

What drives opposition to suicide? Two exploratory studies of normative judgments

Justin F. Landy* Pritika Shah†

Abstract

The act of suicide is commonly viewed as wrong in some sense, but it is not clear why this is. Based on past empirical research and philosophical theorizing, we test ten different explanations for why suicide is opposed on normative grounds. Using a within-subjects design, Study 1 showed that seven out of ten manipulations had significant effects on normative judgments of suicide: time left to live, lack of close social relationships, a history of prior immoral behavior, the manner in which the suicide is committed, painful, incurable medical issues, impulsive decision-making, and the actor's own moral-religious background. However, in all cases, the act of suicide was still considered wrong, overall. Using a between-subjects design, Study 2 tested the combined effect of the seven significant manipulations from Study 1. In combination, the seven manipulations eliminated opposition to suicide, on average. Implications for moral psychology and suicide prevention are discussed.

Keywords: suicide, moral judgment, normative judgment, exploratory methods, strong inference

1 Introduction

The act of suicide poses a conundrum. On the one hand, it is widely considered to be wrong in some way, and great effort and expense is devoted to suicide prevention efforts.

*Department of Psychology and Neuroscience, Nova Southeastern University, 3301 College Avenue, Fort Lauderdale, FL 33314. E-mail: jlandy@nova.edu. <https://orcid.org/0000-0002-4082-9505>.

†Nova Southeastern University, ORCID 0000-0002-7314-7244.

We thank Ria Bajaria, Fatima Bahar, Rose-Arrielle Jerome, and Rebekah Sager for their assistance with this research, and Jonathan Baron, two anonymous reviewers, and the members of the Reasoning, Person-Perception, and Morality Lab for their helpful feedback on this paper.

The two studies reported here represent the entirety of the research we have conducted on this topic. We report how we determined our sample sizes, all data exclusions, all manipulations, and all measures in both studies. Full materials, data, and analysis scripts for both studies are available on the Open Science Framework at <https://osf.io/mjvgq>.

Copyright: © 2021. The authors license this article under the terms of the Creative Commons Attribution 3.0 License.

For instance, the 2020 presidential budget allocated \$19 million to the National Suicide Prevention Lifeline Program and about \$35,427,000 to Youth Suicide Early Intervention and Prevention (NASMHPD, 2020). Yet, on the other hand, it is not clear what underlies this opposition. Suicide is prohibited by most major religious traditions, but, as Hecht (2013, p. 6) aptly puts it, “our [secular] society today has no coherent argument against suicide”.¹ That is, despite the obvious fact that people widely oppose suicide, we do not know very much about why they do so, as very few empirical studies have investigated normative judgments of suicide.²

This dearth of research is somewhat surprising, considering the prevalence and societal costs of suicide. In 2019, suicide was the tenth most common cause of death among Americans, and the second most common cause of death among Americans aged 10–34 years old. Over 47,000 Americans died by their own hand that year, more than twice the number that died by homicide (National Institute of Mental Health, 2021). Thus, suicide is a common occurrence, but research on people’s normative judgments of it is sorely lacking. From a pragmatic, applied perspective, understanding what considerations underlie opposition to suicide may be useful for designing new interventions to discourage it, a possibility to which we return in the General Discussion.

Before reviewing existing empirical work on normative judgments of suicide, it is necessary to define what we mean by the term. Arriving at a coherent definition of suicide turns out to be more difficult than it might appear, and considerable philosophical work has been devoted to the problem (see Fairbairn, 1995; Wreen, 1998). For our purposes, we will sidestep the philosophical issues, and follow the National Institute of Mental Health by defining suicide as “death caused by self-directed injurious behavior with intent to die as a result of the behavior” (National Institute of Mental Health, 2021). Notably, this excludes such edge cases as religious martyrdom (in which the intent is typically not to die, but to avoid violating one’s faith, with death being the only way to do so), so-called altruistic suicide (which is enacted out of a desire to help someone else, with one’s own death being merely a means to this end; e.g., a soldier who throws himself on a grenade to protect others), accidental self-killing (in which there is obviously no intent to die), and deliberately manipulating a situation so that someone else ends one’s life – colloquially, “suicide by cop” (because the cause of death is not self-directed injurious behavior).

