

Supporting Materials for

**Crossing the Line:
On the Distinction Between Harm Aversion and Harm
Avoidance in Moral Choice**

Contents	Page
I. Scenario Wordings	2
II. Pre-Screen & Attention Check Questions	8
III. Study S1: Ruling out an Identifiable Victim Effect as an Alternative Explanation	10
IV. Study S2: Increasing vs. Diminishing Marginal Harm Conceptual Replicate – Abortion Study	13
V. Study S3: Increasing vs. Diminishing Marginal Harm Conceptual Replicate – Life Support Study	16
VI. Study S4: Ruling out Value Uncertainty as an Alternate Explanation	18
VII. References	22

Scenario Wordings

Introduction Study

Avoidable Harm condition:

Consider the following two proposals, A and B. Both proposals are aimed at creating farmland to feed hungry families in Africa:

- Proposal A is to build 1 acre of farmland for an African village. No trees will need to be cut down to create this farmland. This farmland will provide enough food to feed 100 hungry families.
- Proposal B is to build 5 acres of farmland for an African village. A total of one acre of scarce tropical rainforest trees will need to be cut down to create these 5 acres of farmland. This farmland will provide enough food to feed 500 hungry families.

Which proposal would you implement if you had the choice?

- Proposal A: Feed 100 families. Cut down 0 trees.
- Proposal B: Feed 500 families. Cut down 1 acre of tropical rainforest trees.

Unavoidable Harm condition:

Consider the following two proposals, A and B. Both proposals are aimed at creating farmland to feed hungry families in Africa:

- Proposal A is to build 1 acre of farmland for an African village. A total of one acre of scarce tropical rainforest trees will need to be cut down to create this 1 acre of farmland. This farmland will provide enough food to feed 100 hungry families
- Proposal B is to build 5 acres of farmland for an African village. A total of two acres of scarce tropical rainforest trees will need to be cut down to create these 5 acres of farmland. This farmland will provide enough food to feed 500 hungry families.

Which proposal would you implement if you had the choice?

- Proposal A: Feed 100 families. Cut down 1 acre of tropical rainforest trees.
- Proposal B: Feed 500 families. Cut down 2 acres of tropical rainforest trees.

Study 1

Avoidable Harm condition:

Imagine that you are a doctor working in a research hospital. You are faced with the following decision:

Option A: A child is terminally ill and on life support and is expected to die in about a week. You can pull life support from the child which will save \$60,000 that will be used to spend on cancer research.

Option B: You can choose not to pull life support from the child. You will not save any money for cancer research.

Which would you choose to do?

- Pull life support from 1 child and save \$60,000 for cancer research.
- Do not pull life support from 1 child. Save no money for cancer research.

Unavoidable Harm condition:

Imagine that you are a doctor working in a research hospital. You are faced with the following decision:

Option A: A child is terminally ill and on life support and is expected to die in about a week. You can pull life support from the child which will save \$800 that will be used to spend on cancer research.

Option B: Two children are terminally ill and on life support and are each expected to die in about a week. You can pull life support from both children which will save \$50,000 that will be used to spend on cancer research.

Which would you choose to do?

- Pull life support from 1 child and save \$800 for cancer research
- Pull life support from 2 children and save \$50,000 for cancer research

Study 2

All conditions:

Are you a registered member of one of the following political parties?

- Democratic Party
- Republican Party
- Neither

—Page Break—

As additional compensation for your participation in this survey, you have the opportunity to make a donation decision. You will be entered into a raffle, and if you win the raffle, we will donate money in accordance with your decision.

The decision will determine how much money will be donated to the following two nonprofits:

- One dedicated to supporting the Democratic nominee in the upcoming presidential election (Priorities USA)
- One dedicated to supporting the Republican nominee in the upcoming presidential election (WinRed)

Which option would you choose?

