

Preprint, penultimate draft. Please cite final draft in *Consciousness and Cognition*

Ignorance and Moral Judgment: Testing the Logical Priority of the Epistemic

Parker Crutchfield

Scott Scheall

Cristal Cardoso Sao Mateus

Hayley Dawn Brown

Mark Rzeszutek

ABSTRACT: It has recently been argued that a person's moral judgments (about both their own and others' actions) are constrained by the nature and extent of their relevant ignorance and, thus, that such judgments are determined in the first instance by the person's epistemic circumstances. It has been argued, in other words, that the epistemic is logically prior to other normative (e.g., ethical, prudential, pecuniary) considerations in human decision-making, that these other normative considerations figure in decision-making only after (logically and temporally) relevant ignorance has constrained the decision-maker's menu of options. If this is right, then a person's moral judgments in some set of circumstances should vary with their knowledge and ignorance of these circumstances. In this study, we test the hypothesis of the logical priority of the epistemic. We describe two experiments in which subjects' knowledge and ignorance of relevant consequences were manipulated. In the second experiment, we also compared the effect of ignorance on moral judgments with that of personal force, a factor previously shown to influence moral judgments. We found broad empirical support for the armchair arguments that epistemic considerations are logically prior to normative considerations.

KEYWORDS: logical priority of the epistemic, epistemic burdens, moral judgment, trolley problem, personal force

1. INTRODUCTION

A philosophical argument (“from the armchair”) has recently been offered for the *logical priority of the epistemic* (Scheall and Crutchfield 2020). According to this thesis, epistemic considerations “pre-consciously shape our incentive structures” (Scheall and Crutchfield 2020, 2). More carefully, the *epistemic burden* of some prospective objective of human action is “simply everything that the actor must know (that and how), which the actor does not already know, in order to realize the objective deliberately as a result of related actions directed to its realization, i.e., not spontaneously, or otherwise in virtue of luck or fortune” (Scheall 2019). The epistemic burden of some goal is the knowledge (including both *knowledge-that* and *knowledge-how*) that an actor still needs to acquire to deliberately realize the goal. According to the logical priority of the epistemic, epistemic burdens function as criteria for the pre-conscious culling of courses of action and the ranking of remaining options in a consciously tractable incentive structure. Persons never consciously consider every conceivable action that might be taken under given circumstances. Rather, the options that a person consciously considers are always some proper subset of the conceivable actions that could be taken in the given decision context. Some courses of action are consciously considered, while others never enter consciousness. What determines whether a particular course of action constitutes an option for a decision-maker and, if so, where it is initially ranked relative to other options in their incentive structure? According to Scheall and Crutchfield (2020), the answer is the decision-maker’s epistemic burdens with regard to contextually relevant courses of action.

Courses of action that seem to bear impossibly heavy epistemic burdens are typically not counted as options in the menu over which an actor deliberates and from which she eventually makes a conscious choice, while courses of action that seem to bear comparatively heavy epistemic burdens are systematically discounted in the actor's menu relative to options that appear less epistemically burdensome (Scheall and Crutchfield 2020, 2).

It is only after (logically and temporally) courses of action have been pre-consciously culled and remaining options ranked in an initial menu of options that other normative considerations (e.g., moral or prudential) enter the conscious decision-making frame and a choice is ultimately made.

Scheall and Crutchfield (2020) offer two philosophical arguments in support of the logical priority of the epistemic.

First, introspection seems to indicate the reality of the phenomena. Introspection seems to reveal that the options a person confronts in any given decision context emerge in consciousness pre-culled and pre-ranked according to the nature and extent of the person's relevant ignorance. Courses of action about which a person is entirely ignorant, i.e., actions that seem to bear impossibly heavy epistemic burdens, simply do not appear as options in the preference ranking that initially emerges in the person's consciousness, while actions about which the person is to some degree ignorant seem to be ranked in this initial conscious preference ranking according to the relative weight of their respective epistemic burdens.

When considering, for example, possibilities for cross-country travel, no human consciously considers *flying like a bird* to be a travel option. The epistemic burden of flying like a bird is so impossibly heavy that it does not count as an option for human beings. On the other

hand, for many humans, the epistemic burdens of *flying on an airplane*, *traveling by train*, *traveling by bus*, and *driving an automobile*, are either non-existent or comparatively easily surmountable. Thus, these modes of travel tend to appear in consciousness as options in many persons' incentive structures, approximately ranked according to their relative epistemic burdens. It is only after a tractable menu of travel options has been pre-consciously determined according to relevant epistemic burdens that the decision-maker consciously applies other normative criteria – concerning, presumably, price, time constraints, convenience, safety, etc. – to decide upon a mode of travel.

Second, Scheall and Crutchfield argue that it is not possible to make sense of the presumed relation between *ought* and *can*, whatever its logical strength (i.e., regardless of whether it is *logical implication*, *presupposing*, *making plausible*, *conversationally implicating*, or some other relation that obtains between *ought* and *can*), unless the epistemic is logically prior to the normative.¹ This is because the word “can” is inherently epistemic.² There is no non-epistemic meaning of “can” – or, at least, none has been proffered in the literature – that does not make the things a person can do a function of the things they know *of* and know *how to do*. Thus,

¹ It is often asserted that *ought implies can*, though other moral philosophers have argued that a weaker relation than *logical implication* obtains between *ought* and *can*. Perhaps *ought* merely *presupposes* or *makes plausible*, or is a *conversational implication of*, *can* (e.g., Hampshire, 1951; Hare, 1951, 1963; Sinnott-Armstrong, 1984; Mizrahi, 2015; Littlejohn, 2009, Vallentyne, 2009; Vogelstein, 2012).

² Recall that, for Scheall and Crutchfield, knowledge encompasses both propositional knowledge-that and non-propositional knowledge-how.

ought implies (or whatever) *knows enough to*.³ In effect, even if there is disagreement about the relation that connects ought and can, everyone who accepts that some relation like those debated in the literature obtains between ought and can is already implicitly committed to the logical priority of the epistemic.

That the epistemic is logically prior to the moral is intuitively plausible. Indeed, as a previous reader of this paper noted, “it is a bit hard for me to imagine any other possibility – how can someone reach a moral judgment about a situation they know nothing about?” Though the thesis might seem obvious, it is nevertheless original – the logical priority of the epistemic has not been explicated in the existing literatures in moral philosophy and moral psychology.⁴ In the

³ Of course, if no relation obtains between ought and can, then this argument loses its force.

However, the argument from introspection would still stand, even if this were the case.

⁴ There is an extensive literature in moral psychology on the significance of *negligence* for moral judgments (see, e.g., Schultz and Wright [1985], Nuñez, Laurent, and Gray [2014], Laurent, Nuñez, and Schweitzer [2016]). However, negligence and ignorance are not cognate notions. Indeed, on most conceptions of negligence, the negligent person *knows* about and is *aware* of (note that *awareness* is an epistemic concept) that which they are negligent, but fails to act in a way supported by this knowledge (Nuñez, Laurent, and Gray [2014]). A negligent person, in other words, fails to use their given knowledge and awareness appropriately. Negligence is *not* typically conceived as a problem of ignorance or inadequate learning. Admittedly, there are alternate conceptions of negligence according to which a negligent person *should have known* about or *should have been aware* of that which they are negligent (Nuñez, Laurent, and Gray [2014]). However, Scheall and Crutchfield (2020) argue that “should have known” claims

present paper, we propose to go beyond Scheall and Crutchfield's armchair philosophical arguments to test the thesis of the logical priority of the epistemic empirically.

2. LOGICAL PRIORITY AND ITS RELATIONSHIP TO THE EMPIRICAL EVIDENCE⁵

To say that one set of considerations is *logically prior* in decision-making to another set of considerations is to say that that the second set can play no role in decision-making unless the first set has played its role. Logical priority is a stronger form of priority than others that might also be relevant to decision-making, such as *psychological priority*. If some set of considerations is logically prior in decision-making to some other set of considerations, it is also psychologically prior. But the opposite need not be true. Moreover, phenomena that are logically prior in decision-making to some others are necessarily so. But phenomena that are merely psychologically prior in decision-making to some others are only contingently so.

Imagine that the evolutionary processes to which the psychologies of decision-makers were subjected were such that courses of action were sorted into a tractable menu of options

implicitly presuppose a "could have known" assumption and, therefore, are not counterexamples to the logical priority of the epistemic. This argument is easily testable: if subjects who assert "should have known" claims tend to revoke such assertions when they are shown that the relevant actor could not have acquired the relevant knowledge, then Scheall and Crutchfield's argument that "should have known" claims imply an unspoken "could have known" assumption would seem established.

⁵ Many thanks to the anonymous referee who encouraged us to elaborate on this aspect of the argument

according to the nature and extent of the sacrifice required to successfully pursue them, rather than by epistemic criteria. Such opportunity-cost considerations would be merely psychologically prior to other considerations – say, the moral worth or prudential value of courses of action – that could have played the same part in decision-making had the contingent processes of evolution been different than they (*ex hypothesi*) were.

Yet, how could decision-makers subject to such contingent evolutionary processes choose between options unless they *knew*, and were *able to evaluate and compare*, the respective opportunity costs of different courses of action? Decision-makers' *knowledge* must come first, in order for non-epistemic criteria to play their assumed role in culling and sorting courses of action into a tractable menu of options. Epistemic considerations are thus logically prior, not merely psychologically prior, in decision-making to non-epistemic considerations.

This being said, empirical evidence *in isolation* would seem impotent to distinguish between the logical priority and the merely psychological priority of epistemic considerations in decision-making. The evidence alone cannot say whether the priority of the epistemic is a necessary or merely a contingent aspect of decision-making. However, Scheall and Crutchfield's philosophical arguments, plus the foregoing argument, in *conjunction* with the empirical evidence discussed below strongly suggest the veracity of the logically stronger thesis.

3. THE EPISTEMIC BURDEN OF MAKING A MORAL JUDGMENT

Above we note that epistemic considerations are logically prior to other normative considerations. If this thesis is true, we would expect that epistemic considerations influence moral decision-making. This predicted influence is why we chose to test the thesis of the logical priority of the epistemic using moral judgments.

Whatever a particular person is predisposed to judge moral, the goal of *making a moral judgment*, i.e., the goal of making a judgment that reflects the person's moral predisposition, can be more or less epistemically burdensome. A person might possess all, some, or none of the knowledge required in some decision context to make a judgment consistent with their underlying moral predisposition. The goal of judging in a way that reflects the decision-maker's moral predisposition is easily deliberately realized when they possess all relevant knowledge. Absent countervailing evidence, we should expect the moral judgments a person makes when relevantly knowledgeable to reflect their moral predisposition. The goal is less easily deliberately realized when the decision-maker lacks some relevant knowledge. We should expect a disconnect to manifest between the judgments a person makes under conditions of partial ignorance and their underlying moral predisposition (as reflected in their moral judgments when relevantly knowledgeable). The goal of judging in line with one's moral predisposition is impossible to deliberately realize when the decision-maker is entirely ignorant of relevant knowledge (though it might nevertheless be non-deliberately or spontaneously realized, e.g., as a matter of luck).

Consequentialism defines a morally appropriate action to be one that engenders positive results (for some sentient beings or other), regardless of any moral rules or duties associated with the action. If the thesis of the logical priority of the epistemic is sound, then whether a person predisposed to consequentialism makes a consequentialist judgment in some decision context depends on the epistemic burden of making a consequentialist judgment in that context. The person predisposed to consequentialism will be able to easily and deliberately make a consequentialist judgment only when they possess all of the knowledge that such a judgment requires. The knowledge required to make a consequentialist judgment would seem to concern,

first and foremost, the different consequences to follow from various courses of action. A person predisposed to consequentialism, confronting a choice between options A and B, can easily and deliberately choose the option that maximizes net consequences only if they know, and can compare, the consequences of choosing A and the consequences of choosing B. Otherwise, the heavier their relevant epistemic burdens – the more ignorant they are of (or the less capable they are of comparing) relevant consequences – the more such a person will struggle to deliberately make a consequentialist judgment. We should then expect a greater disconnect between their judgments in epistemically burdensome contexts and their underlying consequentialist predisposition, as indicated by their judgments when relevantly knowledgeable.

Conversely, *deontological* moral theories make a morally appropriate action one that conforms to a moral rule or duty, such as the duty to *show respect for persons* or to *not harm others*, regardless of the consequences of the action. If the logical priority of the epistemic is sound, then whether a person predisposed to deontology makes a deontological judgment in some decision context depends on the epistemic burden of making a deontological judgment in that context. The person predisposed to deontology will be able to easily and deliberately make a deontological judgment only when they possess all of the knowledge that such a judgment requires. The knowledge required to make a deontological judgment would seem to concern, first and foremost, the moral rules and duties relevant in some decision context and the relevance to, or applicability in, this context of these moral rules and duties. A person predisposed to deontology, confronting a choice between options A and B, can easily and deliberately choose the option most in keeping with relevant moral rules and duties, only if they know relevant moral rules and duties, and their relevance to the given decision context. Otherwise, the more ignorant they are of relevant moral rules and duties, or of the applicability of these rules and duties in the

given decision context, the more such a person will struggle to deliberately make a deontological judgment. We should then expect a greater disconnect in epistemically burdensome contexts between their judgments and their underlying deontological predisposition, as indicated by their judgments when relevantly knowledgeable.

We hypothesize that knowledge relevant in some decision context to the goal of making a moral judgment of a particular (consequentialist or deontological) kind in that context functions as guardrails keeping a persons' moral judgments in line with their moral predisposition. When this relevant knowledge is removed, they struggle to make the kind of judgment they are predisposed to make, one they would easily make in more favorable epistemic circumstances. In effect, in the presence of relevant ignorance, all moral bets are off.

If decision-making proceeds from epistemic to other normative considerations, then a person's transient moral judgments in some set of circumstances should vary with their knowledge and ignorance of these circumstances.⁶ The same person, considering the same

⁶ There has been much work on what one should do when one is uncertain about normative facts. This work is tangentially related to the present work. We are highlighting the descriptive facts pertaining to the influence ignorance has upon the determination and subjective ranking of options. The work on normative uncertainty pertains to what one should do when one is uncertain about right- and wrong-making properties of an act. That one is normatively uncertain may also affect where an option appears in a person's preference ranking, but that it does so is simply an instance of the logical priority of the epistemic at work. In other words, we are presently concerned with describing how people will decide, not determining how they ought to decide, though our descriptive project and others' normative project are related.

circumstances, should make different judgments, depending on what they know about these circumstances. However, precisely because “all moral bets are off” in the presence of relevant ignorance, it is difficult to predict – and we do not offer a model of – exactly *how*, i.e., in what direction, judgments will shift when subjects are relevantly ignorant. We see no *a priori* reason to expect judgments to shift in any particular direction when subjects are made ignorant.

There is an extensive literature, perhaps most closely associated with Daniel Kahneman, Amos Tversky, and the research program of behavioral economics, concerning the various heuristics and biases, rules of thumb, guesses, and life experiences that people rely upon in decision-making (see, e.g., Kahneman and Tversky 1979, Kahneman 2011). In this sense, the results of the present paper are consistent with the behavioral economics research program. In another sense, however, in the sense in which behavioral economics is frequently interpreted – perhaps misinterpreted (see Kahneman, 2011, p. 411) – to evaluate real-world human decision-making and to find it lacking according to the normative criteria implied by expected utility theory (EUT), the present paper is in tension with this common interpretation of behavioral economics. The logical priority of the epistemic implies that decision-making is *difficult* and always reflective of the decision-maker’s often deficient epistemic circumstances. Far from finding decision-makers *at fault* when their decisions fail to meet the high standards of EUT, the logical priority of the epistemic implies that decisions in line with EUT should be expected only in the rarest of epistemic conditions: more or less complete omniscience of decision-relevant circumstances. If this is right, then decision-makers need more and better knowledge, rather than “nudges” (Thaler and Sunstein 2009), which are unlikely to be effective (Maier et al., 2021), in any case, if the knowledge required for their effectiveness is lacking.

Moreover, heuristics, biases, guesses, and rules of thumb may be especially influential in the absence of relevant knowledge (Gurevich, 2019). Indeed, ignorance may enable these aspects of decision-making. This is not merely consistent with the logical priority of the epistemic, it is predicted by it. If relevant knowledge fixes one's options and how they are ranked, lacking this knowledge in a context in which the decision-maker must nevertheless make a choice may enable these cognitive shortcuts.

4. EPISTEMIC BURDENS VS. PERSONAL FORCE

If epistemic considerations are logically prior to other moral considerations in moral decision-making, we should expect epistemic considerations to influence moral decision-making. It is thus useful to investigate these effects, especially in the context of other factors that purportedly influence moral decision-making. This is not to say the thesis of the logical priority of the epistemic is in contention or competition with other theses. But other authors have shown that other non-moral properties influence moral decision-making. Part of our aim in the present paper is to investigate the effect of yet another non-moral property – ignorance – on moral decision-making, and to compare that effect with that of other non-moral properties.