¹In discussing lay normative judgments of suicide, we must, by necessity, delve a bit into the history of philosophical treatments of this act. We are deeply indebted to Jennifer Michael Hecht’s (2013) masterful treatment of this subject, *Stay*, and strongly recommend it to any reader interested in philosophical views on suicide, both ancient and modern, or in a secular argument against suicide.

²Throughout this paper, we use the terms “oppose” and “opposition” to refer to a judgment that suicide is wrong in some sense, and we use the term “normative” as a descriptor for this judgment. Consistent with terminology from Social Domain Theory (e.g., Tisak & Turiel, 1984), we reserve the term “moral” for judgments something is wrong because it harms another person (see Turiel, 1983), and “prudential” for judgments that something is wrong because it harms the actor themselves. Thus, moral and prudential judgments of wrongness are both normative judgments that an act is wrong, but of different sorts. Our measures are intended to capture normative opposition to suicide, broadly construed. Landy et al. (2017) used similar methods to those employed here and utilized these terms in the same fashion.

1.1 Prior Research on Normative Judgments of Suicide and Self-Harm

Some prior research has investigated normative judgments of suicide, though, as noted above, not very much. Other work has examined judgments of self-directed harms, but not suicide, specifically. Overall, the research has been somewhat fragmented, with each study examining one or a handful of explanations for normative opposition to suicide and self-harm. Little work has connected these disparate findings, as of yet.

DeScioli et al. (2012) found that people who declared suicide to be wrong also tended to report that it harmed some victim, typically identified as the actor, or the actor's family. However, it is not clear whether these perceptions of harm drove participants' judgments, or if they represent post-hoc rationalizations to support the initial judgment of wrongness. For their part, the authors of this study favor the latter view, writing that "people often nominated unverifiable victims suggesting that people readily fabricate victims when they are unavailable" (DeScioli et al., 2012, p. 147).

Chakroff et al. (2013) examined judgments of self-directed bad acts (though not suicide, specifically), and found that they display several characteristics associated with violations of "purity". Specifically, bad acts directed toward the self are considered more impure than harmful, they are particularly condemned by socially conservative participants, judgments of these acts are less sensitive to intent (compared to judgments of other-directed acts), and judgments of the actor's moral character were particularly harsh, as compared to judgments of the act itself. Based on these results, the authors conclude that self-directed bad acts are viewed as violations of purity. However, follow-up research tested whether self-harms exhibit three properties typically ascribed to purity violations: first, they elicit disgust (in Western samples), second, they are considered disgusting and wrong across diverse cultures, and third, judgments of their wrongness are relatively insensitive to intent (Kollareth & Russell, 2018). Compared to harms directed at others, self-harms did not generally elicit more disgust, they were considered less wrong, and judgments were sensitive to intent, and these results were similar for American and Indian participants. Thus, these findings are inconsistent with the idea that self-harms are opposed because they violate "purity".

Building on the work of Chakroff et al. (2013), Rottman et al. (2014a) conducted two correlational studies, and found that perceptions that suicide "tainted the soul" of the actor predicted ratings of how wrong the act was, whereas perceptions of harm did not (though see Gray, 2014, and Rottman et al., 2014b for some discussion of this result). They thus argue that suicide is viewed as a purity violation that is wrong for reasons beyond the harm it does. This is an interesting and important insight, but the research is weakened somewhat by its reliance on correlational designs. Follow-up research using experimental methods tested whether suicide shows two hallmarks of purity violations – they elicit disgust, and wrongness judgments of them are insensitive to intentions (Kollareth et al., 2019). Across five studies, suicide generally did not show these purity hallmarks. Thus, the idea that self-harms (including suicide) are violations of "purity" has received somewhat mixed support.

Interviews with suicide attempt survivors, family members of people who died by suicide, and therapists specializing in loss of loved ones by suicide have found that those who attempt suicide are stigmatized as immoral, malicious, untrustworthy, and selfish (among other negative attributes including cowardly, incompetent, and dangerous, see Sheehan et al., 2017), but this research does not tell us *why* this is the case. Lastly, a recent unpublished master's thesis found that suicide is considered more wrong when it has no clear cause and when the deceased person is survived by dependents, and less wrong when the deceased had a particularly difficult and tragic life (Ramírez Sierra, 2020).