Avoidable Harm condition

Democratic Party Participant Choice Set

- Donate \$2.00 to the Democratic nonprofit, and donate \$0 to the Republican nonprofit
- Donate \$4.00 to the Democratic nonprofit, and donate \$1.00 to the Republican nonprofit

Republican Party Participant Choice Set

- Donate \$2.00 to the Republican nonprofit, and donate \$0 to the Democratic nonprofit
- Donate \$4.00 to the Republican nonprofit, and donate \$1.00 to the Democratic nonprofit

Unavoidable Harm condition:

Democratic Party Participant Choice Set

- Donate \$2.00 to the Democratic nonprofit, and donate \$2.00 to the Republican nonprofit
- Donate \$4.00 to the Democratic nonprofit, and donate \$3.00 to the Republican nonprofit

Republican Party Participant Choice Set

- Donate \$2.00 to the Republican nonprofit, and donate \$2.00 to the Democratic nonprofit
- Donate \$4.00 to the Republican nonprofit, and donate \$3.00 to the Democratic nonprofit

Study 3

Avoidable Harm condition:

There is a charity that provides lunches to children in Africa. Consider two possible courses of action that this charity is considering:

Proposal A is to start using Nutrient A in their farming, which will allow them to grow enough additional crops to provide lunch to 1 million children in Africa who would otherwise go without lunch. Nutrient A is safe for humans, and will not have any impact on the ozone layer.

Proposal B is to start using Nutrient B in their farming, which will allow them to grow enough additional crops to provide lunch to 4 million children in Africa who would otherwise go without lunch. Nutrient B is safe for humans, but the charity's use of Nutrient B will cause a 0.5-inch hole in the ozone layer.

Which proposal would you rather be implemented?

- Proposal A: Provide lunches to 1 million hungry children. Create no hole in the ozone layer.
- Proposal B: Provide lunches to 4 million hungry children. Create a 0.5-inch hole in the ozone layer.

Unavoidable Harm condition:

There is a charity that provides lunches to children in Africa. Consider two possible courses of action that this charity is considering:

Proposal A is to start using Nutrient A in their farming, which will allow them to grow enough additional crops to provide lunch to 1 million children in Africa who would otherwise go without lunch. Nutrient A is safe for humans, but the charity's use of Nutrient A will cause a 1-inch hole in the ozone layer.

Proposal B is to start using Nutrient B in their farming, which will allow them to grow enough additional crops to provide lunch to 4 million children in Africa who would otherwise go without lunch. Nutrient B is safe for humans, but the charity's use of Nutrient B will cause a 1.5-inch hole in the ozone layer.

Which proposal would you rather be implemented?

- Proposal A: Provide lunches to 1 million hungry children. Create a 1-inch hole in the ozone layer.
- Proposal B: Provide lunches to 4 million hungry children. Create a 1.5-inch hole in the ozone layer.

Study 4

All conditions:

A trolley without passengers and without a conductor is traveling at full speed down a track. On the track there are eleven people, who will all surely be killed if nothing changes.

Avoidable Harm condition:

There is one person standing on the side of the track. If you push this one other person onto the track (which will cause that person to die), you know that you'll slow down the trolley enough to save three of the people on the track.

Which would you choose to do?

- Do not push the person onto the track.
- Push the person onto the track. This person will die, but you will save three people's lives.

Diminishing Marginal Returns to Harm condition:

There are two people standing on the side of the track. If you push one person onto the track (which will cause that person to die), you know that you'll slow down the trolley enough to save three of the people on the track. If you push the second person onto the track (which will cause the second person to die), you know that you'll slow down the trolley enough to save five people on the track.

Which would you choose to do?

- Push one person onto the track. This person will die, but you will save three people's lives.
- Push both people onto the track. They both will die, but you will save five people's lives.

Increasing Marginal Returns to Harm condition:

There are two people standing on the side of the track. If you push one person onto the track (which will cause that person to die), you know that you'll slow down the trolley enough to save two of the people on the track. If you push the second person onto the track (which will cause the second person to die), you know that you'll slow down the trolley enough to save five people on the track.

Which would you choose to do?

- Push one person onto the track. This person will die, but you will save two people's lives.
- Push both people onto the track. They both will die, but you will save five people's lives.