One significant thread to emerge from research on the factors that influence moral judgment is the notion that deontological judgments appear to be more “emotional” or “personal,” and consequentialist judgments more “cognitive” (Haidt, 2001). Greene et al. (2001, 2008a, 2008b, 2009) argue that moral judgment is affected by factors that elicit emotional responses. Furthermore, Cushman et al. (2006) have shown that, in circumstances such as those described in the well-known Trolley Problem, *physical distance* affects whether one makes a deontological or a consequentialist judgment. Indeed, Greene et al. (2009) suggest that many of

these other factors such as physical distance and intention matter only to the extent that they involve the use of personal force. When faced with a moral dilemma, people are more likely to make a deontological judgment when the consequentialist alternative promotes better consequences but requires using personal force against another person. The need to use personal force is another factor that can dislodge a person's momentary moral judgments from their underlying moral predisposition, from the judgments they are predisposed to make when personal force is not required. And since there is evidence that other factors reduce to personal force, when investigating the roots of momentary moral decisions, it is useful to compare the effects of ignorance with other potential fundamental factors.

So, on one hand, there is evidence that personal force and other factors that elicit emotional responses affect a person's moral judgments. On the other hand, a plausible argument can be made from the armchair that epistemic considerations are logically fundamental in human decision-making, including in the moral judgments that persons make about their own and others' actions. Although moral psychologists have extensively investigated the significance of emotion-evoking factors for moral judgment, the effects of relevant ignorance have heretofore gone unexamined. The present research explores whether knowledge and ignorance affect moral judgment, and, if so, how the effects of ignorance on moral judgment interact with those of personal force. In particular, we are interested in whether the effects of ignorance are more fundamental for moral judgment than the effects of emotion-eliciting factors like personal force.

Like much of the research just discussed, we observed participants' responses to the Trolley Problem.⁷ In this earlier research, however, experimental subjects were always

⁷ It has recently been argued that such “sacrificial moral dilemmas” as one encounters in the standard Trolley Problem (and its many variants) are inadequate to the measurement of a subject's inclination toward utilitarian (i.e., consequentialist) rather than deontological moral judgments (Kahane [2015]). A person's reactions to the Trolley Problem capture only a part, and perhaps not the most important part, of the psychology of utilitarian judgments (Kahane et al [2015], Kahane et al [2018]). It is important to emphasize that our sole concern in the present paper is with testing the thesis of the logical priority of the epistemic, i.e., with uncovering any evidence for (or against) the notion that a person's transient moral judgments depend on their relevant knowledge and ignorance, and are likely to change when their relevant knowledge and ignorance changes. Given this goal, a simple Trolley Problem setup in which subjects' relevant knowledge is easily manipulated seems appropriate. Were we uniquely concerned with the psychology of utilitarian (or, for that matter, deontological) judgments, a more sophisticated measure of subjects' preference for one kind of judgment rather than the other, such as the two-dimensional scale developed and validated by Kahane et al (2018) would be appropriate. However, if it is sound, then the thesis of the logical priority of the epistemic applies to utilitarians, other (non-utilitarian) consequentialists, deontologists of all stripes, traditional religious moralists, libertarians, and virtue ethicists. In other words, it does not matter for our purposes whether Trolley Problem experiments capture none, some, or all of the psychological process underlying utilitarian judgments. All that matters is whether such experiments can provide evidence concerning the effects of relevant ignorance on momentary moral judgments.

knowledgeable about the consequences of the different options they confronted. If Scheall and Crutchfield (2020) are right, altering subjects' epistemic circumstances – in particular, given the constraints of the Trolley Problem, what they know about the consequences of the available options – should affect their momentary moral judgments. Thus, we varied research subjects' epistemic circumstances, their ignorance of the consequences of different courses of action, and observed their judgments.

The present study includes two experiments. The first investigates whether ignorance of relevant consequences affects moral judgment. The second study investigates the interaction of ignorance and personal force to determine whether the influence of one or the other on moral judgment is more fundamental.

5. METHODS

4.1 General

Both experiments were approved by the institutional review board at [redacted for blind review]. Participants were recruited through Amazon's Mechanical Turk (MTurk). MTurk has been validated as a tool for quality data collection (Buhrmester et al., 2011; Casler et al., 2013; Holden et al., 2013). The study posting was only visible to participants who resided in the USA, had completed 100 MTurk tasks, and had at least a 95% approval rating. Participants completed an electronic consent form and were directed to a Qualtrics survey. The survey began with a comprehension question. Only participants who answered the comprehension question correctly continued to complete the survey. Participants were compensated \$2.00.

Upon correctly responding to the comprehension question (see below), participants were presented with several trolley scenarios. For each scenario, participants could choose whether to do something to divert the train from its current path or to do nothing to divert the train. The questions were timed and participants informed of this. If a participant did not make a choice before time ran out, they were coded as having chosen to do nothing. After completing the trolley scenarios, participants completed a short demographic questionnaire. To access the study, participants were presented with a scenario in which an apple was on one track and a banana was on the other. They were informed that pressing the button would make the train switch tracks causing the apple to be smushed and that doing nothing would cause the banana to be smushed. Participants were then asked to smush the apple. If they responded with choosing to do nothing, they were excluded from the study. Because both experiments were relatively short (approximately 15 questions), no attention check was included.

All statistical analysis was conducted using R 4.03 (R Core Team, 2020). Because of the use of a repeated measures design in both experiments, mixed-effects regressions were used with participants as the random factor. The lme4 package (Bates et al., 2015) was used for all mixed-effects regressions. For questions requiring a binary response (e.g., push/do not push), logistic mixed-effects regressions were used. For questions that were on a scale or using a continuous variable, linear mixed-effects regressions were used. Homogeneity of variance was assessed using Levene's test for responses to questions assessed with a continuous variable (i.e., acceptability of action). Wald tests were conducted on all mixed-effects regressions using the car package (Fox & Weisberg, 2019) using type-II sums of squares. Post-hoc comparisons were conducted using the emmeans package (Lenth, 2020). Multiple comparisons were corrected for using the Holm p-value adjustment method. In all cases, type-I error rate was set at .05. Given

the novelty and the nature of the phenomena under investigation, simulations were conducted to determine an adequate sample size to determine effect. The sample size was therefore calculated based on sample sizes common in the literature. A sample size of 175 would be near the median for those studies that were statistically significant. We aimed for 175-200 subjects to complete the survey following exclusion from the comprehension question. Simulated datasets were created to determine expected power for low to medium sized effect sizes. Simulated data of samples with 175 subjects and what are considered between small to medium effects sizes for odds ratios, (i.e., 1.72, 1.93, and 2.18; see Chen et al., 2010) resulted in $\beta = .591, .768, \text{ and } .894$ respectively while $\alpha = .05$.

4.2 First Experiment

Two hundred fifty-four participants started the survey. After excluding participants who incorrectly answered the comprehension question, a final sample of 193 participants was obtained.

In addition to the comprehension question, participants were presented with a total of eight trolley scenarios. These scenarios were split among five variations in which ignorance of the consequences of either switching tracks or staying on the present track was manipulated. For scenarios that included puzzles that participants had to solve to learn the missing knowledge, participants were first asked whether they wanted to solve the puzzle or do nothing. If they opted to solve the puzzle, they were directed to it and asked to enter the correct solution to switch tracks or to submit “NA” to do nothing. Participants were supplied with written descriptions and

visual depictions of each variation. Brief descriptions of the variations are provided below and the full scenarios can be found in Appendix B.

The Baseline scenario read as follows:

You are standing at the point where one train track splits into two tracks. A runaway train is approaching. On the right track are five rail workers. The brakes have stopped working, so the train can't stop. Next to you is a button that switches tracks.

If you do nothing, the train will hit the five rail workers, who will die.

However, if you press the button, the train will switch to the left track. On the left track is one rail worker. The train will hit the rail worker, who will die, but the train will stop and the five rail workers on the right track will live.

What do you do?

Participants had the option to press the button or do nothing.

The other scenarios manipulated participants' epistemic burdens with respect to some or all of the consequences of switching to the left track, or staying on the right track. In some scenarios, these burdens were insurmountable: participants could not remedy their ignorance. In other scenarios, participants could surmount their epistemic burdens by solving a puzzle that would provide the missing knowledge.

In Complete Ignorance, participants were told that there may or may not be people on either or both of the tracks. They were not told whether and, if so, how many, people were on each track. In Partial Ignorance Both Tracks, participants were told that there were people on both tracks, but were left ignorant as to how many. There was no way for participants to learn the missing knowledge in either scenario. The order of the two scenarios was randomized across participants.

In Partial Ignorance Switch Track, participants knew the number of people on the stay (right) track, but were ignorant of whether and, if so, how many, people were on the switch (left) track. They had no way of learning this information. In Partial Ignorance Stay Track, participants knew the number of people on the switch track, but not whether and, if so, how many, people were on the stay track. Again, participants could not learn this missing knowledge. The order of the two scenarios was randomized across participants

In Surmountable Ignorance Switch Track, participants knew the number of people on the stay track, but were initially ignorant of whether and, if so, how many, people were on the switch track. However, they had the option of solving a puzzle to learn how many people were present on the switch track. In Surmountable Ignorance Stay Track, participants knew the number of people on the switch track, but were initially ignorant of whether and, if so, how many, people were on the stay track. However, they had the option of solving a puzzle to learn how many people were present on the stay track. The puzzle was drawn from a sample SAT question. The order of the two scenarios was randomized across participants.

For the Know-How Ignorance scenario, participants were presented with the Baseline scenario, but told that the button that controls the rail switch was not functioning and that, if they

wanted to switch tracks, they would have to solve a logical reasoning puzzle in order to repair it. The logical reasoning puzzle was drawn from a sample law school admissions test (LSAT).

Participants were presented with the scenarios in the following order: Surmountable Ignorance, then Baseline, then, in a randomized order across participants, Complete Ignorance, Partial Ignorance, and Know-How Ignorance. Participants had 45 seconds to choose in scenarios without puzzles and 90 seconds to choose in scenarios with puzzles.⁸

4.3 Data Organization and Analysis

Due to the several ways a participant could respond to different questions, responses were simplified to whether a participant chose to switch tracks or to stay on the same track (i.e., switch or stay). For purposes of analysis, a switch was coded as 1 while a stay was coded as a 0. Therefore, all proportions and estimates in a positive direction indicate scenarios where participants were more likely to switch, while those in a negative direction indicate scenarios where participants were more likely to stay. Whether a participant solved a puzzle or if their answer was correct was not considered for purposes of analysis. To establish an individual's "moral baseline" (or, as above, their "moral predisposition"), those who responded to the Baseline scenario by switching were classified as "consequentialists." Those who did not switch

⁸ We imposed the time constraints to more accurately simulate the pressure of time, were one actually deciding whether to switch tracks with a train barreling toward the rail spur. The time constraint may interact with a subject's decision, but the absence of such a constraint may also have interacted.

were classified as “deontologists.”⁹ In the regression analysis, participants were treated as random effects while each question was treated as a fixed effect. Because of issues of convergence when assessing interaction between questions and moral baseline, the traditional trolley problem is not included in the regression analyses that include moral baseline.

4.4 Second Experiment

Participants who completed the first study were excluded from the second. Two-hundred twenty-two participants began the survey. After excluding those who answered the comprehension question incorrectly, 178 participants were included in the final sample for analysis.

⁹ The interpretation of a track-switcher as having made a consequentialist choice and of a track-stayer as having made a deontological choice was maintained across all experimental variations. That is, regardless of the specific circumstances of a particular scenario, a participant who chose to switch tracks chose to instrumentally harm any persons on the switch track and, thus, could not be interpreted to have made a deontological judgment under the given circumstances. At the same time, a participant who refused to switch tracks chose to not be a party to instrumental harm, regardless of consequences (about some or all of which they were ignorant, remediable or otherwise, in different scenarios), and, thus, could not be interpreted to have made a consequentialist judgment under the given circumstances.

This study included variations of the footbridge trolley problem in which ignorance (whether the participant knew if there were people and, if so, how many, on one or more of the tracks, and / or whether participants knew if their actions would successfully stop the train before striking any persons) and personal force (whether participants were asked to choose either pushing a button or pushing a person to stop the train) were manipulated. Participants were provided with eight different scenarios in addition to the comprehension question. Two scenarios were presented, randomized across participants, for each variation on the Trolley Problem (excluding Comprehension). In one scenario, participants were asked to push a button that would release a scaffolding, causing a rail worker to fall on to the track, and, in the other scenario, to push a rail worker on to the track. Brief descriptions of the variations are provided below and the full scenarios can be found in Appendix C.

In Baseline Button, participants knew the consequences of both options. In order to stop the train, they had to push a button that would cause a rail worker to fall onto the tracks, stopping the train. In Baseline Push, participants knew the consequences of both options. In order to stop the train, they had to choose to push a rail worker off the footbridge onto the tracks, stopping the train.

Complete Ignorance Button reads as follows (Complete Ignorance Push differs only in that the participant must choose to push the worker off the footbridge).

Complete Ignorance Button: An empty runaway trolley is speeding down a set of tracks toward an unknown number of rail workers. There is a footbridge above the tracks in between the runaway trolley and the workers on the tracks. You are standing on this footbridge. Standing on some scaffolding attached to the bridge is another rail worker and

his heavy equipment.¹⁰ If you do nothing, the trolley will proceed down the tracks and hit any workmen on the tracks, who will die.

However, it may be possible to avoid the deaths of any workers that are on the tracks. The switch to release the scaffolding that the other rail worker is standing on is right next to you. If you push the switch, the rail worker and his equipment will fall onto the tracks. Though this worker will certainly die, there is some likelihood that the combined weight of the worker's body and his equipment will be heavy enough to stop the trolley, but you don't know what this likelihood is. The probability that pushing the switch stops the train is more than 0%, but less than 100%.

Do you push the switch?

In Partial Ignorance Track Workers-Button, participants were presented with the scenario above, except that they were informed that, if they pressed the button, the rail worker's body and heavy equipment would certainly stop the trolley upon falling onto the tracks. Similarly, in Partial Ignorance Track Workers-Push, participants were informed that, if they pushed the worker from the footbridge, the worker's body and equipment would certainly stop the trolley. In Partial Ignorance Footbridge Worker-Button, participants were presented with the scenario above, except that they were told the number of workers on the track while remaining ignorant

¹⁰ The inclusion of heavy equipment in the scenario is supposed to make it more plausible that that the body (plus equipment) is of sufficient mass to stop a runaway trolley.

whether, if they pressed the button, the worker's body and equipment would stop the trolley. In Partial Ignorance Footbridge Worker-Push, participants were presented with the scenario above, except that they were told the number of workers on the track while remaining ignorant whether, if they pushed the worker from the footbridge, the worker's body and heavy equipment would stop the trolley.

In the Comprehension test participants were presented with a scenario in which they were told an apple was on the tracks and would be smushed by the trolley and that there was a banana on the scaffolding, which if pushed off would be smushed along with the apple. Participants were then asked to smush both the banana and apple.

Participants were presented with scenarios in the following order: Complete Ignorance (Button/Push), then, in a randomized order across participants, Partial Ignorance Track Workers (Button/Push) and Partial Ignorance Footbridge Worker (Button/Push), then Baseline. Following each scenario, participants were asked to indicate how morally acceptable it is to push the button/worker on a scale from 1 (not at all acceptable) to 7 (entirely acceptable). There was a 30-second time limit for making selections in the scenarios and no time limit for indicating acceptability.

4.5 Data Organization and Analysis

Participants were assigned a moral baseline much like in the first experiment, using Baseline Button. Coding of responses was similar, with pushing the worker/pressing the button

being coded as 1, while not pushing the worker/pressing the button was coded as 0.¹¹ Regression analysis for binary response options was the same as in the first experiment, with the moral baseline question excluded as part of the analysis. For acceptability, the moral baseline question was included in the regression analysis because there were no issues of convergence when assessing interactions. Because there were two factors, force and knowledge, they were separable for purposes of analysis. Therefore, this experiment was analyzed as a 2 (force) x 4 (knowledge) design.

6. RESULTS

5.1 First Experiment

The final sample for analysis totaled 176 participants. Of those participants, 151 were classified as predisposed to consequentialism and 25 were classified as predisposed to deontology. The majority of participants were White (81%) and college educated (78%). A little more than half of participants identified as male (63%), and the average participant age was 37.14 (SD=10.82). See Supplemental Table 14 for a breakdown of demographic variables across classification.

¹¹ Interpretation of these choices was much the same as the first study. Pushing the button/worker constitutes instrumental harm and cannot be considered a deontological choice, while not pushing the button/worker means avoiding committing instrumental harm and cannot be assimilated to a consequentialist decision.