1.2 The Present Research

Thus, very little empirical research has investigated why ordinary people oppose the act of suicide. Because so little is known on this subject, we adapt the methodology used by Landy et al. (2017), in an attempt to make rapid progress toward answering this question. Landy et al. (2017) examined normative opposition to the use of performance-enhancing drugs (PEDs) in an exploratory manner. Specifically, they experimentally tested eleven possible explanations for this opposition without committing *a priori* to any particular explanation or set of explanations. They concluded that opposition to PEDs is fully explainable by perceptions of fairness, the existence of rules and laws prohibiting them, and the risk they pose to the user. Put another way, they found that when everyone in a competition used PEDs (i.e., there was no unfair advantage), there were no rules prohibiting them, and they carried no risk, opposition to PED use disappeared, at least on average. This “deep dive” methodology, like Platt's (1964) “strong inference”, is capable of falsifying multiple, competing explanations for a phenomenon in a single program of research. We employ a similar approach here, testing ten possible reasons why people might oppose the act of suicide, which we draw from philosophy, prior empirical research, and public health resources. We describe each of these possible reasons in turn below. We did not specify any *a priori* predictions about which reasons would ultimately turn out to matter to people's judgments, but rather tested all of them in an exploratory manner.

Time Left. Perhaps the most obvious reason for opposing suicide on normative grounds is that it destroys future life-years that would have been worth living. This idea is consistent with a vast swath of prior research showing that people place greater value on the lives of younger individuals than older individuals, at least in part because the young have more expected life-years left to live (Cropper et al., 1994; Dolan et al., 2005; Goodwin & Landy, 2014; Johannesson & Johannesson, 1997; Lewis & Charny, 1989; Li et al., 2010; Ratcliffe, 2000; Rodríguez & Pinto, 2000; Tsuchiya et al., 2003). Thus, one factor that might influence opposition to suicide is the age of the actor; opposition might be reduced if the actor is already quite old, and therefore likely does not have many life-years left anyway.

Social Harm. Suicide often does harm to other people besides the actor. As the United Health Foundation puts it, “When someone dies by suicide, like any cause of death, the loss is felt by many. The ripples of loss spread from close family and friends to community members, acquaintances and even people the deceased did not know” (United Health Foundation, 2020). This idea was presaged centuries ago by Thomas Aquinas’s argument that suicide injures the deceased’s community (see Beauchamp, 1989). And, as noted above, empirical research has found that people nominate the deceased’s family as victims harmed by their suicide (DeScioli et al., 2012). If the harm done to close others partially drives opposition to suicide, then this leads to the prediction that suicide should be considered less wrong when the actor does not have any close others *to* harm by their action.

Soul Taint. As discussed above, Rottman et al. (2014a) found using correlational designs that perceptions that the actor’s soul was “tainted” predicted opposition to suicide. We aimed to test this idea experimentally. We considered several possible ways of doing so, including simply stipulating for participants that souls do not exist (and therefore cannot be tainted), or stipulating that souls may or may not exist, but cannot be tainted by the act of suicide. We ultimately rejected such approaches, because we did not want to require participants to suspend their preexisting metaphysical beliefs by fiat. Instead, we tested the novel prediction that, if a person has lived a life of criminality and deviance, then that person’s soul should *already* be considered “tainted”, and therefore suicide should have less of a tainting effect on the soul of such a person than on the soul of the average person. Thus, according to this explanation, the suicide of a “tainted” criminal should be considered less wrong than the suicide of an average person.

Self-Harm. Turiel’s Social Domain Theory (SDT) postulates that actions can be “wrong” in several different ways. One way is by violating “prudential” rules or norms. Prudential violations are considered wrong because of the harm they cause to the actor (Tisak & Turiel, 1984). Examples of prudential violations include recreational drug use, when assessed by non-users (Nucci, Guerra & Lee, 1991). Suicide is clearly the ultimate form of self-harm, and therefore the perception of this harm may partially drive opposition. It is logically impossible to develop an experimental scenario in which someone commits suicide without harming themselves, so we manipulated the amount of harm done to the self by manipulating *how* the act of suicide was carried out – painlessly, or slowly and painfully. If suicide is seen as an especially serious prudential violation, then it should be considered *even more wrong* when it is done in a way that causes the actor great pain.