Study 5***Avoidable Harm condition:***

Consider the following two proposals, A and B. Both proposals are aimed at building water treatment facilities for people in Africa:

- Proposal A is to build 1 water treatment facility for an African village. No trees will need to be cut down to create this 1 facility. This facility will provide in-home water to people in 1 village who would otherwise need to walk one mile to get clean water.
- Proposal B is to build 5 water treatment facilities - each facility will provide in-home water to people in different African villages. A total of one acre of scarce tropical rainforest trees will need to be cut down to create these 5 facilities. These 5 facilities will provide in-home water to people in 5 different villages who would otherwise need to walk one mile to get clean water.

Which proposal would you rather be implemented?

- Proposal A: Build 1 facility. Cut down 0 trees.
- Proposal B: Build 5 facilities. Cut down 1 acre of scarce tropical rainforest trees.

Unavoidable Harm condition:

Consider the following two proposals, X and Y. Both proposals are aimed at building water treatment facilities for people in Africa:

- Proposal X is to build 1 water treatment facility for an African village. A total of one acre of scarce tropical rainforest trees will need to be cut down to create this 1 facility. This facility will provide in-home water to people in 1 village who would otherwise need to walk one mile to get clean water.
- Proposal Y is to build 5 water treatment facilities - each facility will provide in-home water to people in different African villages. A total of two acres of scarce tropical rainforest trees will need to be cut down to create these 5 facilities. These 5 facilities will provide in-home water to people in 5 different villages who would otherwise need to walk one mile to get clean water.

Which proposal would you rather be implemented?

- Proposal X: Build 1 facility. Cut down 1 acre of scarce tropical rainforest trees.
- Proposal Y: Build 5 facilities. Cut down 2 acres of scarce tropical rainforest trees.

Pre-Screen & Attention Check Questions

Introduction Study

Pre-Screen Question: Participants were first asked to write down the third word in the following statement: “A rolling stone gathers no moss.” Those who did not answer this question correctly were not allowed to participate in the study.

Attention Check: At the end of the survey, participants were asked the following attention check question: “What was the goal of the programs you were asked to evaluate in this study?”

- Provide education for children
- Build shelter for the homeless
- Build farmland for hungry families
- Provide medicine for the sick

Study 1

Pre-Screen Question: Participants were first asked to write down the third word in the following statement: “A rolling stone gathers no moss.” Those who did not answer this question correctly were not allowed to participate in the study.

Attention Check: At the end of the survey, participants were asked the following attention check question: “Which of the following options were you presented with?”

- Do not pull life support for 1 child. Save no money for cancer research.
- Pull life support from 2 children and save \$50,000 for cancer research

Study 2

Pre-Screen Question: Participants were first asked to write down the third word in the following statement: “A rolling stone gathers no moss.” Those who did not answer this question correctly were not allowed to participate in the study.

Pre-Screen Question: Participants were then asked to identify which political party they were registered with: “Democratic Party” “Republican Party” or “Neither.” Those who selected “Neither” were told they did not qualify for the study.

Attention Check: At the end of the survey, participants were asked the following attention check question: “What type of non-profits did you make a decision about in this survey?”

- Ones that support the homeless
- Ones that supports political parties
- Ones that support the environment
- Ones that support the arts

Study 3

Pre-Screen Question: Participants were first asked to write down the third word in the following statement: “A rolling stone gathers no moss.” Those who did not answer this question correctly were not allowed to participate in the study.

Attention Check: At the end of the survey, participants were asked the following attention check question: “What was the goal of the programs you were asked to evaluate in this study?”

- Provide medicine for children in Africa
- Provide homes for children in Africa
- Provide lunches to children in Africa
- Provide child care for children in Africa

Study 4

Pre-Screen Question: Participants were first asked to write down the third word in the following statement: “A rolling stone gathers no moss.” Those who did not answer this question correctly were not allowed to participate in the study.

Attention Check: At the end of the survey, participants were asked the following attention check question: “Which of the following decisions were you presented with?”