Using a simple model not incorporating moral baseline, the effects of ignorance on moral decision-making were significant $\chi^2(7) = 173.31, p < .0001$. Post-hoc comparisons indicated that, with the exception of responses to Partial Ignorance Switch Track, responses to all questions were significantly different from baseline (see Supplemental Table 1 for all post-hoc comparisons based on the simple model). Estimated lever pull proportions (i.e., track switches) can be found in Figure 1.

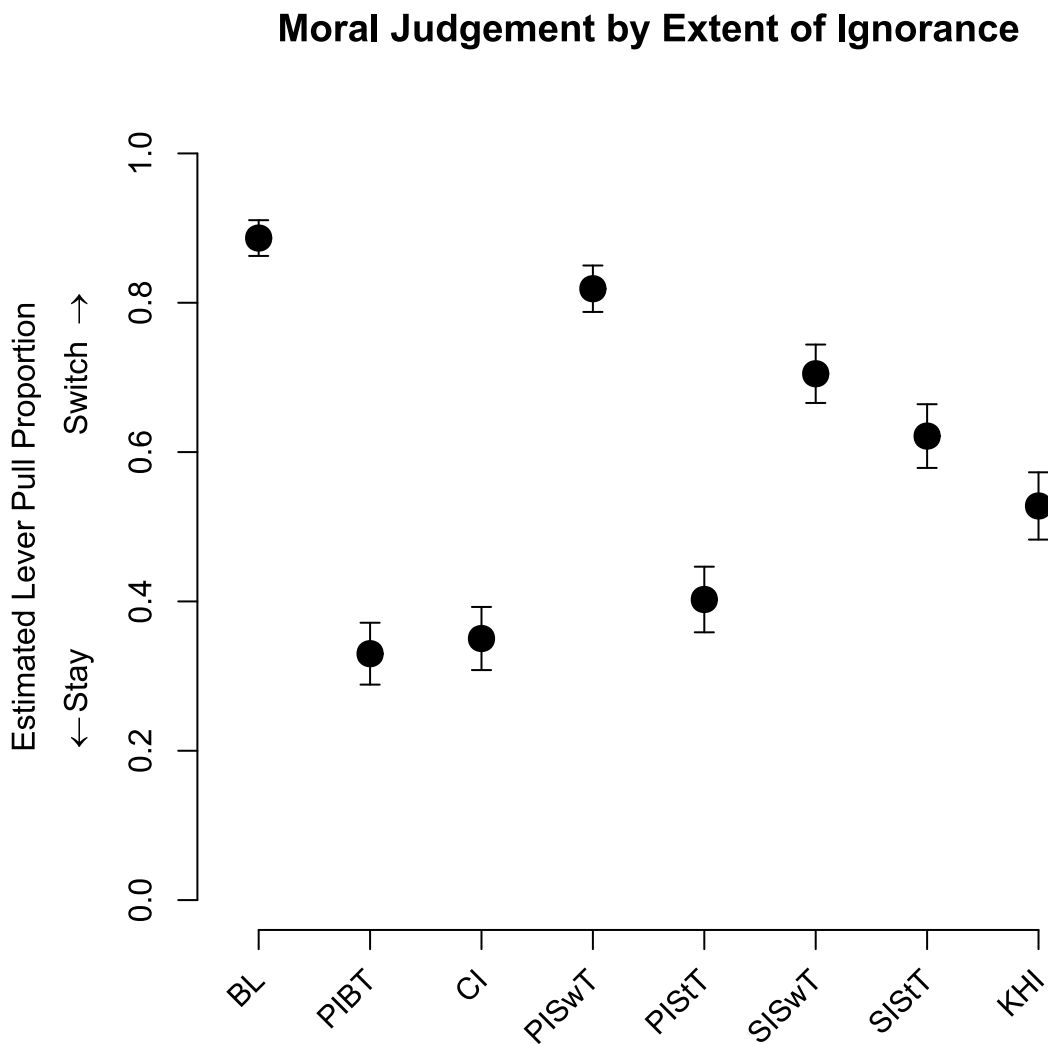


Figure 1. Plot of estimated proportions of those who pulled the lever (i.e., switched the track of the trolley) in Experiment 1. Error bars indicated standard error of the estimate. BL: Baseline trolley problem. PIBT: Partial Ignorance Both Tracks. CI: Complete Ignorance. PISwT: Partial Ignorance Switch Track. PIStT: Partial Ignorance Stay Track. SISwT: Surmountable Ignorance Switch Track. SIStT: Surmountable Ignorance Stay Track. KHI: Know-How Ignorance.

When moral baseline was incorporated, there were significant differences in moral judgment based on ignorance involved, $\chi^2(6) = 103.276$, $p < .0001$, moral baseline, $\chi^2(1) = 15.867$, $p < .0001$, as well as a significant interaction between ignorance and moral baseline, $\chi^2(6) = 26.$, $p = .0002$. Complete post-hoc comparisons for differences based on the nature and extent of ignorance and moral baseline are available in Supplemental Tables 2 and 3. The estimated lever pull proportions can be found in Figure 2.

Moral Judgement by Nature and Extent of Ignorance

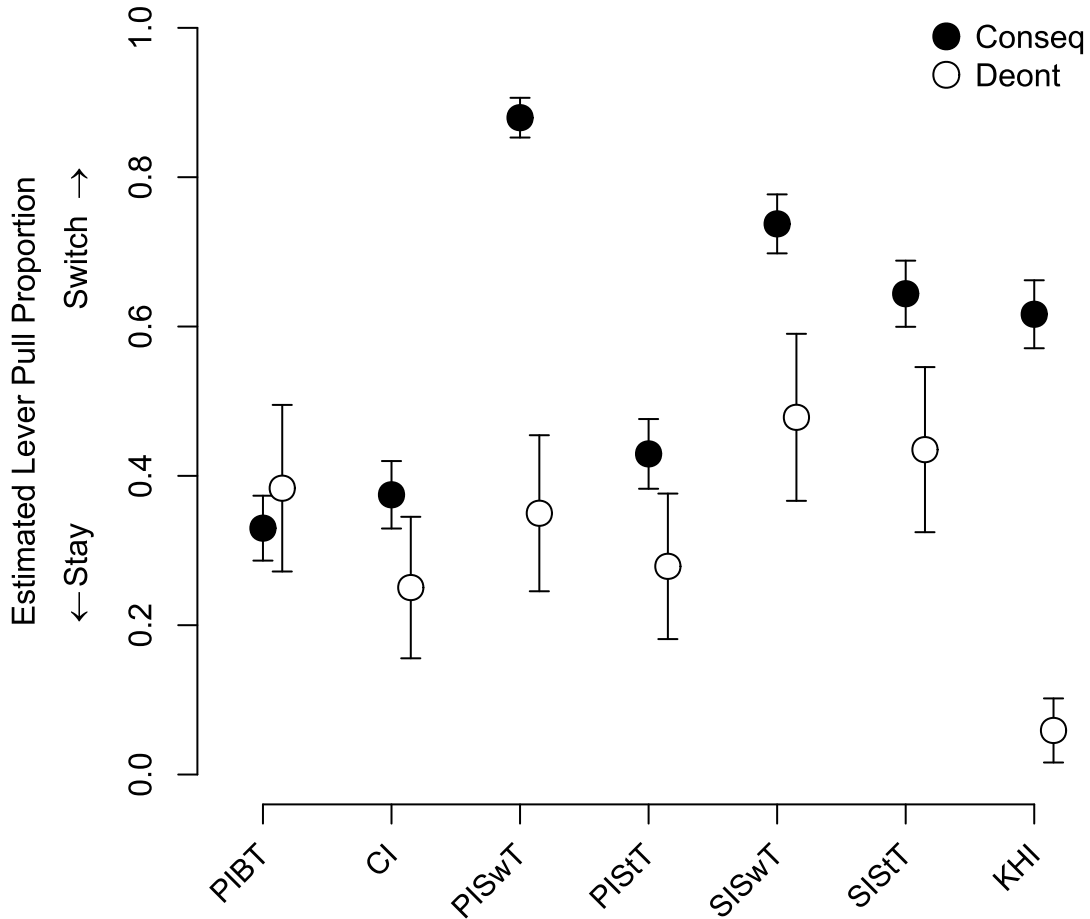


Figure 2. Plot of estimated proportions of those who pulled the lever (i.e., switched the track of the trolley) in Experiment 1 by moral baseline. Error bars indicate standard error of the estimate.

Conseq: Participants identified in Baseline as predisposed to consequentialism. Deont:

Participants identified in Baseline as predisposed to deontology. BL: Baseline trolley problem.

PIBT: Partial Ignorance Both Tracks. CI: Complete Ignorance. PISwT: Partial Ignorance Switch

Track. PIStT: Partial Ignorance Stay Track. SISwT: Surmountable Ignorance Switch Track.

SIStT: Surmountable Ignorance Stay Track. KHI: Know-How Ignorance.

The scenarios in which there was a significant difference between consequentialists and deontologists were Partial Ignorance Switch Track, Surmountable Ignorance Switch Track, and Know-How Ignorance. Generally, those predisposed to consequentialism were more influenced by the nature and extent of ignorance assumed in a scenario, while the proportion of those inclined to deontology who chose to switch tracks under conditions of ignorance remained constant, except in the case of Know-How Ignorance.

5.2 Second Experiment

The final sample for analysis totaled 178 participants. Of those participants, 121 were classified as predisposed to consequentialism and 57 were classified as predisposed to deontology. The majority of participants were White (58%) or Black or African American (33%) and college educated (89%). More than half of participants identified as male (70%), and the average participant age was 35.97 (SD=9.77). See Supplemental Table 15 for a breakdown of demographic variables across classification.

Using a simple model (i.e., without moral baseline) to assess participant moral judgment, both factors of force, $\chi^2(1) = 15.212$, $p < .0001$, and ignorance, $\chi^2(3) = 24.247$, $p < .0001$, were significant. However, there was no interaction detected between the two, $\chi^2(3) = 1.026$, $p = .795$. This model included Baseline Button, the scenario used to categorize a person as predisposed to either consequentialism or deontology. Estimated pull lever proportions can be found in Figure 3.

Moral Judgement by Extent of Ignorance

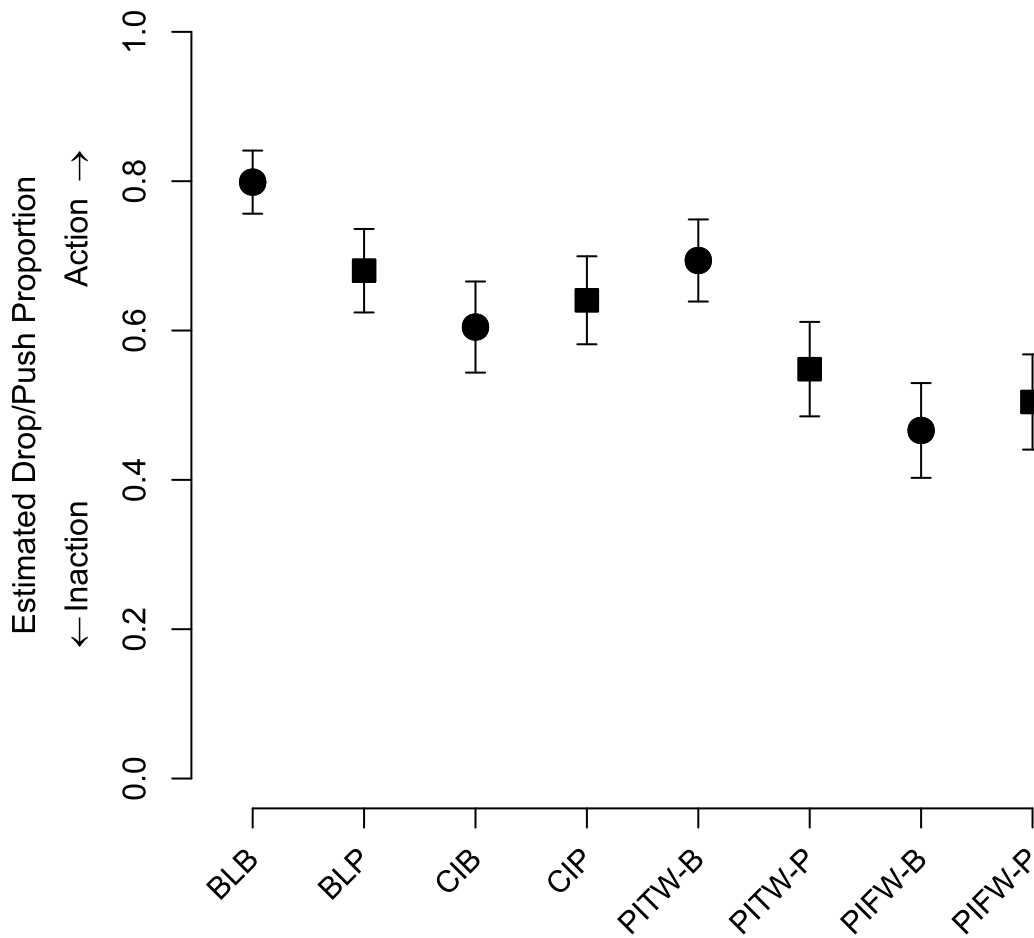


Figure 3. Proportion of participants who chose to push or drop the worker in Experiment 2. Circles represent when a decision did not involve personal force (button press), and squares indicate questions when personal force was involved (push worker). Error bars indicate standard error of the estimate. BLB: Baseline Button. BLP: Baseline Push. PIFW-B: Partial Ignorance Footbridge Worker-Button. PIFW-P: Partial Ignorance Footbridge Worker-Push. PITW-B: Partial Ignorance Track Workers-Button. PITW-P: Partial Ignorance Track Workers-Push. CIB: Complete Ignorance Button. CIP: Complete Ignorance Push.

Using a model that included moral baseline, force, $\chi^2(1) = 10.561$, $p = .0012$, ignorance, $\chi^2(3) = 11.670$, $p = .0085$, and moral baseline, $\chi^2(1) = 50.762$, $p < .0001$ were significant factors in moral judgment. Interactions between ignorance and force remained non-significant, $\chi^2(2) = 0.788$, $p = .675$, however there was a significant interaction between ignorance and moral baseline, $\chi^2(3) = 8.223$, $p = .042$. There was no interaction between force and moral baseline, $\chi^2(1) = 2.167$, $p = .141$, or between all factors, $\chi^2(2) = 0.723$, $p = .695$. Figure 4 shows the moral judgment based on moral baseline, for each trolley scenario. For post-hoc comparisons that could be estimated, there were no significant differences in moral judgment by baseline deontologists for force or ignorance (see Supplemental Tables 4 & 5). When baseline consequentialists were not ignorant, they were more likely to push the worker off the footbridge (see Supplemental Tables 6–8).

Moral Judgement by Nature and Extent of Ignorance

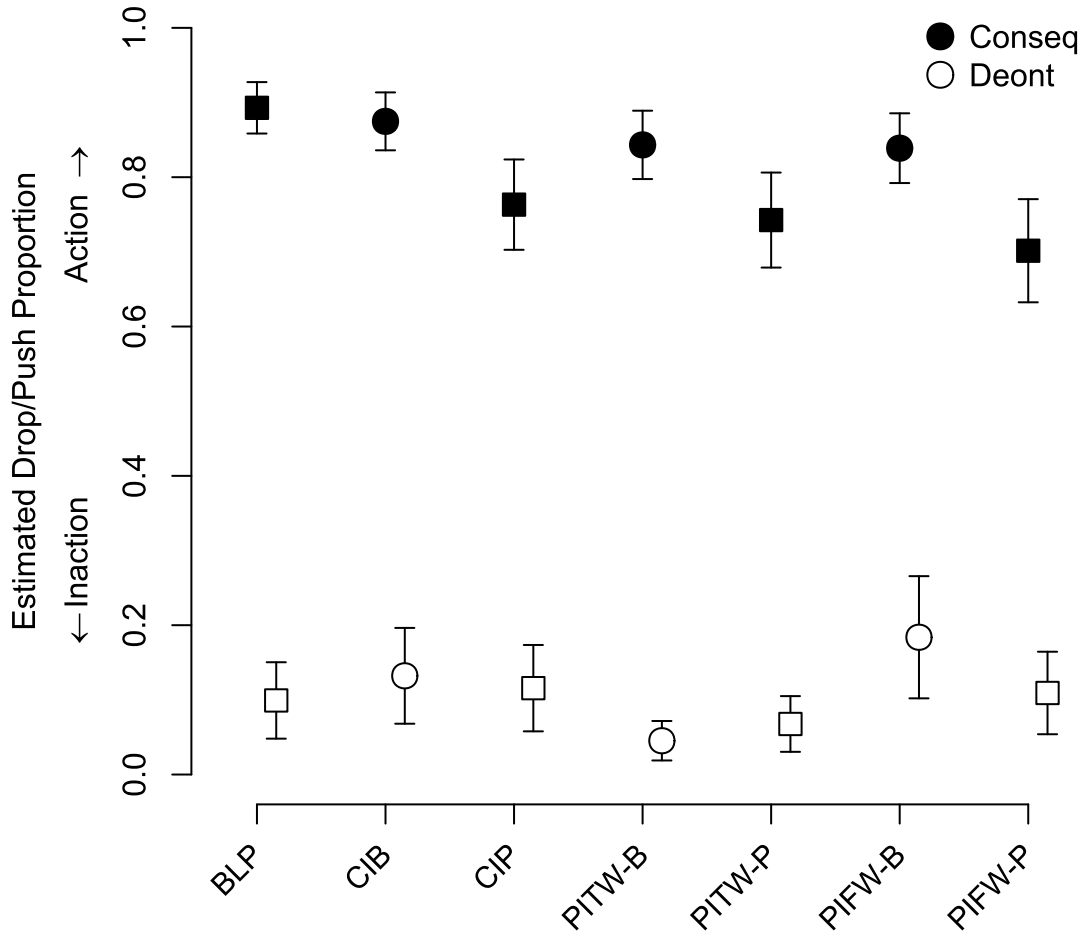


Figure 4. Proportion of participants who chose to push or drop the worker in Experiment 2 separated by moral baseline. Circles represent when a decision did not involve personal force (button press), and squares indicate questions when personal force was involved (push worker). Error bars indicate standard error of the estimate. Conseq: baseline consequentialist. Deont: baseline deontologist. Note that all differences between consequentialists and deontologists are significant. BLB: Baseline Button. BLP: Baseline Push. PIFW-B: Partial Ignorance Footbridge Worker-Button. PIFW-P: Partial Ignorance Footbridge Worker-Push. PITW-B: Partial Ignorance

Track Workers-Button. PITW-P: Partial Ignorance Track Workers-Push. CIB: Complete Ignorance Button. CIP: Complete Ignorance Push. Note that all differences between consequentialists and deontologists are significant.