Severity of Circumstances. Suicide is usually enacted in response to some present difficulties in one’s life. Perhaps, the severity of those difficulties matters to normative judgments of suicide. Judgments of someone who commits suicide in response to *extremely* difficult circumstances may be less harsh than of someone who does so in response to only *mod-*

erately difficult circumstances. Indeed, some prior empirical evidence supports this idea (Ramírez Sierra, 2020).

Self-Euthanasia. Deliberately ending one's life in order to escape a chronic or incurable medical condition – what might be termed “self-euthanasia” – constitutes a special case of the “severity of circumstances” explanation above. To illustrate what distinguishes the two, consider that it would be unthinkable for a doctor to end a person's life because the person was facing financial trouble, no matter how severe, but it might be permissible for a doctor to end a patient's life to relieve their suffering from an incurable ailment. Killing *oneself* in this case is just doing what (some) doctors might otherwise do. Indeed, philosophical treatments of the morality of suicide frequently treat cases of debilitating medical conditions differently from other impetuses to suicide. For instance, Plato is opposed to ending one's life in most cases, but in the *Republic*, he advocates for withholding medical treatment from individuals who are so incapacitated that they are no longer able to lead worthwhile, productive lives.³ Renaissance-era philosopher Michel de Montaigne drew on Pliny's argument that certain especially painful diseases permit suicide and concluded that extreme pain and suffering are the “most excusable incitements” to suicide (Ferngren, 1989, p. 161). Hume expressed a similar belief that “life is not worth living if ill health seriously and irreversibly undermines one's potential for enjoyment” (Beauchamp, p. 201). Moreover, the very existence of a separate word (“euthanasia”) for cases in which one wishes to die to escape painful, incurable medical problems may suggest that people do not consider them to be cases of “suicide” at all, but rather something qualitatively different. Indeed, Hecht (2013, p. 24) notes that a suicide that “simply hastens an inevitable and otherwise painful end” “may not even be suicide”. Therefore, we test whether killing oneself to escape a painful, incurable medical condition is considered more morally acceptable than other types of suicide, separately from our test of whether the severity of one's (non-medical) circumstances matters.

Prior Exposure to Suicide. Ever since Emile Durkheim's (1979) landmark treatment of the subject, it has been recognized that suicide is “contagious”; prior exposure to suicide by someone else is a risk factor for committing suicide oneself (see, e.g., Coleman, 1987; Hecht, 2013; Wilcox et al., 2010). But might such prior exposure also affect normative judgments of suicide? If a person has been previously exposed to suicide, perhaps their own suicide is partially attributed to this prior exposure, reducing the amount of blame placed on the actor themselves.

Attempted Suicide. Fortunately, not all suicide attempts are successful. Perhaps the *outcome* of a suicide attempt influences moral judgments. From a crudely “utilitarian”

³This is necessarily an oversimplification of Plato's views. For a thorough treatment of the nuances of Plato's arguments about suicide in general, and severe medical cases in particular, see Cooper (1989).

perspective, an attempted (but unsuccessful) suicide is preferable to a successful suicide, because no life is actually lost.⁴ Prior empirical research on “moral luck” has found that when identical acts lead to differently negative consequences (e.g., a drunk driver crashes into a mailbox, or crashes into a pedestrian), at least some sorts of normative judgments track the severity of the consequences, with worse consequences generally eliciting greater desire for punishment (e.g., Cushman, 2008; Martin & Cushman, 2015, 2016). Thus, it may be the case that attempted suicides are judged to be less wrong than successful ones because they do not share the same negative consequences.

Impulsiveness. Prior research has found that the amount of thought put into a decision can affect normative judgments about it. Specifically, thinking carefully before committing an immoral act sometimes makes the act seem less wrong (Pizarro et al., 2003; Starmans & Bloom, 2016). The primary reason for this seems to be that impulsive decisions reflect less decision conflict (i.e., less of a competing desire to *not* commit the immoral act). It seems plausible that suicide is considered less wrong when one has thought long and hard about it, as compared to when one kills oneself impulsively. However, the mechanism underlying this difference (if it is found) may have less to do with inferred decision conflict, and more to do with an inference that one is treating one’s life as “cheap” or unimportant, a thing to be disposed of with little consideration (see Boyle, 1989). We return to this point in the General Discussion. On the other hand, bad acts are sometimes seen as *worse* when committed after thorough deliberation. Consider, for instance, that premeditated murder is usually subject to harsher punishment than more impulsive crimes. It is possible that suicide is similarly made worse when it has been thought hard about and deliberated over.