- Whether to push 1 person or no one onto a trolley track
- Whether to push 1 person or two people onto a trolley track

Study 5

Pre-Screen Question: Participants were first asked to write down the third word in the following statement: “A rolling stone gathers no moss.” Those who did not answer this question correctly were not allowed to participate in the study.

Attention Check: At the end of the survey, participants were asked the following attention check question: “What was the goal of the programs you were asked to evaluate in this study?”

- Building water treatment facilities
- Building hospitals
- Building police stations

Study S1: Ruling Out an Identifiable Victim Effect as an Alternative Explanation

We conducted this study for two reasons. First, we examine whether the results from Study 1 are due to an identifiable victim effect (Kogut & Ritov, 2011). Specifically, Study 1 found that participants were averse to pulling life support from a single child when they could avoid pulling life support altogether. However, they were more likely to pull life support from two children than from a single child. In this study, we examined whether an aversion to pulling life support from specifically one child drove Study 1's results.

Specifically, this study included two different *Avoidable Harm* conditions as well as the *Unavoidable Harm* condition: An *Avoidable Harm – Single Child* condition (in which participants chose between removing life support from one child to save \$60,000 for cancer research versus not doing so), an *Avoidable Harm – Two Children* condition (in which participants chose between removing life support from two children to save \$60,000 for cancer research versus not doing so), and the *Unavoidable Harm* condition (in which participants chose between removing life support from one child to save \$800 for cancer research versus removing life support from two children to save \$60,000 for cancer research). We expected that participants would be more willing to commit greater harm for greater benefits in the *Unavoidable Harm* condition relative to each of the two *Avoidable Harm* conditions.

Second, we also examined whether Study 1's results could have been due to an idiosyncratic preference for the greater harm/greater benefit option available in the *Unavoidable Harm* condition. If this were the case, then participants would be more likely to choose to pull life support from two children to save \$60,000 rather than to not pull life support at all. Note that the *Avoidable Harm – Two Children* condition presents participants with this exact choice. Thus, if an idiosyncratic preference for pulling life support from two children to save \$60,000 drove

our results, then we would expect that participants would be willing to do so even if they could avoid committing any harm. However, we expect that participants faced with this decision would still be averse to committing any harm in this condition.

Method

We preregistered the design, sample size, and analysis plan for this study on AsPredicted.org (<https://aspredicted.org/blind.php?x=se8pk5>). We aimed to recruit 600 participants from Mechanical Turk and recruited a total of 615 participants. Per our pre-registration plan, we removed 53 participants who failed an attention check resulting in a final sample of 562 participants (mean age = 38; 57% female).

Participants were randomly assigned to either an *Avoidable Harm – Single Child* condition, an *Avoidable Harm – Two Children* condition, or an *Unavoidable Harm* condition. All participants assumed the role of a doctor who faced a decision regarding whether or not to pull life support to save money that would be used for cancer research. In the *Avoidable Harm – Single Child* condition, participants could pull life support for one dying child in order to save \$60,000 for cancer research or choose not to do so. In the *Avoidable Harm – Two Children* condition, participants could pull life support for two dying children in order to save \$60,000 for cancer research or choose not to do so. In the *Unavoidable Harm* condition, participants instead faced a decision about whether to pull life support for one dying child in order to save \$800 for cancer research or to pull life support for two dying children in order to save \$60,000 for cancer research.

Results

Consistent with our pre-registered hypothesis, participants were more likely to prefer to commit greater harm in exchange for greater benefits in the *Unavoidable Harm* condition (63.8%) than in the *Avoidable Harm – Single Child* condition (42.6%), $\chi^2(df = 1, N = 368) = 16.55, p < .001, \phi = .21$ and the *Avoidable Harm – Two Children* condition (27.8%), $\chi^2(df = 1, N = 379) = 50.55, p < .001, \phi = .36$. Moreover, a significantly greater proportion of participants chose to commit greater harm for greater benefit in the *Avoidable Harm – Single Child* condition than in the *Avoidable Harm – Two Children* condition, $\chi^2(df = 1, N = 377) = 9.05, p = .003, \phi = .16$. Thus, it does not appear that a particular aversion to pulling life support from a single identifiable child could explain participants reluctance to do so. These results are also inconsistent with the possibility that an idiosyncratic preference for pulling life support from two children to save \$60,000 drove Study 1's results.