For acceptability, a simple model comparing ignorance, force, and their interaction was assessed. Both force, $\chi^2(1) = 11.283$, $p = .0008$, and ignorance, $\chi^2(3) = 33.114$, $p < .0001$, were significant factors for judgments of acceptability, but the interaction between the two was not significant, $\chi^2(3) = 0.384$, $p = .9435$ (see Supplemental Tables 9 & 10 for all contrasts).

Estimated acceptability for actions from this model can be found in Figure 5.

Acceptability of Action by Extent of Ignorance

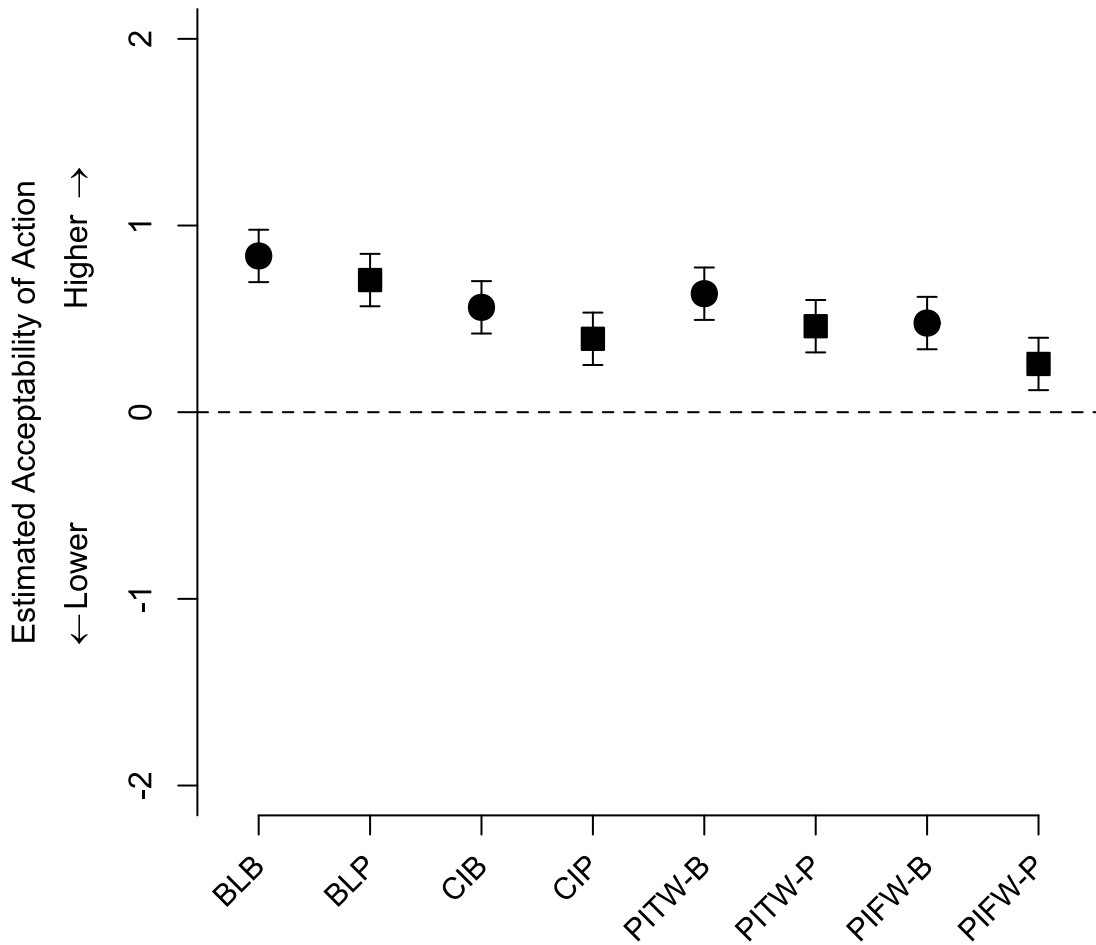


Figure 5. Moral acceptability by question and moral baseline of pushing/dropping the worker to stop the trolley. Circles represent when a decision did not involve personal force (button press), and squares indicate questions when personal force was involved (push worker). Error bars indicate standard error of the estimate. The horizontal dashed line indicates neutrality towards the action. BLB: Baseline Button. BLP: Baseline Push. PIFW-B: Partial Ignorance Footbridge Worker-Button. PIFW-P: Partial Ignorance Footbridge Worker-Push. PITW-B: Partial Ignorance

Track Workers-Button. PITW-P: Partial Ignorance Track Workers-Push. CIB: Complete Ignorance Button. CIP: Complete Ignorance Push. Y-axis was recoded to center 4 (neutral) as 0.

When incorporating moral baseline, all the factors of force, $\chi^2(1) = 11.308$, $p = .0008$, ignorance, $\chi^2(3) = 33.189$, $p < .0001$, and moral baseline, $\chi^2(1) = 60.8567$, $p < .0001$, were significant. None of the interactions of force and ignorance, $\chi^2(3) = 0.385$, $p = .9433$, force and moral baseline, $\chi^2(1) = 1.721$, $p = .1896$, ignorance and moral baseline, $\chi^2(3) = 7.441$, $p = .0591$, or all three combined, $\chi^2(3) = 0.641$, $p = .8869$, were significant. For these models, all questions were used in the regression. Figure 6 is a visual representation of estimated acceptability based on question and moral baseline. Baseline deontologists did not have statistically significant differences in acceptability regardless of the question (see Supplemental Tables 10–12), whereas those predisposed to consequentialism rated acceptability of pushing/dropping the worker according to the nature and extent of ignorance involved, but not according to force (see Supplemental Tables 10–12).

Acceptability of Action by Nature and Extent of Ignorance

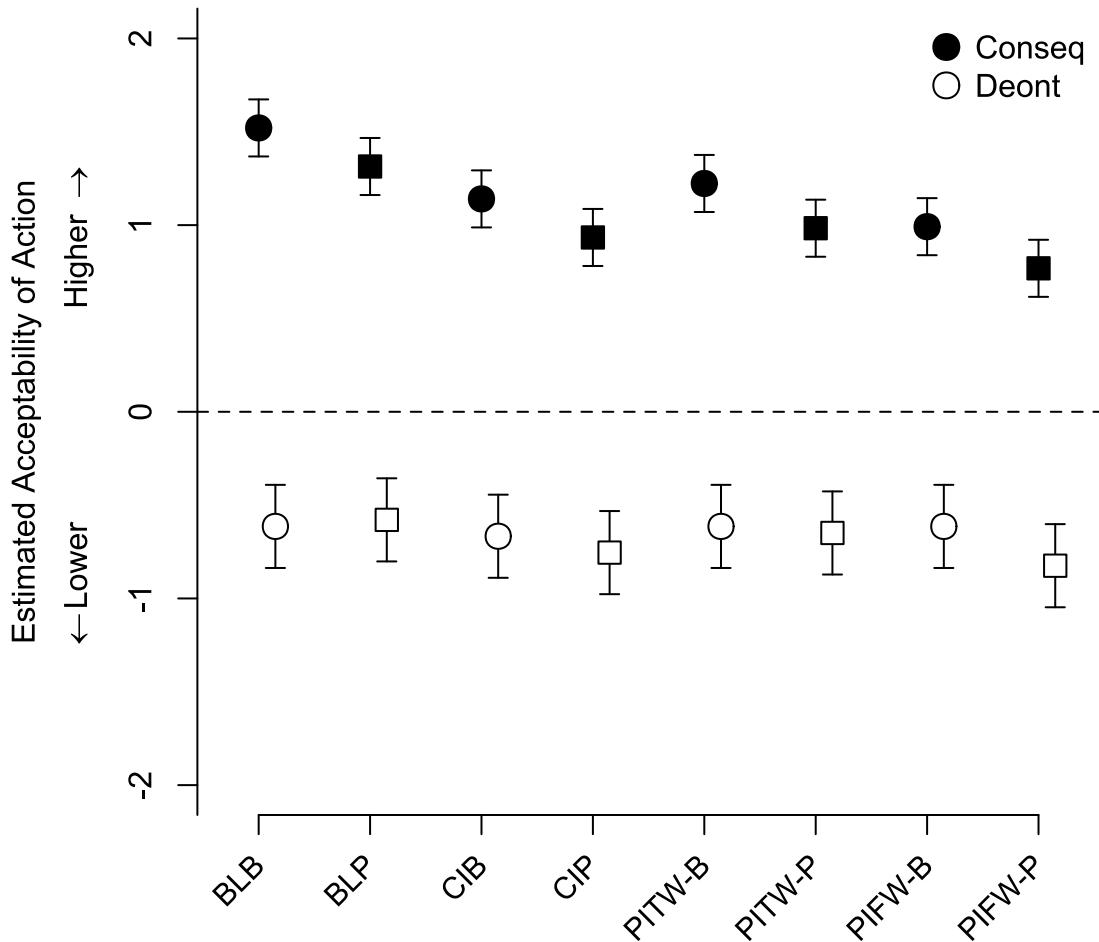


Figure 6. Moral acceptability by question and moral baseline of pushing/dropping the worker to stop the trolley based on moral baseline. Circles represent when a decision did not involve personal force (button press), and squares indicate questions when personal force was involved (push worker). Error bars indicate standard error of the estimate. The horizontal dashed line indicates neutrality towards the action. Conseq: baseline consequentialist. Deont: baseline deontologist. BLB: Baseline Button. BLP: Baseline Push. PIFW-B: Partial Ignorance Footbridge

Worker-Button. PIFW-P: Partial Ignorance Footbridge Worker-Push. PITW-B: Partial Ignorance Track Workers-Button. PITW-P: Partial Ignorance Track Workers-Push. CIB: Complete Ignorance Button. CIP: Complete Ignorance Push. All differences between consequentialists and deontologists were significant. Y-axis was recoded to center 4 (neutral) as 0.

7. DISCUSSION

The thesis of the logical priority of the epistemic predicts that epistemic considerations influence moral decision-making. Furthermore, given that the main claim is that epistemic considerations are fundamental to moral decision-making, we sought to investigate the effects that epistemic considerations have on moral decision-making and to compare those effects with those of other factors suggested by previous authors as similarly fundamental.

We investigated how ignorance of relevant consequences affects the moral judgments that persons make in Trolley Problem settings. This experimental apparatus was chosen for its relative convenience and the ease with which it can display the effects of ignorance on moral judgments, not because we are interested in the psychology of either utilitarian or deontological judgments *per se*. We found compelling evidence for the following claims: (a) that ignorance of consequences affects one's transient moral judgments; (b) that ignorance of consequences in some given set of circumstances affects whether one makes a deontological or a consequentialist judgment in these circumstances; (c) that ignorance of consequences has a greater influence upon those predisposed to make consequentialist judgments in the Baseline scenario, and tends to make their momentary judgments more deontological; and (d) that ignorance of consequences is indeed prior to other considerations and at least as fundamental to moral judgment as emotional

factors such as whether the action requires personal force. That is, the empirical evidence is broadly in line with the priority of the epistemic.

6.1 First Experiment

The first experiment tested whether ignorance of consequences affects one's momentary moral judgments and, if so, in what way. Using the standard Trolley Problem as a baseline, more than 80% of participants responded that they would push the button to switch tracks. Using their responses to this scenario to categorize participants as either baseline consequentialists or baseline deontologists, their responses to other scenarios indicated that ignorance of consequences affected their transitory moral judgments. Compared to their responses to the Baseline scenario, participants were on average less likely to push the button (i.e., to make the consequentialist judgment to switch tracks) the more ignorant they were of the consequences of available options.

According to the thesis of the logical priority of the epistemic, epistemic burdens affect a person's moral judgments, leading their momentary judgments out of line with their underlying moral predisposition. If true, then manipulating epistemic burdens – making them heavier or lighter, or more or less surmountable – should affect subjects' transient moral judgments. In the first study, respondents were given the option of solving a puzzle in order to surmount ignorance of consequences. We found that having to surmount an epistemic burden affected whether one pushed the button (Figure 1, Supplemental Table 1). When having to surmount an epistemic burden related to ignorance of the consequences of pushing the button, respondents were more likely to do nothing. When having to surmount a burden related to ignorance of the consequences

of doing nothing, respondents were more likely to push the button (Figure 1, Supplemental Table 1). In short, respondents were more likely to choose the less epistemically burdensome option. In particular, if epistemic burdens shape decision-making in the way predicted by the thesis of the priority of the epistemic, we should expect to observe in Know-How Ignorance that baseline consequentialists pushed the button less frequently on average, but that, because they were less inclined to push the button in the first place, the need to surmount an epistemic burden by solving the puzzle had no effect on baseline deontologists. This is exactly what we observed.

Ignorance of consequences evidently affects the transitory moral judgments of those otherwise predisposed to consequentialism and deontology. Participants categorized as deontologists given their response to the Baseline scenario were far more likely to act like consequentialists and to switch tracks when ignorant of consequences (Supplemental Tables 2 & 3). There was no difference in the responses of baseline consequentialists and deontologists in circumstances where they were entirely ignorant of the consequences of available options. Contrast these similar responses with subjects' responses in circumstances where their ignorance was more partial, such as those in which subjects knew the consequences of staying on the present track but were ignorant of the consequences of switching. Under conditions of only partial rather than complete ignorance, the difference between baseline consequentialists and deontologists was greater.

These results suggest that ignorance of consequences affects baseline consequentialists and deontologists alike. Knowledge of consequences seems to enable whatever disposes a person to one or the other kind of moral judgment, while ignorance of consequences appears to short-circuit the revelation of such predispositions. In the presence of ignorance of consequences, all

moral bets are off: baseline consequentialists are less inclined to consequentialism and baseline deontologists are less disposed to deontology.

6.2 Second Experiment

In the second study, we sought to investigate further this “enabling” effect of relevant knowledge and to test whether ignorance of consequences was indeed prior to other considerations relevant to moral judgments. Specifically, we tested the interaction between ignorance of consequences and the use of personal force, because it has been shown to be a significant factor in moral decision-making. Greene et al. (2009) showed that personal force is itself prior to other factors, perhaps by virtue of eliciting the effects that emotional factors have upon moral judgment (2008a).

The results from the second study support the claims that ignorance affects transient moral judgments, that ignorance of consequences makes transient moral judgments more deontological, and that ignorance of consequences more strongly affects baseline consequentialists than baseline deontologists (Figure 2). As in the first study, ignorance of consequences made the transitory judgments of baseline consequentialists more deontological and those of baseline deontologists more consequentialist. However, we also tested the interaction between personal force and ignorance. Although both ignorance and personal force lowered the likelihood that a person would make the momentary consequentialist choice to push the rail worker on to the tracks, the effects of ignorance and personal force appeared to be independent of each other (Supplemental Tables 4 & 5). This shows that ignorance of consequences is a significant factor and no less fundamental than personal force in moral judgment.

Although personal force affected moral judgments, it was not a significant factor in the extent to which a particular choice was judged morally acceptable. It influenced subjects' judgments about what to do, but not their judgments about whether the choice was morally acceptable. Ignorance of consequences, however, did have such an effect. Ignorance of consequences appeared to make people less likely to judge either pushing the button or pushing the rail worker onto the tracks as morally acceptable (Supplemental Tables 9 & 10) In other words, personal force affected moral judgment to a lesser degree than did ignorance of consequences.