Moral-Religious Background. Although most religious traditions reject suicide, there are some ethical systems that permit or even embrace it (Hecht, 2013). Perhaps, when judging the permissibility of suicide, people consider the religious or ethical system the actor was raised in. If Person A comes from a religious tradition that prohibits suicide, but Person B comes from a tradition that permits it, people may believe that suicide is wrong for Person A, but less wrong for Person B. Prior research suggests that people generally consider their own *moral* values to be objectively true (Goodwin & Darley, 2008) and universally applicable (Landy, 2016; Turiel, 1983), but that values rooted in religious traditions can be seen as more relative (e.g., Nucci, 1982; Nucci & Turiel, 1993). For instance, Nucci and Turiel (1993) found that Amish-Mennonite children believed that it was wrong for a

⁴Of course, this crude analysis obscures the fact that genuine utilitarianism would evaluate the rightness or wrongness of an action based on *expected* harm done, rather than *actual* harm done (see, e.g., Mulgan, 2020). For instance, a utilitarian would not excuse drunk driving because the driver *just happened* to make it home without injuring anyone, because the plain fact of the matter is that, by driving drunk, the driver very likely *could* have injured or killed someone. We put this issue aside here, because what is at issue is not whether utilitarian philosophers would consider unsuccessful suicide attempts to be less wrong than successful ones (most almost certainly would not), but rather whether laypeople’s normative judgments are sensitive to this difference in outcomes.

woman of their faith to go without a head covering, but that this was not necessarily wrong for women from other religious traditions that do not stipulate head coverings as part of the female wardrobe. People might view the wrongness of suicide as similarly contingent upon the actor's own moral-religious background, such that it is less wrong for people who come from traditions that permit it. Related to this, normative judgments might be more lenient when an agent's own moral-religious background permits suicide simply because the agent is seen as ignorant of the moral facts, and therefore unaware that by killing himself, he intends to do something wrong. Thus, there are multiple reasons to think that the moral-religious background of the actor might matter to normative judgments of suicide.

2 Study 1

In Study 1, we tested whether each of the ten possible explanations above causally contributes to normative opposition to suicide, by experimentally manipulating the presence or absence of each one, in a within-subjects design.

2.1 Method

2.1.1 Participants

Participants in this study were undergraduate students at Nova Southeastern University, who participated in exchange for course credit. We aimed to recruit at least 100 participants the first semester we ran this study. Because we did not reach this target after one semester, we ran the study for a second semester. Before beginning the study, participants were asked if they or anyone close to them had ever attempted suicide. Participants who responded in the affirmative were not permitted to complete the study, due to the potential for emotional distress (though they still received credit for their participation). The final sample consisted of $N = 183$ participants who completed at least a portion of the study (31 male, 15 no response, $M_{Age} = 19.36$, $SD_{Age} = 1.69$).

2.1.2 Materials and Procedure

This study was run during the COVID-19 pandemic, making in-person data collection infeasible, so participants completed the study online on their own devices. The study consisted of two blocks of scenarios. The first block presented 15 scenarios, presented one after the other, each on its own screen, in a new randomized order for each participant. One of these 15 scenarios was our "Standard" scenario: "Mr. K is a 40-year-old man who is very unhappy with his life. He decides to take his own life. Mr. K takes an entire bottle of sleeping pills and dies painlessly." The remaining 14 were variations on this scenario meant to experimentally manipulate the ten possible reasons for normative opposition to suicide discussed above, each of which included a different character ("Mr. B", "Mr. F", etc.). In all cases below, we refer to the scenario that we predicted would elicit less normative opposition

to suicide as the “experimental” scenario, and the scenario that we predicted would elicit more normative opposition as the “control” scenario. The Standard scenario served as the control scenario in five of ten comparisons, and as the experimental scenario in a sixth, with the remaining 14 scenarios each serving as one condition in one comparison. After each scenario, participants responded to a single normative judgment question, adapted from the measure used by Landy et al. (2017, Study 13): “How right or wrong is what [character name] did?”. Responses were made on a bipolar 15-point Likert scale ranging from “The wrong thing to do” (coded as -7) to “The right thing to do” (coded as 7), with the midpoint labeled “Neither right nor wrong” (coded as 0).