**Study S2:
Increasing versus Diminishing Marginal Returns to Harm
Conceptual Replicate – Abortion Study**

Method

We pre-registered the design, sample size, and analysis plan for this study on AsPredicted (<https://aspredicted.org/blind.php?x=vb83wh>). We aimed to recruit 500 participants from Mechanical Turk and collected a total of 510 participants. Per our pre-registration plan, we removed 51 participants who failed an attention check, resulting in a final sample of 459 participants (mean age = 36; 56.2% female).

Participants first indicated their stance in the abortion debate by selecting whether they were Pro-Choice, Pro-Life or had no opinion. Per our pre-registration plan, we restricted participation to those who were either Pro-Choice (79.5%) or Pro-Life (20.5%). All participants faced a decision regarding the allocation of money to two non-profits. One of these non-profits was associated with supporting Pro-Choice policies (Planned Parenthood), while the other was associated with Pro-Life policies (Pro-Life Across America).

Participants were then randomly assigned to one of three between-subjects conditions. In all three conditions, participants faced the decision of how to allocate money to these two non-profits. We operationalized benefits as a contribution to the non-profit that was consistent with their stance and harm as a contribution to the non-profit that was inconsistent with their stance. Thus, across conditions, participants faced a conflict between choosing to maximize donations to support their own stance in the abortion debate (i.e., maximize benefits) versus to minimize donations to support the opposing stance in the abortion debate (i.e., minimize harm).

Specifically, in the *Avoidable Harm* condition, participants faced the decision to either (a) donate \$22 to the non-profit consistent with their stance on abortion and \$0 to the non-profit with

the opposing stance or (b) donate \$44 to the non-profit consistent with their stance on abortion and \$12 to the non-profit with the opposing stance. In the *Diminishing Marginal Returns to Harm* condition, participants faced the decision to either (a) donate \$22 to the non-profit consistent with their stance on abortion and \$12 to the non-profit with the opposing stance or (b) donate \$40 to the non-profit consistent with their stance on abortion and \$24 to the non-profit with the opposing stance. Thus, the benefit to harm ratio was *more attractive* for the former option (for every \$1.83 donated to the value-consistent option, the value-inconsistent option received \$1) than the latter option (for every \$1.67 donated to the value-consistent option, the value-inconsistent option received \$1). Finally, in the *Increasing Marginal Returns to Harm* condition, participants faced the decision to either (a) donate \$22 to the non-profit consistent with their stance on abortion and \$12 to the non-profit with the opposing stance or (b) donate \$44 to the non-profit consistent with their stance on abortion and \$20 to the non-profit with the opposing stance. Thus, both options required committing harmful acts. Moreover, the benefit to harm ratio of money was *less attractive* for the former option (for every \$1.83 donated to the value-consistent option, the value-inconsistent option received \$1) than the latter option (for every \$2.20 donated to the value-consistent option, the value-inconsistent option received \$1).

Before making their selection, all participants learned that in addition to their participation payment, they will be entered into a raffle. If they won the raffle, a donation would be made on their behalf in accordance with their selection in the survey.

Finally, after making their selection, participants completed a manipulation check: They evaluated the extent to which a donation to each non-profit would be beneficial or harmful to society on a seven-point scale (-3 = “extremely harmful”, 0 = “neither beneficial nor harmful”, +3 = “extremely beneficial”).

Results

Manipulation Check: Consistent with the intent of the manipulation, participants evaluated a donation to support the non-profit consistent with their own stance on abortion to be more beneficial than harmful to society ($M = +2.02$, $SD = 1.02$; one-sample $t = 35.26$, $p < .001$), whereas a donation to the non-profit inconsistent with their own stance on abortion to be more harmful than beneficial to society ($M = -0.87$, $SD = 1.91$; one-sample $t = -9.73$, $p < .001$).