Both personal force and ignorance of consequences matter to moral judgment, but ignorance of consequences seems to matter more to those predisposed to consequentialism. Although the judgments of baseline deontologists were affected by changes in their ignorance, they were less affected by varying degrees of ignorance than were the judgments of baseline consequentialists.

6.3 General Discussion

Scheall and Crutchfield claim that ignorance constrains moral judgment and that this effect is logically prior to the effects that other factors exert on moral judgment. The results of these two experiments support this thesis and some conclusions about the psychology of moral judgments. Greene et. al (2001, 2009) showed that scenarios that elicit emotional responses, such as scenarios requiring the use of personal force, are more likely to produce deontological judgments. Furthermore, the effects of other emotion-eliciting factors like physical distance can be reduced to the influence of personal force. Our studies indicate that personal force is indeed a

factor in moral judgment, but that ignorance of consequences is at least as, if not more, fundamental in determining one's transitory moral judgments. This being said, the results provide only limited support for the notion that knowledge enables a person's predisposition toward a particular type of moral judgment. Effectively surmounting an epistemic burden seems to move participants back toward their baseline judgments in the first experiment. Although ignorance made baseline consequentialists more deontological and baseline deontologists more consequentialist in the second experiment, there was no variation within different degrees of ignorance. (Figure 2, Supplemental Table 6).

The thesis of the priority of the epistemic is an indispensable part of any explanation of the present observations, but our results do not necessitate a replacement of other models of moral judgment, such as those that Greene et al. motivate. Though ignorance of consequences affected both baseline consequentialists and baseline deontologists, because the influence of personal force was independent of ignorance of consequences, we cannot conclude that epistemic considerations are prior to personal force. However, our results provide reasons to think that personal force is less fundamental to moral judgment than ignorance of consequences. First, we noted above that personal force had no effect on the moral acceptability of an action, but that ignorance had a significant effect. Second, adding personal force to the Baseline (Push) question in the second experiment made baseline deontologists more likely to act as consequentialists and to push the worker onto the track (Figure 2). Thus, there is a group of people who, in possession of full knowledge of consequences, will not press the button, but will push the person. Personal force is supposed to have the opposite effect.

Our results do not rule out any particular account as much as they rule in the significant role of knowledge and ignorance in moral decision-making. Though they do not necessarily cast

doubt upon any particular explanation of variations in moral judgment, they do suggest that explanations that omit knowledge and ignorance as factors are, at best, incomplete. Thus, one plausible model is that epistemic burdens constrain moral judgments to an equal or greater degree than other factors, but that, per Greene et al., baseline consequentialists and baseline deontologists base their moral judgments on different factors. The presence or absence of emotional factors, and knowledge or ignorance of consequences, jointly enable particular kinds of moral judgments.

Together, the two experiments provide strong evidence that epistemic factors are critical in determining one's moral judgments, and only slightly weaker evidence that they are more fundamental than other factors.

6.4 Limitations

There were limitations with the current studies that involve the power to detect effects, notably in Experiment 2, and the order in which questions were presented across both Experiments. While there was no statistically significant difference between how baseline deontologists responded to scenarios in Experiment 2, an observed power simulation was conducted with the *simr* package (Green & MacLeod, 2016). To determine the observed power based on the results. Based on 100 simulations including only baseline deontologists across scenarios, the sample of these predisposed to deontology in Experiment 2 was slightly underpowered to detect an effect ($\beta = .61$ [.51–.71]). Thus, some of the conclusions regarding the effects of force and knowledge on actions in the Trolley Problem for that sample should be considered tentative. Regarding ordering of the scenarios, how questions were presented could have influenced later responses. Because of the novelty of the questions, determining how the

order of questions affected moral decision-making was not our primary goal. This seems like an area for further study. For example, how might exposure to surmountable ignorance scenarios first affect responses to insurmountable ignorance scenarios presented second? However, despite these limitations, general effects in line with predictions of force and ignorance were still identified.

An additional limitation is the assumption that subjects were solving the problems as presented. But it could be that the participants, in spite of being given the relevant information, failed to incorporate all of the given information into their decision-making. In that case, non-epistemic factors (e.g., heuristics, biases, guesses, and rules of thumb) may have influenced their decisions. We did precede each study with a comprehension check, but even if that failed to weed out those disposed to ignore the problem prompts, the influence of non-epistemic factors in the context of ignorance is, as we indicate above, not only compatible with the logical priority of the epistemic, it may a consequence of it. Thus, the assertion that those participants who ignored the question prompts were more heavily influenced in their momentary decisions by non-epistemic factors presumes the logical priority of the epistemic.

8. CONCLUSION

The present work found that ignorance of consequences affects the momentary moral judgments of both those predisposed to consequentialism and those inclined to deontology, but that the effect was more profound for baseline consequentialists. Consequences are the signals or, if you prefer, the guardrails, that baseline consequentialists rely upon to make consequentialist judgments. Remove these guardrails, blind consequentialists to the signals that

good consequentialism requires, and it becomes increasingly difficult to be a good consequentialist.

This line of thought suggests the possibility of another experiment. Moral rules and perceived moral duties are to deontologists what consequences are to consequentialists. Moral principles guide those inclined to deontology to good deontological judgments. If it were possible to blind baseline deontologists to any moral principles relevant in a particular context, it is reasonable to assume that they would react similarly to the loss of their moral guardrails as baseline consequentialists in our two experiments reacted to the loss of their moral guardrails. That is, one would predict that ignorance of relevant moral rules would complicate deontological judgments (and probably consequentialist judgments as well) and incline baseline deontologists further toward consequentialism. Perhaps baseline deontologists would then be more likely to rely on heuristics, biases, guesses, and rules of thumb (Gurevich, 2019) rather than moral rules.

There may also be another way to get at the possible priority of the epistemic over other considerations, such as Greene's emotional factors. In Greene's work, as in our second experiment, participants *knew* that they would or would not have to use personal force. However, if one could induce different momentary moral judgments by obscuring this knowledge – by, say, making subjects ignorant whether a particular course of action would require the use of personal force – then this would seem to provide further evidence for the thesis of the logical priority of the epistemic in human decision-making.

Pending the results of these supplementary experiments, it seems fair to assert that the logical priority of the epistemic is tentatively established.

REFERENCES

- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*, 1–48.
- Chen, H., Cohen, P., & Chen. (2010). How big is a big odds ratio? Interpreting the magnitudes of odds ratios in epidemiological studies. *Communications in Statistics - Simulation and Computation*. 2010, *39*(4), 860–864.
- 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*(12), 1082–1089. <https://doi.org/10.1111/j.1467-9280.2006.01834.x>
- Fox, J., & Weisberg, S. (2019). *An R companion to applied regression* (3rd edition). Thousand Oaks CA: Sage. URL: <https://socialsciences.mcmaster.ca/jfox/Books/Companion/>
- Green, P., & MacLeod, C. J. (2016). simr: An R package for power analysis of generalised linear mixed models by simulation. *Methods in Ecology and Evolution*, *7*(4), 493–498.
- 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 (New York, N.Y.)*, *293*(5537), 2105–2108. <https://doi.org/10.1126/science.1062872>
- Greene, J. D. (2008a). *The secret joke of Kant's soul*. In W. Sinnott-Armstrong (Ed.), *Moral psychology, Vol. 3. The neuroscience of morality: Emotion, brain disorders, and development* (p. 35–80). MIT Press.
- Greene, J. D., Morelli, S. A., Lowenberg, K., Nystrom, L. E., & Cohen, J. D. (2008b). Cognitive load selectively interferes with utilitarian moral judgment. *Cognition*, *107*(3), 1144–1154. <https://doi.org/10.1016/j.cognition.2007.11.004>

- Greene, J. D., Cushman, F. A., Stewart, L. E., Lowenberg, K., Nystrom, L. E., & Cohen, J. D. (2009). Pushing moral buttons: The interaction between personal force and intention in moral judgment. *Cognition*, *111*(3), 364–371.
<https://doi.org/10.1016/j.cognition.2009.02.001>
- Gurevich, G. (2019). Ignorance is strength: May human mind's unique capabilities stem from its limitations? *Consciousness and Cognition*, *69*(September 2018), 1–13.
<https://doi.org/10.1016/j.concog.2019.01.009>
- Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, *108*(4), 814–834. <https://doi.org/10.1037/0033-295X.108.4.814>
- Hampshire, Stuart (1951). Symposium: Freedom of the will. *Aristotelian Society Supplementary Volume 25*: 161-178.
- Hare, R.M. (1951). Symposium: Freedom of the will. *Aristotelian Society Supplementary Volume 25*: 201-216.
- Hare, R.M. (1963). *Freedom and Reason*. Oxford: Clarendon Press.
- Kahane, G. (2015). Sidetracked by trolleys: Why sacrificial moral dilemmas tell us little (or nothing) about utilitarian judgment. *Social Neuroscience*, *10*, 551–560.
<http://dx.doi.org/10.1080/17470919.2015.1023400>
- Kahane, G., Everett, J. A. C., Earp, B. D., Farias, M., & Savulescu, J. (2015). 'Utilitarian' judgments in sacrificial moral dilemmas do not reflect impartial concern for the greater good. *Cognition*, *134*, 193–209. <http://dx.doi.org/10.1016/j.cognition.2014.10.005>

- Kahane, G., Everett, J. A., Earp, B. D., Caviola, L., Faber, N. S., Crockett, M. J., & Savulescu, J. (2018). Beyond sacrificial harm: A two-dimensional model of utilitarian psychology. *Psychological Review*, 125 (2), 131-164.
- Kahneman, D. (2011). *Thinking Fast and Slow*. New York: Farrar, Straus and Giroux
- Kahneman, D. & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47 (2), 263–291. <https://doi.org/10.2307/1914185>.
- Maier, M., Bartoš, F., Stanley, T. D., Shanks, D. R., Harris, A. J. L., & Wagenmakers, E.-J. (2022). No evidence for nudging after adjusting for publication bias. *Proceedings of the National Academy of Sciences*, 119(31), e2200300119. <https://doi.org/10.1073/pnas.2200300119>
- Nuñez, N., Laurent, S., & Gray, J. (2014). Is negligence a first cousin to intentionality? Lay conceptions of negligence and its relationship to intentionality. *Applied Cognitive Psychology*, 28: 55–65.
- Nuñez, N., Laurent S., & Schweitzer, K. (2016). Unintended, but still blameworthy: the roles of awareness, desire, and anger in negligence, restitution, and punishment. *Cognition and emotion*, 30(7), 1271-1288.
- Lenth, R. (2020). emmeans: Estimated Marginal Means, aka Least-Squares Means (Version 1.5.2-1). URL: <https://CRAN.R-project.org/package=emmeans>
- Littlejohn, Clayton (2009). ‘Ought’, ‘can’, and practical reasons. *American Philosophical Quarterly* 46 (4): 363-373
- Mizrahi, Moti (2015). Ought, can, and presupposition: An experimental study. *Methode* 4 (6):232-243.

- R Core Team (2020). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Scheall S. (2019). Ignorance and the Incentive Structure Confronting Policymakers. *Cosmos + Taxis*, 7 (1–2): 39–51.
- Scheall, S., & Crutchfield, P. (2020). The Priority of the Epistemic. *Episteme*, 1-12.
doi:10.1017/epi.2019.56
- Schultz, T. R. & Wright, K. (1985). Concepts of negligence and intention in the assignment of moral responsibility. *Canadian Journal of Behavioral Science* 17 (2): 97-108.
- Sinnott-Armstrong, W. (1984). ‘Ought’ conversationally implies ‘can’. *The Philosophical Review* 93: 249-261.
- Thaler, R & Sunstein, C. (2009). *Nudge: Improving Decisions About Health, Wealth, and Happiness*. New York: Penguin.
- Vallentyne, P. (1989). Two types of moral dilemmas. *Erkenntnis* 30: 301-318.
- Vogelstein, E. (2012). Subjective reasons. *Ethical Theory and Moral Practice* 15: 239-257

APPENDIX

A. Tables

Supplemental Table 1

Experiment 1 contrasts by question (simple model) for switching the track

Contrast	Δ_{LO}	SE	95% CI	<i>z</i>	<i>p</i>
PIBT - BL	-2.77	0.29	[-3.67, -1.86]	-9.59	< .001
CI - BL	-2.67	0.29	[-3.57, -1.78]	-9.32	< .001
CI - PIBT	0.09	0.24	[-0.66, 0.84]	0.38	.707
PISwT - BL	-0.55	0.29	[-1.46, 0.37]	-1.87	.366
PISwT - PIBT	2.22	0.26	[1.39, 3.04]	8.39	< .001
PISwT - CI	2.13	0.26	[1.31, 2.95]	8.09	< .001
PIStT - BL	-2.45	0.28	[-3.34, -1.56]	-8.63	< .001
PIStT - PIBT	0.31	0.24	[-0.43, 1.06]	1.31	.570
PIStT - CI	0.22	0.24	[-0.52, 0.97]	0.94	.698
PIStT - PISwT	-1.90	0.26	[-2.71, -1.09]	-7.33	< .001
SISwT - BL	-1.19	0.28	[-2.06, -0.31]	-4.22	< .001
SISwT - PIBT	1.58	0.25	[0.81, 2.35]	6.41	< .001
SISwT - CI	1.49	0.25	[0.72, 2.25]	6.08	< .001
SISwT - PISwT	-0.64	0.26	[-1.44, 0.17]	-2.47	.109
SISwT - PIStT	1.27	0.24	[0.51, 2.02]	5.23	< .001
SISStT - BL	-1.56	0.28	[-2.43, -0.69]	-5.60	< .001
SISStT - PIBT	1.20	0.24	[0.45, 1.95]	5.01	< .001
SISStT - CI	1.11	0.24	[0.37, 1.86]	4.66	< .001
SISStT - PISwT	-1.01	0.26	[-1.81, -0.22]	-3.97	.001
SISStT - PIStT	0.89	0.24	[0.15, 1.63]	3.77	.002
SISStT - SISwT	-0.38	0.24	[-1.12, 0.37]	-1.57	.502
KHI - BL	-1.95	0.28	[-2.82, -1.07]	-6.95	< .001
KHI - PIBT	0.82	0.24	[0.07, 1.57]	3.43	.007
KHI - CI	0.73	0.24	[-0.01, 1.47]	3.07	.019
KHI - PISwT	-1.40	0.26	[-2.20, -0.60]	-5.45	< .001
KHI - PIStT	0.51	0.24	[-0.23, 1.24]	2.15	.219
KHI - SISwT	-0.76	0.24	[-1.51, -0.01]	-3.18	.015
KHI - SISStT	-0.38	0.23	[-1.11, 0.35]	-1.64	.502

Note. Contrasts for all questions in experiment one without incorporation of moral baseline. Δ_{LO} : Change in log odds between questions. SE: Standard error based on model. 95% CI: the upper and lower bounds of the 95% confidence interval. z: The test statistic for the difference. Multiple comparisons were corrected for using the Holm correction. BL: Baseline trolley problem. PIBT: Partial Ignorance Both Tracks. CI: Complete Ignorance. PISwT: Partial Ignorance Switch Track. PIStT: Partial Ignorance Stay Track. SISwT: Surmountable Ignorance Switch Track. SIStT Surmountable Ignorance Stay Track. KHI: Know How Ignorance.