The second block consisted of manipulation checks. The 15 scenarios were presented again, though the standard control was presented six times, with a different manipulation check question each time, because this was one condition in six of the ten comparisons (e.g., we needed to evaluate both how much longer Mr. K would be expected to live, and, separately, how tainted his soul was). As in the first block, each scenario and manipulation check question was presented on its own screen, one after the other, and the order of presentation was randomized for each participant. Following the second block, participants completed a brief demographics survey (age, gender, political affiliation, religious affiliation, religiosity, and year at school). After completing the demographics survey, participants were debriefed and thanked for their time.

For each of the ten possible explanations detailed above, there were two dependent measures and two relevant manipulation checks. We analyze the ten explanations separately, so we include in each analysis all participants who completed the relevant dependent measures and manipulation checks, even if they did not complete the entire study. Therefore, the exact sample size varies across analyses, due to some partially incomplete responses. For instance, for the “severity of circumstances” explanation, any participant who responded to both the control and experimental scenarios, and their corresponding manipulation checks, would be included in the analysis, even if he/she did not respond to the scenarios relevant to other analyses.

Below, we detail the experimental manipulations, and the manipulation checks.

Time Left. The experimental scenario was identical to the Standard scenario, except for the character’s name (“Mr. F”), and age: Mr. F was described as being 90 years old rather than 40. The Standard scenario served as the control scenario. The comparison of these two scenarios tests whether people view suicide as less wrong for people with fewer expected life years remaining. In the second block of the study, participants saw both scenarios again, and responded to the question “If [Mr. K/Mr. F] had not taken his own life, how much longer would he probably have lived?”, with responses made on a nine-point scale ranging from (1) “A very short time” to (9) “A very long time”. Participants indicated that 40-year-old Mr. K would likely have lived longer ($M = 7.00$, $SD = 1.73$) than 90-year-old Mr. F ($M =$

3.07, $SD = 2.10$; $t(163) = 20.05$, $p < .001$, $d_{RM} = 1.57$),⁵ indicating that the manipulation was successful.⁶ Importantly, this difference in means represents a difference of four scale points on a nine-point Likert scale, *not* a difference of four expected life-years. Mr. K was rated near the high end of the scale (“a very long time” left to live), while Mr. F was rated much closer to the low end (“a very short time” left to live), as would be expected.

Social Harm. The experimental scenario was identical to the Standard scenario, except for the character’s name (“Mr. H”), and a statement indicating that he had no living family or close friends who would miss him if he was gone. The Standard scenario served as the control scenario. This comparison tests whether people view suicide as less wrong when no one else is negatively affected by the act, other than the person committing it (i.e., whether any social harm was done). In the second block, participants saw both scenarios again, and responded to the question “How much did [Mr. K/Mr. H] taking his own life negatively affect other people, besides [Mr. K/Mr. H]?”, with responses made on a nine-point scale ranging from (1) “No negative effects” to (9) “Extreme negative effects”. Participants indicated that less harm was done to others by Mr. H (who had no family or close friends, $M = 3.98$, $SD = 2.54$) than by Mr. K (whose social status was not described, $M = 6.89$, $SD = 2.16$; $t(158) = 13.69$, $p < .001$, $d_{RM} = 1.09$), indicating that the manipulation was successful.

Soul Taint. The experimental scenario was identical to the Standard scenario, except for the character’s name (“Mr. N”), and a statement that he had lived a life of criminality and deviance and hurt many people during his life. Our reasoning was that Mr. N’s soul should be seen as *already* tainted, and therefore less *further* tainted by an act of suicide. The Standard scenario served as the control scenario. Thus, this comparison was intended to test whether people view suicide as less wrong when it cannot have the effect of tainting one’s soul. In the second block, participants viewed both scenarios again and responded to the question “Was the purity of [Mr. K’s/Mr. N’s] soul tainted as a result of him taking his own life?”, which we directly adapted from Rottman et al. (2014), with responses made on a nine-point scale ranging from (1) “Not at all tainted” to (9) “Extremely tainted”. However, participants indicated that Mr. N’s soul was *more* tainted ($M = 4.90$, $SD = 2.73$) than Mr. K’s ($M = 4.45$, $SD = 2.41$; $t(151) = 2.31$, $p = .022$, $d_{RM} = 0.19$), indicating that our manipulation was unsuccessful. Yet, as we shall see below, this manipulation did nonetheless affect judgments of moral wrongness. We return to this puzzling pattern of results below.