Monetary Allocation Decisions: Consistent with our pre-registered hypothesis, participants were less likely to commit greater harm for greater benefits in the *Avoidable Harm* condition (26.6%) than in both the *Diminishing Marginal Returns to Harm* condition (40.8%; $\chi^2(1) = 7.01$, $p = .008$; $\phi = .15$) and the *Increasing Marginal Returns to Harm* condition (57.7%; $\chi^2(1) = 30.58$, $p < .001$; $\phi = .32$). Additionally, participants were more likely to commit greater harm for greater benefits in the *Increasing Marginal Returns to Harm* condition than in the *Diminishing Marginal Returns to Harm* condition, $\chi^2(1) = 8.63$, $p = .003$; $\phi = .17$.

Study S3:
Increasing versus Diminishing Marginal Returns to Harm
Conceptual Replicate – Life Support Study

Method

We recruited 300 participants from Mechanical Turk (mean age = 36; 61% female). All participants assumed the role of a doctor facing a decision about the magnitude of harm they were willing to inflict in order to save money for medical research. In the *Avoidable Harm* condition, participants faced a decision about whether to pull life support for one dying child in order to save \$60,000 for cancer research or to do nothing. In the *Diminishing Marginal Returns to Harm* condition, participants faced a decision about whether to pull life support for one dying child in order to save \$40,000 for cancer research, or to pull life support for two dying children in order to save \$60,800 for cancer research. Thus, the benefit to harm ratio of pulling life support for two children (\$30,400 saved per child) was less favorable than the benefit to harm ratio of pulling life support for a single child (\$40,000 saved per child). Finally, in the *Increasing Marginal Returns to Harm* condition, participants faced a decision about whether to pull life support for one dying child in order to save \$800 for cancer research or two dying children in order to save \$60,800 for cancer research. Thus, in this condition, the benefit to harm ratio of pulling life support for two children (\$30,400 saved per child) was more favorable than the benefit to harm ratio of pulling life support for a single child (\$800 saved per child).

Results

Participants were less likely to exchange greater harm for greater benefits in the *Avoidable Harm* condition (26.9%) than in both the *Diminishing Marginal Returns to Harm* condition (51.0%; $\chi^2(1) = 12.26, p < .001; \phi = .25$) and the *Increasing Marginal Returns to*

Harm condition (77.0%; $\chi^2(1) = 51.18, p < .001; \phi = .50$). Additionally, participants were less likely to exchange greater harm for greater benefits in the *Diminishing Marginal Returns to Harm* condition than in the *Increasing Marginal Returns to Harm* condition ($\chi^2(1) = 14.38, p < .001; \phi = .26$).

Study S4: Ruling out Value Uncertainty as an Alternate Explanation

Readers may wonder whether one reason why individuals are reluctant to make value-tradeoffs when it is possible for an actor to avoid committing any harm is that actors are uncertain about the value of harm. In other words, perhaps people in the *Avoidable Harm* conditions are simply uncertain about whether the value-tradeoff that they are presented with represents a good exchange relative to an outside option. In contrast, it could be that the presence of the smaller unavoidable harm in the *Unavoidable Harm* condition supplied participants with context to understand what would be considered a reasonable harm/benefit tradeoff to make.

We believe this presents an unlikely explanation for our findings for three reasons. First, Study 4, Study S2, and Study S3 found that participants continued to be more likely to inflict greater harm when the harm produced diminishing marginal benefits. In other words, even when the value of committing additional harm was defined to be relatively inefficient at producing benefits, participants were still more likely to commit additional harm in order to receive additional benefits relative to when they could avoid harm altogether. In other words, if participants were reluctant to commit harm when all harm is avoidable because they are uncertain about its value, then they ought to similarly avoid exchanging additional harm for additional benefits when they know that doing so represents a relatively poor deal. Yet, we find that the preference to avoid harm when harm is completely avoidable is so strong that people prefer to do so to a greater extent than when they are faced with diminishing marginal returns to committing more harm.