Supplemental Table 2

Experiment 1 contrasts by moral baseline and question (moral baseline model) for switching the tracks

Contrast	Δ_{LO}	SE	95% CI	<i>z</i>	<i>p</i>
Consequentialists					
CI - PIBT	0.20	0.26	[-0.58, 0.97]	0.77	> .999
PISwT - PIBT	2.70	0.31	[1.77, 3.63]	8.79	< .001
PISwT - CI	2.50	0.30	[1.58, 3.42]	8.24	< .001
PIStT - PIBT	0.42	0.25	[-0.35, 1.20]	1.67	.456
PIStT - CI	0.23	0.25	[-0.54, 0.99]	0.91	> .999
PIStT - PISwT	-2.27	0.30	[-3.19, -1.36]	-7.55	< .001
SISwT - PIBT	1.74	0.27	[0.93, 2.55]	6.50	< .001
SISwT - CI	1.55	0.26	[0.74, 2.35]	5.85	< .001
SISwT - PISwT	-0.96	0.30	[-1.88, -0.04]	-3.16	.013
SISwT - PIStT	1.32	0.26	[0.52, 2.11]	5.03	< .001
SISwT - PIBT	1.30	0.26	[0.52, 2.09]	5.03	< .001
SISwT - CI	1.10	0.26	[0.33, 1.88]	4.33	< .001
SISwT - PISwT	-1.40	0.30	[-2.30, -0.49]	-4.69	< .001
SISwT - PIStT	0.88	0.25	[0.11, 1.65]	3.46	.005
SISwT - SISwT	-0.44	0.26	[-1.23, 0.35]	-1.69	.456
KHI - PIBT	1.18	0.26	[0.40, 1.97]	4.59	< .001
KHI - CI	0.99	0.26	[0.21, 1.76]	3.88	.001
KHI - PISwT	-1.52	0.30	[-2.42, -0.61]	-5.08	< .001
KHI - PIStT	0.76	0.25	[-0.01, 1.53]	3.00	.019
KHI - SISwT	-0.56	0.26	[-1.35, 0.23]	-2.14	.193
KHI - SISwT	-0.12	0.25	[-0.89, 0.65]	-0.47	> .999
Deontologists					
CI - PIBT	-0.62	0.65	[-2.58, 1.34]	-0.96	> .999
PISwT - PIBT	-0.14	0.61	[-2.01, 1.72]	-0.24	> .999
PISwT - CI	0.48	0.64	[-1.46, 2.42]	0.75	> .999
PIStT - PIBT	-0.48	0.63	[-2.39, 1.44]	-0.75	> .999
PIStT - CI	0.15	0.66	[-1.84, 2.14]	0.22	> .999
PIStT - PISwT	-0.33	0.62	[-2.22, 1.56]	-0.53	> .999
SISwT - PIBT	0.39	0.61	[-1.45, 2.23]	0.64	> .999
SISwT - CI	1.01	0.63	[-0.91, 2.93]	1.60	> .999

SISwT - PISwT	0.53	0.60	[-1.28, 2.34]	0.89	> .999
SISwT - PIStT	0.86	0.62	[-1.01, 2.73]	1.40	> .999
SIStT - PIBT	0.21	0.61	[-1.63, 2.06]	0.35	> .999
SIStT - CI	0.84	0.63	[-1.09, 2.76]	1.32	> .999
SIStT - PISwT	0.36	0.60	[-1.46, 2.17]	0.60	> .999
SIStT - PIStT	0.69	0.62	[-1.18, 2.56]	1.12	> .999
SIStT - SISwT	-0.17	0.59	[-1.96, 1.61]	-0.30	> .999
KHI - PIBT	-2.29	0.87	[-4.94, 0.35]	-2.63	.161
KHI - CI	-1.67	0.89	[-4.37, 1.02]	-1.89	.951
KHI - PISwT	-2.15	0.87	[-4.78, 0.48]	-2.48	.233
KHI - PIStT	-1.82	0.88	[-4.48, 0.85]	-2.07	.649
KHI - SISwT	-2.68	0.86	[-5.30, -0.07]	-3.11	.039
KHI - SIStT	-2.51	0.86	[-5.13, 0.11]	-2.91	.072

Note. Contrasts for all questions in experiment one by moral baseline. Because of convergence issues, the baseline trolley question is not included. Δ_{LO} : Change in log odds between questions. SE: Standard error based on model. 95% CI: the upper and lower bounds of the 95% confidence interval. z: The test statistic for the difference. Multiple comparisons were corrected for using the Holm correction. PIBT: Partial Ignorance Both Tracks. CI: Complete Ignorance. PISwT: Partial Ignorance Switch Track. PIStT: Partial Ignorance Stay Track. SISwT: Surmountable Ignorance Switch Track. SIStT Surmountable Ignorance Stay Track. KHI: Know How Ignorance.

Supplemental Table 3

Experiment 1 contrasts by question and moral baseline (moral baseline model) for switching the tracks

Question	Contrast	Δ_{LO}	SE	95% CI	<i>z</i>	<i>p</i>
PIBT	Deont - Conseq	0.23	0.51	[-0.77, 1.23]	0.46	.647
CI	Deont - Conseq	-0.58	0.54	[-1.64, 0.47]	-1.08	.279
PISwT	Deont - Conseq	-2.61	0.53	[-3.64, -1.58]	-4.97	< .001
PIStT	Deont - Conseq	-0.67	0.52	[-1.69, 0.35]	-1.28	.201
SISwT	Deont - Conseq	-1.12	0.49	[-2.08, -0.15]	-2.27	.023
SISStT	Deont - Conseq	-0.85	0.49	[-1.81, 0.11]	-1.74	.081
KHI	Deont - Conseq	-3.24	0.80	[-4.81, -1.68]	-4.07	< .001

Note. Contrasts for moral baseline and all questions. Because of convergence issues, the baseline trolley question is not included. Δ_{LO} : Change in log odds between questions. SE: Standard error based on model. 95% CI: the upper and lower bounds of the 95% confidence interval. *z*: The test statistic for the difference. Multiple comparisons were corrected for using the Holm correction. PIBT: Partial Ignorance Both Tracks. CI: Complete Ignorance. PISwT: Partial Ignorance Switch Track. PIStT: Partial Ignorance Stay Track. SISwT: Surmountable Ignorance Switch Track. SISStT: Surmountable Ignorance Stay Track. KHI: Know How Ignorance. Deont: Baseline deontologists. Conseq: Baseline consequentialists.

Supplemental Table 4

Experiment 2 Contrasts by Knowledge and Force (simple model) for pushing/dropping the worker

Knowledge	Contrast	Δ_{LO}	SE	95% CI	<i>z</i>	<i>p</i>
BL	Push - Button	-0.58	0.30	[-1.16, 0.00]	-1.97	.049
PIFW	Push - Button	-0.58	0.28	[-1.13, -0.02]	-2.04	.042
PITW	Push - Button	-0.35	0.28	[-0.89, 0.20]	-1.24	.215
CI	Push - Button	-0.74	0.28	[-1.30, -0.19]	-2.64	.008

Note. Contrasts for all questions in experiment one without incorporation of moral baseline. Δ_{LO} : Change in log odds between questions. SE: Standard error based on model. 95% CI: the upper and lower bounds of the 95% confidence interval. *z*: The test statistic for the difference. Multiple comparisons were corrected for using the Holm correction. BL: Baseline, all outcomes are known. PIFW: Partial Ignorance Footbridge Worker. PITW: Partial Ignorance Track Workers. CI: Complete ignorance.

Supplemental Table 5

*Experiment 2 Contrasts by Force and Knowledge (simple model)
for pushing/dropping the worker*

Contrast	Δ_{LO}	SE	95% CI	z	p
<u>Button</u>					
PIFW -BL	-0.63	0.30	[-1.41, 0.15]	-2.12	.135
PITW -BL	-1.07	0.29	[-1.85, -0.30]	-3.65	.002
PITW - PIFW	-0.45	0.28	[-1.20, 0.30]	-1.57	.346
CI -BL	-0.72	0.30	[-1.50, 0.06]	-2.44	.074
CI - PIFW	-0.09	0.29	[-0.84, 0.66]	-0.32	.750
CI - PITW	0.36	0.28	[-0.39, 1.10]	1.26	.417
<u>Push</u>					
PIFW -BL	-0.62	0.28	[-1.37, 0.13]	-2.19	.113
PITW -BL	-0.84	0.28	[-1.59, -0.09]	-2.95	.016
PITW - PIFW	-0.22	0.28	[-0.95, 0.52]	-0.77	> .999
CI -BL	-0.88	0.28	[-1.63, -0.13]	-3.10	.012
CI - PIFW	-0.26	0.28	[-0.99, 0.48]	-0.93	> .999
CI - PITW	-0.04	0.28	[-0.78, 0.69]	-0.15	> .999

Note. Contrasts for all questions in experiment one without incorporation of moral baseline. Δ_{LO} : Change in log odds between questions. SE: Standard error based on model. 95% CI: the upper and lower bounds of the 95% confidence interval. z: The test statistic for the difference. Multiple comparisons were corrected for using the Holm correction. BL: Baseline, all outcomes are known. PIFW: Partial Ignorance Footbridge Worker. PITW: Partial Ignorance Track Workers. CI: Complete ignorance.

Supplemental Table 6

Experiment 2 contrasts by moral baseline, force, and knowledge (moral baseline model) for pushing/dropping the worker

Moral Baseline	Contrast	Δ_{LO}	SE	95% CI	z	p
				<u>Button</u>		
	PIFW -BL	NA	NA	NA	NA	NA
	PITW -BL	NA	NA	NA	NA	NA
	PITW - PIFW	-0.26	0.37	[-1.22, 0.70]	-0.72	1.000
	CI -BL	NA	NA	NA	NA	NA
	CI - PIFW	-0.29	0.36	[-1.26, 0.67]	-0.81	1.000
	CI - PITW	-0.03	0.36	[-0.98, 0.91]	-0.09	1.000
Consequentialists				<u>Push</u>		
	PIFW -BL	-0.95	0.36	[-1.91, 0.01]	-2.61	.036
	PITW -BL	-1.06	0.36	[-2.02, -0.10]	-2.92	.018
	PITW - PIFW	-0.11	0.34	[-1.02, 0.80]	-0.32	1.000
	CI -BL	-1.27	0.36	[-2.23, -0.31]	-3.48	.003
	CI - PIFW	-0.32	0.34	[-1.22, 0.59]	-0.92	1.000
	CI - PITW	-0.20	0.34	[-1.11, 0.70]	-0.60	1.000
				<u>Button</u>		
	PIFW -BL	NA	NA	NA	NA	NA
	PITW -BL	NA	NA	NA	NA	NA
	PITW - PIFW	-1.17	0.64	[-2.85, 0.51]	-1.84	.332
	CI -BL	NA	NA	NA	NA	NA
	CI - PIFW	0.39	0.60	[-1.18, 1.96]	0.66	1.000
	CI - PITW	1.56	0.64	[-0.12, 3.24]	2.45	.086
Deontologists				<u>Push</u>		
	PIFW -BL	0.17	0.61	[-1.44, 1.78]	0.28	1.000
	PITW -BL	-0.42	0.62	[-2.06, 1.23]	-0.67	1.000
	PITW - PIFW	-0.59	0.62	[-2.23, 1.05]	-0.95	1.000
	CI -BL	0.11	0.61	[-1.51, 1.72]	0.18	1.000
	CI - PIFW	-0.06	0.61	[-1.67, 1.54]	-0.11	1.000
	CI - PITW	0.52	0.62	[-1.12, 2.17]	0.84	1.000

Note. Contrasts for all questions in experiment one without incorporation of moral baseline. Δ_{LO} : Change in log odds between questions. SE: Standard error based on model. 95% CI: the upper and lower bounds of the 95% confidence interval. z: The test statistic for the difference. Multiple comparisons were corrected for using the Holm correction. BL: Baseline, all outcomes

are known. PIFW: Partial Ignorance Footbridge Worker. PITW: Partial Ignorance Track Workers. CI: Complete ignorance. NAs indicate comparisons that were missing due to the exclusion of the baseline question.

Supplemental Table 7

Experiment 2 contrasts by moral baseline, knowledge, and force (moral baseline model) for pushing/dropping the worker

Moral Baseline	Knowledge	Contrast	Δ_{LO}	SE	95% CI	<i>z</i>	<i>p</i>
Consequentialists	BL	Push - Button	NA	NA	NA	NA	NA
	PIFW	Push - Button	-0.77	0.36	[-1.48, -0.07]	-2.15	.031
	PITW	Push - Button	-0.62	0.35	[-1.31, 0.07]	-1.77	.077
	CI	Push - Button	-0.80	0.35	[-1.48, -0.11]	-2.27	.024
Deontologists	BL	Push - Button	NA	NA	NA	NA	NA
	PIFW	Push - Button	-0.15	0.61	[-1.34, 1.03]	-0.25	.801
	PITW	Push - Button	0.43	0.64	[-0.84, 1.69]	0.66	.508
	CI	Push - Button	-0.61	0.60	[-1.79, 0.57]	-1.01	.312

Note. Contrasts for all questions in experiment one without incorporation of moral baseline.

Δ_{LO} : Change in log odds between questions. SE: Standard error based on model. 95% CI: the upper and lower bounds of the 95% confidence interval. *z*: The test statistic for the difference.

Multiple comparisons were corrected for using the Holm correction. BL: Baseline, all outcomes are known. PIFW: Partial Ignorance Footbridge Worker. PITW: Partial Ignorance Track Workers.

CI: Complete ignorance. NAs indicate comparisons that were missing due to the exclusion of the baseline question.

Supplemental Table 8

Experiment 2 contrasts by knowledge, force, and moral baseline (moral baseline model) for pushing/dropping the worker

Knowledge	Force	Contrast	Δ_{LO}	SE	95% CI	<i>z</i>	<i>p</i>
BL	Button	Deont - Conseq	NA	NA	NA	NA	NA
	Push	Deont - Conseq	-4.33	0.70	[-5.69, -2.96]	-6.20	< .001
PIFW	Button	Deont - Conseq	-3.83	0.68	[-5.16, -2.49]	-5.62	< .001
	Push	Deont - Conseq	-3.20	0.67	[-4.52, -1.89]	-4.76	< .001
PITW	Button	Deont - Conseq	-4.73	0.73	[-6.15, -3.31]	-6.52	< .001
	Push	Deont - Conseq	-3.68	0.70	[-5.05, -2.32]	-5.29	< .001
CI	Button	Deont - Conseq	-3.14	0.66	[-4.43, -1.85]	-4.75	< .001
	Push	Deont - Conseq	-2.95	0.67	[-4.27, -1.64]	-4.40	< .001

Note. Contrasts for all questions in experiment one without incorporation of moral baseline. Δ_{LO} : Change in log odds between questions. SE: Standard error based on model. 95% CI: the upper and lower bounds of the 95% confidence interval. *z*: The test statistic for the difference. Multiple comparisons were corrected for using the Holm correction. BL: Baseline, all outcomes are known. PIFW: Partial Ignorance Footbridge Worker. PITW: Partial Ignorance Track Workers. CI: Complete ignorance. Deont: Baseline deontologists. Conseq: Baseline consequentialists. NAs indicate comparisons that were missing due to the exclusion of the baseline question.

Supplemental Table 9

Experiment 2 contrasts by knowledge and force (simple model) for acceptability of action

Knowledge	Contrast	Δ_M	SE	95% CI	<i>t</i>	<i>p</i>
BL	Push - Button	-0.13	0.10	[-0.33, 0.07]	-1.26	.209
PIFW	Push - Button	-0.17	0.10	[-0.37, 0.03]	-1.64	.102
PITW	Push - Button	-0.17	0.10	[-0.38, 0.03]	-1.69	.091
CI	Push - Button	-0.22	0.10	[-0.42, -0.02]	-2.13	.033

Note. Contrasts for all questions in experiment one without incorporation of moral baseline.

Δ_M : Change in means between questions. SE: Standard error based on model. 95% CI: the upper and lower bounds of the 95% confidence interval. *z*: The test statistic for the difference.

Multiple comparisons were corrected for using the Holm correction. BL: Baseline, all outcomes are known. PIFW: Partial Ignorance Footbridge Worker. PITW: Partial Ignorance Track Workers. CI: Complete ignorance.

Supplemental Table 10

Experiment 2 contrasts by force and knowledge (simple model) for acceptability of action

Contrast	Δ_M	SE	95% CI	<i>t</i>	<i>p</i>
<u>Button</u>					
PIFW -BL	-0.28	0.10	[-0.55, 0.00]	-2.68	.038
PITW -BL	-0.20	0.10	[-0.47, 0.07]	-1.97	.198
PITW - PIFW	0.07	0.10	[-0.20, 0.34]	0.71	.826
CI -BL	-0.36	0.10	[-0.63, -0.09]	-3.50	.003
CI - PIFW	-0.08	0.10	[-0.36, 0.19]	-0.82	.826
CI - PITW	-0.16	0.10	[-0.43, 0.11]	-1.53	.379
<u>Push</u>					
PIFW -BL	-0.31	0.10	[-0.59, -0.04]	-3.06	.011
PITW -BL	-0.25	0.10	[-0.52, 0.02]	-2.40	.066
PITW - PIFW	0.07	0.10	[-0.20, 0.34]	0.66	.512
CI -BL	-0.45	0.10	[-0.72, -0.18]	-4.37	< .001
CI - PIFW	-0.13	0.10	[-0.41, 0.14]	-1.31	.380
CI - PITW	-0.20	0.10	[-0.47, 0.07]	-1.97	.148

Note. Contrasts for all questions in experiment one without incorporation of moral baseline.

Δ_M : Change in means between questions. SE: Standard error based on model. 95% CI: the upper and lower bounds of the 95% confidence interval. *z*: The test statistic for the difference.

Multiple comparisons were corrected for using the Holm correction. BL: Baseline, all outcomes are known. PIFW: Partial Ignorance Footbridge Worker. PITW: Partial Ignorance Track Workers. CI: Complete ignorance.