⁵For within-subjects comparisons, we report the repeated-measures Cohen’s d , d_{RM} , as our effect size measure. Briefly, this effect size is calculated as the mean of participants’ difference scores, divided by the standard deviation of difference scores. Thus, the interpretation of this statistic is somewhat different than the more familiar independent-groups Cohen’s d , d_{IG} (see Morris & DeShon, 2002, for more detail).

⁶Like most of our manipulation checks, this one employs a continuous response scale. Therefore, there is no clear criterion by which to assess whether any particular participant “passed” or “failed” the manipulation check. We therefore did not exclude participants from analysis based on their responses, but rather used the manipulation checks to determine whether our manipulations were successful, overall. With one exception, the manipulation checks indicated that they were.

Self-Harm. The Standard scenario served as the experimental scenario for this comparison. The control scenario was identical to the Standard scenario, except for the character's name ("Mr. M"), and the method of suicide: Mr. M killed himself slowly and painfully, by hanging himself. This comparison tests whether people view suicide as less wrong when it causes less pain, rather than more pain, for the person committing it. In the second block of the study, participants saw both scenarios again, and responded to the question "How much pain and harm did [Mr. M/Mr. K] cause to himself?", with responses made on a nine-point scale ranging from (1) "Very little pain and harm" to (9) "Very much pain and harm". Participants correctly indicated that Mr. M, who hanged himself, experienced more pain and harm ($M = 8.48$, $SD = 1.32$) than Mr. K, who took a bottle of sleeping pills ($M = 3.78$, $SD = 2.85$; $t(160) = 18.07$, $p < .001$, $d_{RM} = 1.42$), indicating that the manipulation was successful.

Severity of Circumstances. The two scenarios in this comparison were both based on the Standard scenario, but explicitly described the difficulty of the circumstances that had pushed the character to suicide. In the experimental scenario, the character ("Mr. B") had recently lost his job with no prospects of finding a new one and was having serious trouble making ends meet. In the control scenario, the character ("Mr. P") was having difficulty at work and some money problems. This comparison tests whether people view suicide as less wrong when one is facing more severe difficulties in their life prior to the act. In the second block, participants saw both scenarios again and responded to the question "How difficult was [Mr. B's/Mr. P's] overall situation in life before he took his own life?", with responses made on a nine-point scale ranging from (1) "Not at all difficult" to (9) "Extremely difficult". Participants indicated that Mr. B's situation was more difficult ($M = 6.28$, $SD = 1.74$) than Mr. P's ($M = 5.43$, $SD = 1.81$; $t(159) = 6.78$, $p < .001$, $d_{RM} = 0.54$), indicating that the manipulation was successful.

Self-Euthanasia. The two scenarios in this comparison were both based on the Standard scenario, except that they contained information about the character's health. In the experimental scenario, the character ("Mr. J") was described as dealing with a painful, untreatable, chronic illness. In the control scenario, the character ("Mr. E") was described as very healthy, physically. This comparison tests whether people view suicide as less wrong when it is done to escape from a serious, untreatable health problem. In the second block, participants saw both scenarios again and responded to the same manipulation check question as in the severity of circumstances scenarios above, as we considered committing suicide to escape a health problem to be a special case of committing suicide to escape difficult circumstances more generally. Participants indicated that Mr. J's situation was more difficult ($M = 7.90$, $SD = 1.75$) than Mr. E's ($M = 3.51$, $SD = 2.04$; $t(165) = 20.17$, $p < .001$, $d_{RM} = 1.57$), indicating that the manipulation was successful.

Prior Exposure to Suicide. The experimental scenario was identical to the Standard scenario, except for the character's name ("Mr. U"), and a statement stipulating that Mr. U took his own life "just like his father did several years ago". The Standard scenario served as the control scenario. This comparison tests whether people view suicide as less wrong when a person has been exposed to it previously, perhaps because the person is seen as less to blame for the act. In the second block, participants saw both scenarios again, and responded to the question, "How much prior exposure to suicide by other people did [Mr. U/Mr. K] have before he took his own life?", with responses made on a nine-point scale ranging from (1) "Very little prior exposure" to (9) "Very much prior exposure". Participants indicated that Mr. U had had more prior exposure to suicide ($M = 8.14$, $SD = 1.73$) than Mr. K ($M = 3.43$, $SD = 2.08$; $t(161) = 21.97$, $p < .001$, $d_{RM} = 1.73$), indicating that the manipulation was successful.