Second, Study 5 employed a within-subjects design, which exposed participants to both the *Avoidable Harm* dilemma and the *Unavoidable Harm* dilemma. As a result, even when

participants evaluated both dilemmas, we still find a preference-inconsistency across the two dilemmas. To evaluate this alternative explanation further, we conducted an additional test to evaluate whether the order of the options interacted with our results. To the extent that the *Unavoidable Harm* condition decreases participants' uncertainty about the value of harm, then seeing this condition first ought to affect how participants respond to the *Avoidable Harm* condition. However, when we ran a binary logistic regression with Scenario Order, Condition and their interaction predicting responses, we found no Scenario Order \times Condition interaction, $B = .42$, $SE = .39$, $p = .29$. Moreover, we replicate our main results when we examine just participants who saw the *Unavoidable Harm* condition first. In this case, those indicating that harming the rainforest violates a protected value were still less likely to maximize harm when harm was avoidable (35.8%) than when harm was unavoidable (53.9%), $z = 4.26$, $p < .001$, $\phi = .24$. Thus, seeing the *Unavoidable Harm* condition first does not eliminate the effect.

Finally, we conducted an additional study to examine whether setting a precedent ahead of being faced with an ethical dilemma increases participants' likelihood of committing harm when it is completely avoidable. Specifically, if participants were uncertain about an appropriate value of harm in the *Avoidable Harm* condition, then highlighting a past decision that set a precedent for an acceptable harm/benefit tradeoff ought to decrease participants' uncertainty with regards to the value of harm. However, if participants hold a strong preference to avoid committing any harm, then even when a precedent is set, participants would still be reluctant to commit harm if it is completely avoidable.

Method

We recruited 314 participants from Mechanical Turk (mean age = 37; 41% male) to complete an experiment in exchange for payment. All participants read an adapted version of the classic trolley problem, in which a trolley without passengers and without a conductor was traveling at full speed down a track towards 11 people who would all die if nothing changed.

In the *No Precedent / Avoidable Harm* condition, participants read that there was one person standing on the side of the track, and that pushing this one person onto the track would cause this person to die but would slow down the train enough to save five people on the track. These participants were then asked whether they would push the person and save five people or not.

In the *Precedent / Avoidable Harm* condition, participants faced the same decision as above. However, before making their decision, participants read that the last time a similar situation occurred, a bystander decided to push one person onto a track (thereby killing this one person) in order to save two other people, thus establishing a precedent of harming one person to save two others.

Finally, in the *Unavoidable Harm* condition, participants read that there were two people standing on the side of the track, and that pushing one of the people onto the track would cause this person to die but would slow down the train enough to save two of the people on the track. They further read that pushing the second person onto the track would also cause this second person to die but would slow down the train enough to save an additional five people on the track. These participants then indicated whether they would push one person onto the track to save two people, or push two people onto the track to save a total of seven people.

In sum, participants' decisions in all conditions altered whether or not one person's life would be sacrificed to save five others. Despite the fact that this tradeoff was objectively

identical in all conditions, we predicted that participants in the *Unavoidable Harm* condition would be more likely to commit greater harm in order to achieve greater benefits than would participants in the two *Avoidable Harm* conditions.

Results

A binary logistic regression of the dummy-coded condition variable on participants' decisions revealed that participants more often sacrificed the one person who they could have spared in the *Unavoidable Harm* condition (63.5%) than in the *No Precedent / Avoidable Harm* condition (31.0%; $b = 1.72$, $Wald = 30.31$, $p < .001$; $\phi = .33$) and the *Precedent / Avoidable Harm* condition (23.8%; $b = 1.36$, $Wald = 21.32$, $p < .001$; $\phi = .40$), the latter of which did not differ ($b = .36$, $Wald = 1.39$, $p = .24$). Thus, these results indicate that a precedent does not increase the likelihood of committing harm when harm is completely avoidable. In conjunction with the previously described studies, these results provide converging evidence inconsistent with the possibility that value uncertainty can explain the preference inconsistency across avoidable and unavoidable harms.

References

- Kogut, T., & Ritov, I. (2011). The identifiable victim effect: Causes and boundary conditions. *The science of giving: Experimental approaches to the study of charity*, 133-148.