1111/j.1468-0017.2011.01416.x](https://doi.org/10.1111/j.1468-0017.2011.01416.x).
30. ^a ^b ^cMoore A, Clark B, Kane M (2008). "Who Shalt Not Kill? Individual Differences in Working Memory Capacity, Executive Control, and Moral Judgment." *Psychol Sci*. **19**:549–557. doi:[10.1111/j.1467-9280.2008.02122.x](https://doi.org/10.1111/j.1467-9280.2008.02122.x).
31. ^a ^bGreene J, Nystrom L, Engell A, Darley J, Cohen J (2004). "The Neural Bases of Cognitive Conflict and Control in Moral Judgment." *Neuron*. **44**:389–400. doi:[10.1016/j.neuron.2004.09.027](https://doi.org/10.1016/j.neuron.2004.09.027).
32. ^aFogelin R (1985). "The Logic of Deep Disagreements." *Inform Logic*. **7**:1–8.
33. ^a ^bKölbel M (2004). "Faultless Disagreement." *Proc Aristot Soc New Ser*. **104**:53–73.
34. ^a ^bPritchard D (2018). "Wittgensteinian Hinge Epistemology and Deep Disagreement." *Topoi*. **40**(5):1117–1125. doi:[10.1007/s11245-018-9612-y](https://doi.org/10.1007/s11245-018-9612-y).

35. ^{a, b, c}Awad E, Dsouza S, Shariff A, Rahwan I, Bonnefon JF (2020). "Universals and Variations in Moral Decisions Made in 42 Countries by 70,000 Participants." *Proc Natl Acad Sci USA*. **117**(5):2332–2337. doi:[10.1073/pnas.1911517117](https://doi.org/10.1073/pnas.1911517117).
36. ^ΔKamm F (2016). *The Trolley Problem Mysteries*. Tanner Lecture on Human Values 2013. New York: Oxford University Press.
37. ^ΔAlexander L, Moore M (2021). "Deontological Ethics." *The Stanford Encyclopaedia of Philosophy*. <https://plato.stanford.edu/archives/fall2021/entries/ethics-deontological/>.
38. ^ΔNelkin D, McKenzie C, Rickless S, Ryazanov A (2023). "Trolley Problems Reimagined: Sensitivity to Ratio, Risk, and Comparisons." In: Aguiar F, Viciana H, Gaitan A, editors. *Experiments in Moral and Political Philosophy*. Routledge.
39. ^ΔMcMahan J (2009). *Killing in War*. Oxford: Oxford University Press.
40. ^ΔThomson J (1990). *The Realm of Rights*. Cambridge (MA): Harvard University Press.
41. ^ΔHirayama T (1981). "Non-Smoking Wives of Heavy Smokers Have a Higher Risk of Lung Cancer: A Study from Japan." *Br Med J (Clin Res Ed)*. **282**:183e5.
42. ^ΔRanalli C (2018). "What Is Deep Disagreement?" *Topoi*. **40**:983–998.
43. ^ΔAhlenius H, Tännsjö T (2012). "Chinese and Westerners Respond Differently to the Trolley Dilemmas." *J Cogn Cult*. **12**(3–4):195–201. doi:[10.1163/15685373-12342073](https://doi.org/10.1163/15685373-12342073).
44. ^ΔKörner A, Deutsch R (2023). "Deontology and Utilitarianism in Real Life: A Set of Moral Dilemmas Based on Historic Events." *Pers Soc Psychol Bull*. **49**(10):1511–1528. doi:[10.1177/01461672221103058](https://doi.org/10.1177/01461672221103058).

Declarations

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.