Supplemental Table 11

Experiment 2 contrasts by moral baseline, force, and knowledge (moral baseline model) for acceptability of action

Moral Baseline	Contrast	Δ_M	SE	95% CI	<i>t</i>	<i>p</i>	
<u>Button</u>							
Consequentialists	PIFW -BL	-0.38	0.13	[-0.71, -0.05]	-3.05	.012	
	PITW -BL	-0.30	0.13	[-0.63, 0.03]	-2.39	.068	
	PITW - PIFW	0.08	0.13	[-0.25, 0.41]	0.66	.507	
	CI -BL	-0.53	0.13	[-0.86, -0.20]	-4.24	< .001	
	CI - PIFW	-0.15	0.13	[-0.48, 0.18]	-1.19	.466	
	CI - PITW	-0.23	0.13	[-0.56, 0.10]	-1.86	.191	
	<u>Push</u>						
	PIFW -BL	-0.38	0.13	[-0.71, -0.05]	-3.05	.012	
	PITW -BL	-0.33	0.13	[-0.66, 0.00]	-2.65	.032	
	PITW - PIFW	0.05	0.13	[-0.28, 0.38]	0.40	.691	
	CI -BL	-0.55	0.13	[-0.87, -0.22]	-4.38	< .001	
	CI - PIFW	-0.17	0.13	[-0.49, 0.16]	-1.33	.370	
CI - PITW	-0.21	0.13	[-0.54, 0.11]	-1.72	.255		
<u>Button</u>							
Deontologists	PIFW -BL	-0.05	0.18	[-0.53, 0.43]	-0.29	> .999	
	PITW -BL	0.00	0.18	[-0.48, 0.48]	0.00	> .999	
	PITW - PIFW	0.05	0.18	[-0.43, 0.53]	0.29	> .999	
	CI -BL	0.00	0.18	[-0.48, 0.48]	0.00	> .999	
	CI - PIFW	0.05	0.18	[-0.43, 0.53]	0.29	> .999	
	CI - PITW	0.00	0.18	[-0.48, 0.48]	0.00	> .999	
	<u>Push</u>						
	PIFW -BL	-0.18	0.18	[-0.66, 0.30]	-0.97	> .999	
	PITW -BL	-0.07	0.18	[-0.55, 0.41]	-0.39	> .999	
	PITW - PIFW	0.11	0.18	[-0.37, 0.59]	0.58	> .999	
	CI -BL	-0.25	0.18	[-0.73, 0.23]	-1.35	> .999	
	CI - PIFW	-0.07	0.18	[-0.55, 0.41]	-0.39	> .999	
CI - PITW	-0.18	0.18	[-0.66, 0.30]	-0.97	> .999		

Note. Contrasts for all questions in experiment one without incorporation of moral baseline. Δ_M : Change in means between questions. SE: Standard error based on model. 95% CI: the upper and lower bounds of the 95% confidence interval. *z*: The test statistic for the difference. Multiple comparisons were corrected for using the Holm correction. BL: Baseline, all outcomes

are known. PIFW: Partial Ignorance Footbridge Worker. PITW: Partial Ignorance Track Workers. CI: Complete ignorance. All questions were included in this model.

Supplemental Table 12

Experiment 2 contrasts by moral baseline, knowledge, and for (moral baseline model) for acceptability of action

Moral Baseline	Knowledge	Contrast	Δ_M	SE	95% CI	<i>t</i>	<i>p</i>
Consequentialists	BL	Push - Button	-0.21	0.13	[-0.45, 0.04]	-1.66	.098
	PIFW	Push - Button	-0.21	0.13	[-0.45, 0.04]	-1.66	.098
	PITW	Push - Button	-0.24	0.13	[-0.48, 0.00]	-1.92	.055
	CI	Push - Button	-0.22	0.13	[-0.47, 0.02]	-1.79	.074
Deontologists	BL	Push - Button	0.04	0.18	[-0.32, 0.39]	0.19	.847
	PIFW	Push - Button	-0.09	0.18	[-0.44, 0.27]	-0.48	.629
	PITW	Push - Button	-0.04	0.18	[-0.39, 0.32]	-0.19	.847
	CI	Push - Button	-0.21	0.18	[-0.57, 0.15]	-1.16	.246

Note. Contrasts for all questions in experiment one without incorporation of moral baseline. Δ_M : Change in means between questions. 95% CI: the upper and lower bounds of the 95% confidence interval. *t*: The test statistic for the difference. Multiple comparisons were corrected for using the Holm correction. BL: Baseline, all outcomes are known. PIFW: Partial Ignorance Footbridge Worker. PITW: Partial Ignorance Track Workers. CI: Complete ignorance. All questions were included in this model.

Supplemental Table 13

Experiment 2 contrasts by knowledge, force, and moral baseline (moral baseline model) for acceptability of action

Knowledge	Force	Contrast	Δ_M	SE	95% CI	<i>t</i>	<i>p</i>
BL	Button	Deont - Conseq	-2.13	0.27	[-2.67, -1.60]	-7.90	< .001
	Push	Deont - Conseq	-1.89	0.27	[-2.43, -1.36]	-7.01	< .001
PIFW	Button	Deont - Conseq	-1.81	0.27	[-2.34, -1.27]	-6.69	< .001
	Push	Deont - Conseq	-1.69	0.27	[-2.22, -1.16]	-6.25	< .001
PITW	Button	Deont - Conseq	-1.84	0.27	[-2.37, -1.30]	-6.80	< .001
	Push	Deont - Conseq	-1.63	0.27	[-2.17, -1.10]	-6.05	< .001
CI	Button	Deont - Conseq	-1.61	0.27	[-2.14, -1.07]	-5.95	< .001
	Push	Deont - Conseq	-1.59	0.27	[-2.13, -1.06]	-5.90	< .001

Note. Contrasts for all questions in experiment one without incorporation of moral baseline. Δ_M : Change in means between questions. 95% CI: the upper and lower bounds of the 95% confidence interval. *t*: The test statistic for the difference. Multiple comparisons were corrected for using the Holm correction. BL: Baseline, all outcomes are known. PIFW: Partial Ignorance Footbridge Worker. PITW: Partial Ignorance Track Workers. CI: Complete ignorance. Deont:

Baseline deontologists. Conseq: Baseline consequentialists. All questions were included in this model.

Supplemental Table 14
Experiment 1 Demographic Variables

	Total Sample (n=176)		Consequentialists (n=151)		Deontologists (n=25)	
	Mean/%	SD	Mean/%	SD	Mean/%	SD
Age	37.14	10.82	37.2	11.95	36.76	19.73
Male	63.10%		62.90%		64%	
Female	35.20%		35.10%		36%	
Non-binary	0.60%		0.70%		0%	
Prefer not to respond	1.1%		1.30%		0%	
White	80.70%		82.10%		72%	
Black or African American	8%		6%		20%	
Asian	5.10%		4.60%		8%	
American Indian or Alaska Native	2.30%		2.60%		0%	
Other	1.10%		1.30%		0%	
Two or more races	2.80%		3.30%		0%	
Spanish, Hispanic, or Latin	15.90%		13.20%		32%	
College Education	78.40%		78.80%		76%	

Political View	3.89	2.06	3.91	2.14	3.76	2.39
-----------------------	------	------	------	------	------	------

Supplemental Table 15

Experiment 2 Demographic Variables

	Total Sample (n=178)		Consequentialists (n=121)		Deontologists (n=57)	
	Mean/%	SD	Mean/%	SD	Mean/%	SD
Age	35.97	9.77	35.02	9.39	37.96	11.27
Male	70.2%		74.4%		61.4%	
Female	29.2%		24.8%		38.6%	
Prefer not to respond	0.6%		0.8%		0%	
White	58.4%		57.9%		59.6%	
Black or African American	32.6%		38.8%		19.3%	
Asian	6.2%		2.5%		14.0%	
American Indian or Alaska Native	0.6%		0.8%		0%	
Other	0.6%		0%		1.8%	
Two or more races	1.7%		0%		5.3%	
Spanish, Hispanic, or Latinx	30.9%		35.5%		21.1%	
College Education	89.3%		92.6%		82.5%	
Political View	4.61	1.98	4.88	1.98	4.04	1.88

B. Experiment 1

1. Baseline (BL)

You are standing at the point where one train track splits into two tracks. A runaway train is approaching. On the right track are five rail workers. The brakes have stopped working, so the train can't stop. Next to you is a button that switches tracks.

If you do nothing, the train will hit the five rail workers, who will die.

However, if you press the button, the train will switch to the left track. On the left track is one rail worker. The train will hit the rail worker, who will die, but the train will stop and the five rail workers on the right track will live.

2. Partial Ignorance Both Tracks (PIBT)

You are standing at the point where one train track splits into two tracks. A runaway train is approaching. There are some rail workers on the right tracks, but you don't know how many. The brakes have stopped working, so the train can't stop. Next to you is a button that switches tracks.

If you do nothing, the train will continue down the right track and hit any rail workers on the track, who will die.

However, if you press the button, the train will switch to the left track. On the left track is at least one rail worker, but maybe more. The train will hit whoever is on the left track, who will die, but the train will stop and the rail workers on the right track will live.

3. **Complete Ignorance (CI)**

You are standing at the point where one train track splits into two tracks. A runaway train is approaching. You don't know whether there are any rail workers on the right track, or, if there are, how many. You have no way of finding out. The brakes have stopped working, so the train can't stop. Next to you is a button that switches tracks.

If you do nothing, the train will continue down the right track and hit any rail workers on the track, who will die.

However, if you press the button, the train will switch to the left track. You don't know whether there are any rail workers on the left track, or, if there are, how many. The train will hit whoever is on the left track, who will die, but the train will stop and any rail workers on the right track will live.

4. **Partial Ignorance Switch Track (PISwT)**

You are standing at the point where one train track splits into two tracks. A runaway train is approaching. There are six rail workers on the right track. The brakes have stopped working, so the train can't stop. Next to you is a button that switches tracks.

If you do nothing, the train will hit the six rail workers, who will die.

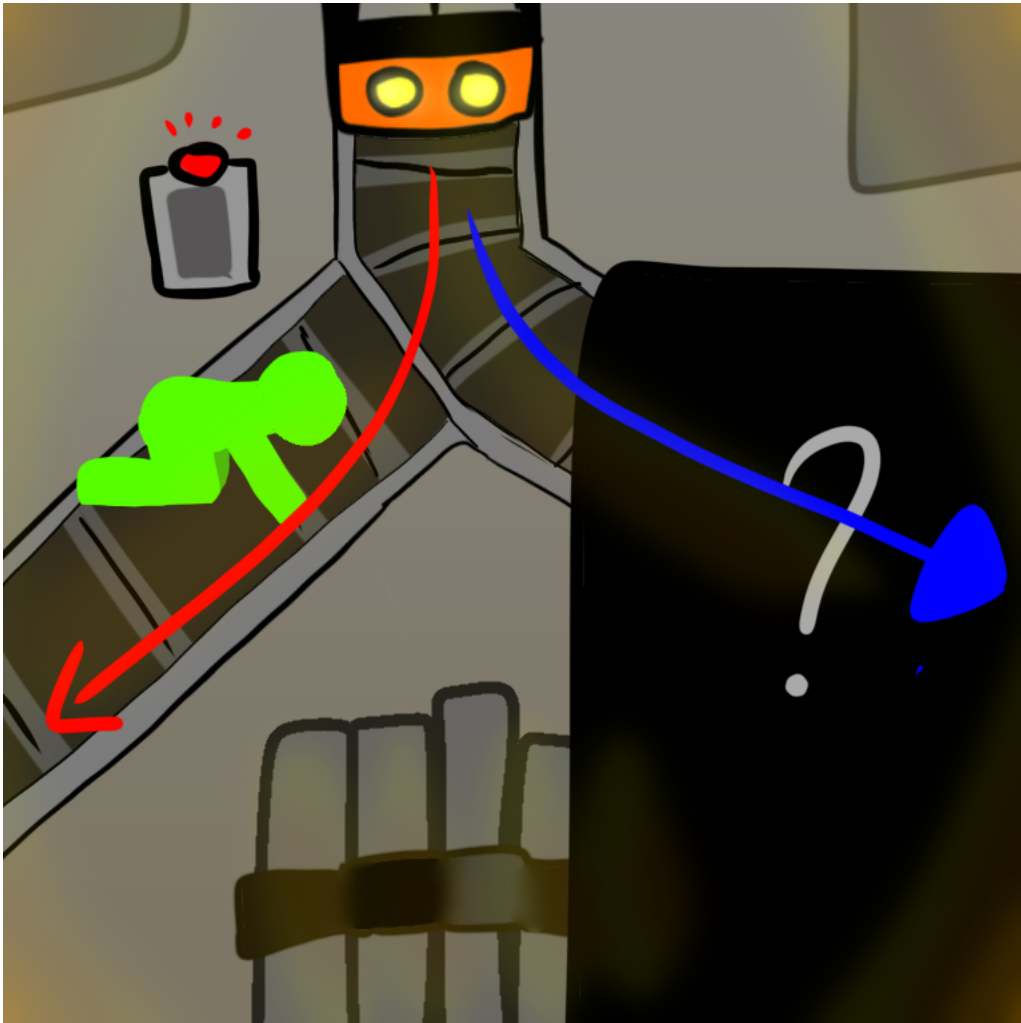
However, if you press the button, the train will switch to the left track. You don't know whether there are any rail workers on the left track, or, if there are, how many. You have no way of finding out. The train will hit whoever is on the left track, who will die, but the train will stop and the rail workers on the right track will live.

5. **Partial Ignorance Stay Track (PIStT)**

You are standing at the point where one train track splits into two tracks. A runaway train is approaching. You don't know whether there are any rail workers on the right track, or, if there are, how many. You have no way of finding out. The brakes have stopped working, so the train can't stop. Next to you is a button that switches tracks.

If you do nothing, the train will continue down the right track and hit any rail workers on the track, who will die.

However, if you press the button, the train will switch to the left track. There is one rail worker on the track. The train will hit the rail worker on the left track, who will die, but the train will stop and the rail workers on the right track, if any, will live.



6. Surmountable Ignorance Switch Track (SISwT)

You are standing at the point where one train track splits into two tracks. A runaway train is approaching. There are six rail workers on the right track. The brakes have stopped working, so the train can't stop. Next to you is a button that switches tracks.

If you do nothing, the train will continue down the right track and hit the rail workers, who will die.

However, if you press the button, the train will switch to the left track. You don't know whether there are any rail workers on the track, or, if there are, how many. The train will hit whoever is on the left track, who will die, but the train will stop and the rail workers on the right track will live.

You can learn whether and, if so, how many rail workers are on the left track by solving this puzzle:

Puzzle 1: First, find the object that fits in the sequence. Then, add the numbers associated with the shaded regions of that object. That's the number of rail workers on the left track, who will die if you press the button.

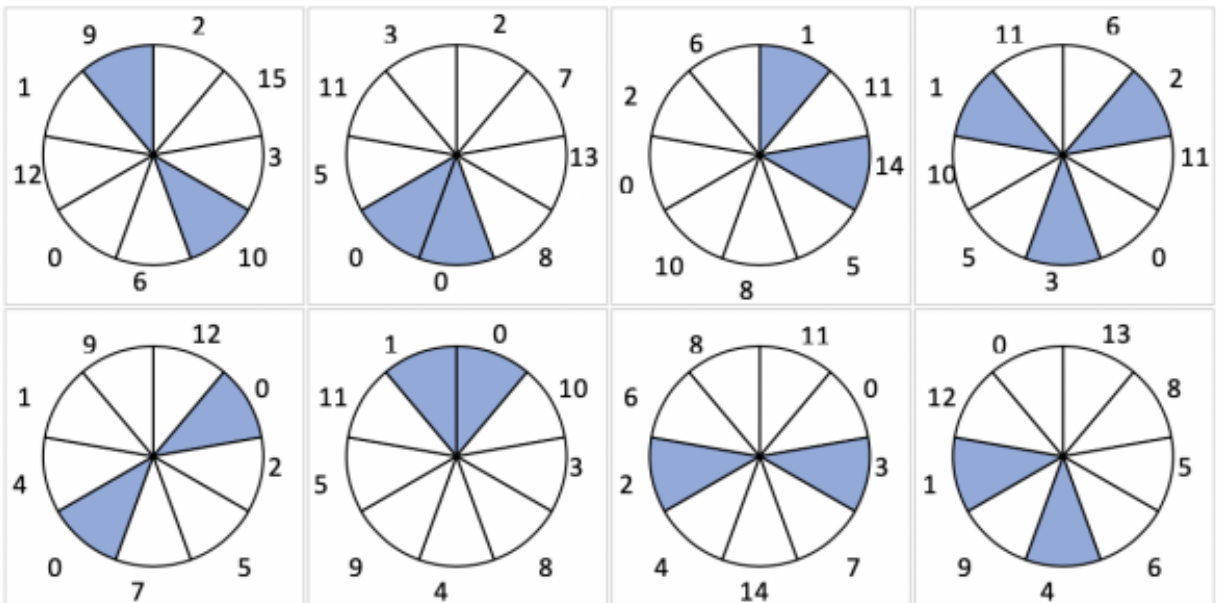
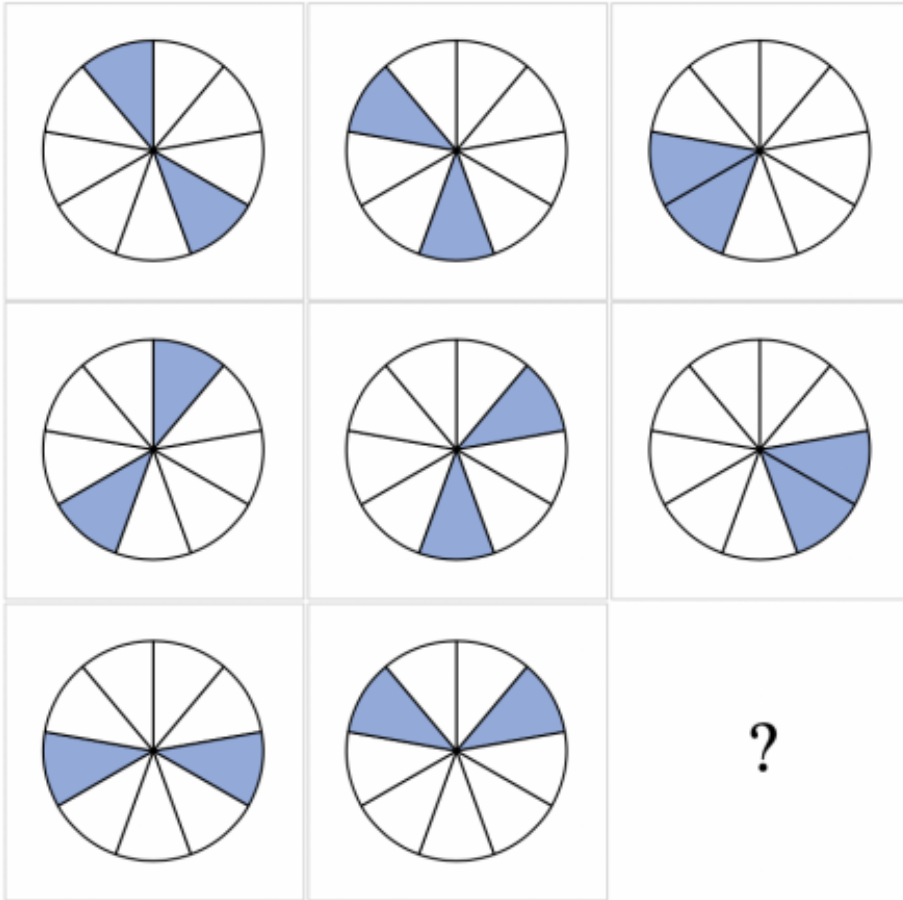
7. Surmountable Ignorance Stay Track (SIStT)

You are standing at the point where one train track splits into two tracks. A runaway train is approaching. You don't know whether there are any rail workers on the right track, or, if there are, how many. The brakes have stopped working, so the train can't stop. Next to you is a button that switches tracks.

You can learn the number of rail workers on the right track by solving the puzzle below

If you do nothing, the train will continue down the right track and hit any rail workers on the track, who will die.

However, if you press the button, the train will switch to the left track. There is one rail worker on the track. The train will hit the rail worker on the left track, who will die, but the train will stop and any rail workers on the right track will live.



Puzzle 2: First, find the object that fits in the sequence. Then, add the numbers associated with the shaded regions of that object. That's the number of rail workers on the right track, who will die if you do not press the button.

8. Know How Ignorance (KHI)

You are standing at the point where one train track splits into two tracks. A runaway train is approaching. On the right track are five rail workers. The brakes have stopped working, so the train can't stop. Next to you is a button that switches tracks.

If you do nothing, the train will hit the five rail workers, who will die.

However, if you press the button, the train will switch to the left track. On the left track is one rail worker. The train will hit the rail worker, who will die, but the train will stop and the five rail workers on the right track will live.

You notice that the electrical panel below the button has been tampered with, and that the wires are all unplugged. The inside of the panel has instructions for re-wiring the button. For the button to work, you need to correctly re-wire the button.

There are exactly eight wires--Red, Orange, Yellow, Green, Blue, Indigo, Violet, and White—that must be connected in the correct order. No wire connects more than once, and no two wires connect to the same spot. The following conditions apply:

Green and Violet both connect before Yellow.

Indigo connects before Blue.

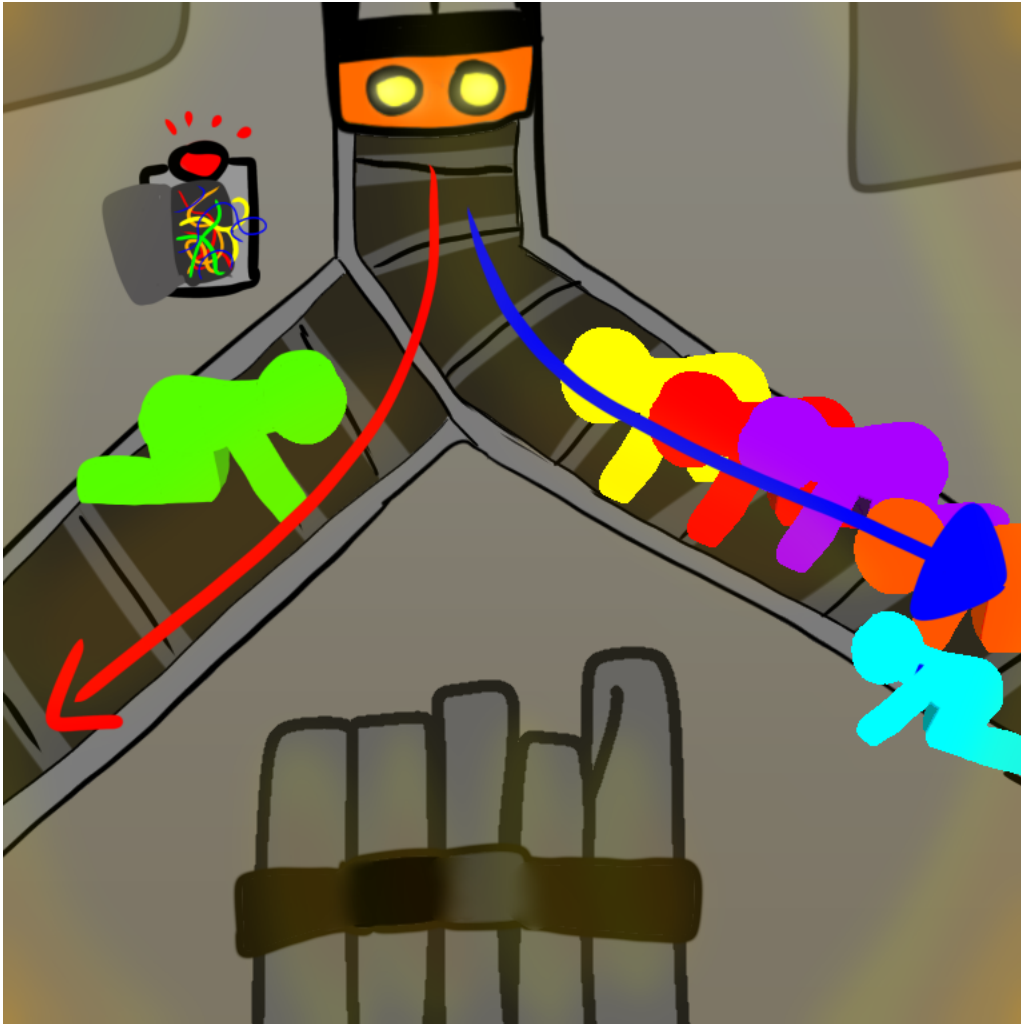
Violet connects before Orange.

White connects after Indigo

Red connects before White but after Yellow.

Which of the following could be the order of the connections from first to last?

- (A) Green, Violet, Yellow, Red, Blue, Indigo, White, Orange
- (B) Violet, Orange, Red, Green, Indigo, Yellow, White, Blue
- (C) Green, Violet, Orange, Yellow, Red, White, Indigo, Blue
- (D) Violet, Green, Orange, Yellow, Indigo, Red, White, Blue
- (E) Violet, Orange, Yellow, Indigo, Green, Blue, Red, White



9. Comprehension

You are standing at the point where one train track splits into two tracks. A runaway train is approaching. On the right track is a banana. The brakes have stopped working, so the train can't stop. Next to you is a button that switches tracks.

If you do nothing, the train will hit the banana, smushing it.

However, if you press the button, the train will switch to the left track. On the left track is an apple. The train will hit the apple, smushing it.

C. Experiment 2

1. Complete Ignorance-Push (CIP)

An empty runaway trolley is speeding down a set of tracks toward an unknown number of rail workers. There is a footbridge above the tracks in between the runaway trolley and the workers on the tracks. You are standing on this footbridge right next to another rail worker and his heavy equipment. If you do nothing, the trolley will proceed down the tracks and hit any workmen on the tracks, who will die.

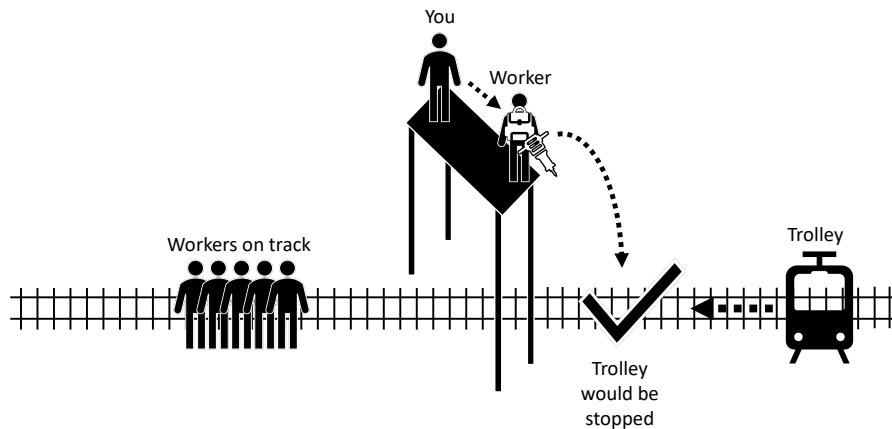
However, it may be possible to avoid the deaths of the workers on the tracks. You are close enough to the rail worker on the bridge, to use your hands to push him and his equipment off. Though this worker will certainly die, there is some likelihood that the combined weight of the worker's body and his equipment will be heavy enough to stop the trolley, but you don't know what this likelihood is. The probability that pushing him off stops the train is more than 0%, but less than 100%.

2. Baseline-Push (BLP)

An empty runaway trolley is speeding down a set of tracks toward five rail workers. There is a footbridge above the tracks in between the runaway trolley and the workers on the tracks. You are standing on this footbridge right next to another rail worker and his heavy equipment. If you do nothing, the trolley will proceed down the tracks and hit the workmen on the tracks, who will die.

However, it is possible to avoid the deaths of the workers on the tracks. You are close enough to the rail worker on the bridge, to use your hands to push him and his equipment off. Though this worker will certainly die, the combined weight of the worker's body and his equipment will be heavy enough to stop the trolley.

Baseline-Push

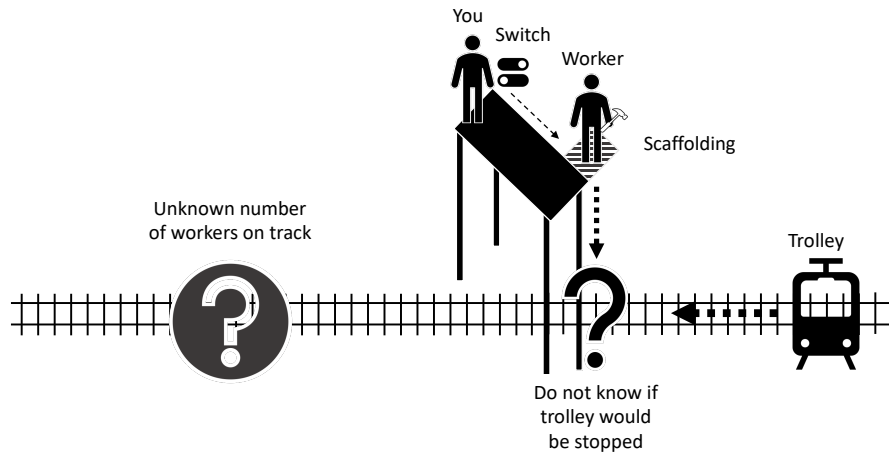


3. Complete Ignorance-Button (CIB)

An empty runaway trolley is speeding down a set of tracks toward an unknown number of rail workers. There is a footbridge above the tracks in between the runaway trolley and the workers on the tracks. You are standing on this footbridge. Standing on some scaffolding attached to the bridge is another rail worker and his heavy equipment. If you do nothing, the trolley will proceed down the tracks and hit any workmen on the tracks, who will die.

However, it may be possible to avoid the deaths of any workers that are on the tracks. The switch to release the scaffolding that the other rail worker is standing on is right next to you. If you push the switch, the rail worker and his equipment will fall onto the tracks. Though this worker will certainly die, there is some likelihood that the combined weight of the worker's body and his equipment will be heavy enough to stop the trolley, but you don't know what this likelihood is. The probability that pushing the switch stops the train is more than 0%, but less than 100%.

Complete Ignorance-Button



4. Baseline-Button (BLB)

An empty runaway trolley is speeding down a set of tracks toward five rail workers. There is a footbridge above the tracks in between the runaway trolley and the workers on the tracks. You are standing on this footbridge. Standing on some scaffolding attached to this bridge is another rail worker and his heavy equipment. If you do nothing, the trolley will proceed down the tracks and hit the workmen on the tracks, who will die.

However, it is possible to avoid the deaths of the workers on the tracks. The switch to release the scaffolding that the other rail worker is standing on is right next to you. If you push the switch, the rail worker and his equipment will fall onto the tracks. Though this worker will certainly die, the combined weight of the worker's body and his equipment will be heavy enough to stop the trolley.

5. Partial Ignorance Track Workers-Push (PITW-P)

An empty runaway trolley is speeding down a set of tracks toward an unknown number of rail workers. There is a footbridge above the tracks in between the runaway trolley and the workers on the tracks. You are standing on this footbridge right next to another rail worker and his heavy equipment. If you do nothing, the trolley will proceed down the tracks and hit any workmen on the tracks, who will die.

However, it is possible to avoid the deaths of any workers that are on the tracks. You are close enough to the rail worker on the bridge, to use your hands to push him and his equipment off. Though this worker will certainly die, the combined weight of the worker's body and his equipment will be heavy enough to stop the trolley.

6. Partial Ignorance Track Workers-Button (PITW-B)

An empty runaway trolley is speeding down a set of tracks toward an unknown number of rail workers. There is a footbridge above the tracks in between the runaway trolley and the workers on the tracks. You are standing on this footbridge. Standing on some scaffolding attached to this bridge is another rail worker and his heavy equipment. If you do nothing, the trolley will proceed down the tracks and hit any workmen on the tracks, who will die.

However, it is possible to avoid the deaths of any workers that are on the tracks. The switch to release the scaffolding that the other rail worker is standing on is right next to you. If you push the switch, the rail worker and his equipment will fall onto the tracks. Though this worker will certainly die, the combined weight of the worker's body and his equipment will be heavy enough to stop the trolley.

7. Partial Ignorance Footbridge Worker-Button (PIFW-B)

An empty runaway trolley is speeding down a set of tracks toward five rail workers. There is a footbridge above the tracks in between the runaway trolley and the workers on the tracks. You are standing on this footbridge. Standing on some scaffolding attached to this bridge is another rail worker and his heavy equipment. If you do nothing, the trolley will proceed down the tracks and hit the workmen on the tracks, who will die.

However, it may be possible to avoid the deaths of the workers on the tracks. The switch to release the scaffolding that the other rail worker is standing on is right next to you. If you push the switch, the rail worker and his equipment will fall onto the tracks. Though this worker will certainly die, there is some likelihood that the combined weight of the worker's body and his equipment will be heavy enough to stop the trolley, but you don't know what this likelihood is. The probability that pushing the switch stops the train is more than 0%, but less than 100%.

8. Partial Ignorance Footbridge Worker-Push (PIFW-P)

An empty runaway trolley is speeding down a set of tracks toward five rail workers. There is a footbridge above the tracks in between the runaway trolley and the workers on the tracks. You are standing on this footbridge right next to another rail worker and his heavy equipment. If you do nothing, the trolley will proceed down the tracks and hit the workmen on the tracks, who will die.

However, it may be possible to avoid the deaths of the workers on the tracks. You are close enough to the rail worker on the bridge, to use your hands to push him and his equipment off. Though this worker will certainly die, there is some likelihood that the combined weight of the worker's body and his equipment will be heavy enough to stop the trolley, but you don't know what this likelihood is. The probability that pushing him off stops the train is more than 0%, but less than 100%.