Attempted Suicide. The experimental scenario was identical to the Standard scenario, except for the character name ("Mr. C"), and the outcome: Mr. C woke up the morning after his suicide attempt, having tried unsuccessfully to kill himself. The Standard scenario served as the control scenario. This comparison tests whether people view *attempted* (but unsuccessful) suicides as less wrong than *successful* suicides, presumably because less harm is actually done. In the second block, participants saw both scenarios again, and responded to the question "Did [Mr. C/Mr. K] successfully take his own life, or unsuccessfully attempt to take his own life?", with two response options: "Successfully took his own life" and "Unsuccessfully tried to take his own life". One hundred fifty-four out of 162 participants correctly answered both questions, which is significantly different from chance responding by a McNemar chi-squared test ($\chi^2(1) = 146.16$, $p < .001$), indicating that the manipulation was successful.

Impulsiveness. Both scenarios were based on the Standard scenario, except that they specified how impulsively or deliberately the character made the decision to commit suicide. In the experimental scenario, the character ("Mr. V"), committed suicide after thinking hard about it for several months, whereas in the control scenario, the character ("Mr. I"), committed suicide impulsively without giving it much thought. This comparison thus tests whether people view suicide as less wrong when it is a carefully considered, deliberate decision, rather than an impulsive one. In the second block, participants saw both scenarios again, and responded to the question "How hard did [Mr. V/Mr. I] think about his decision to take his own life?", with responses made on a nine-point scale ranging from (1) "Not hard at all" to (9) "Extremely hard". Participants indicated that Mr. V had thought harder about his decision ($M = 7.50$, $SD = 1.89$) than Mr. I ($M = 1.84$, $SD = 1.76$; $t(164) = 28.35$, $p < .001$, $d_{RM} = 2.21$), indicating that the manipulation was successful.

Moral-Religious Background. Both scenarios were based on the Standard scenario, except that they contained information about the character's religious upbringing and his religion's teachings on the normative status of suicide. In the experimental scenario, the character ("Mr. Z") had been raised in a religious home where he was taught that suicide was a personal choice. In the control scenario, the character ("Mr. S") had been raised in a religious home where he was taught that suicide is wrong. This comparison thus tests whether the ethical beliefs that a person was raised with affect how wrong people consider it to be for that person to commit suicide. In the second block, participants saw both scenarios again and responded to the question "What does [Mr. Z's/Mr. S's] religion teach about whether suicide is permissible or not?", with responses made on a nine-point scale ranging from (1) "Suicide definitely not permissible" to (9) "Suicide definitely permissible". Participants indicated that Mr. Z's religion considered suicide to be more permissible ($M = 6.17$, $SD = 2.56$) than Mr. S's ($M = 1.77$, $SD = 1.74$; $t(163) = 17.50$, $p < .001$, $d_{RM} = 1.37$), indicating that the manipulation was successful.

2.2 Results and Discussion

Descriptive and inferential statistics for all comparisons in this study are presented in Table 1. Briefly, participants considered suicide to be less wrong when the actor had little (versus much) time left to live, when the actor had no family or friends and therefore caused minimal social harm, when the actor had lived a life of criminality and deviance, when the method of suicide was painless, when the actor had a painful chronic illness (therefore making the act of suicide an example of self-euthanasia), when the actor had thought hard about the decision (as opposed to making it impulsively), and when the actor's own religious background taught that suicide was a personal choice. The severity of the actor's (financial) circumstances, prior exposure to suicide, and whether the suicide attempt was successful or unsuccessful did not significantly affect normative judgments.

One of these results in particular deserves further elaboration. The manipulation of social harm may have inadvertently also manipulated perceptions of loneliness, which might be considered a special case of severity of circumstances, and a potential justification for suicide. To examine this possibility, we conducted a within-subjects mediation analysis using the MEMORE macro for SPSS (Montoya & Hayes, 2017), with within-subjects condition (control vs. experimental) as the predictor variable, the difference between the manipulation checks as the mediator, and the difference between the wrongness judgments as the outcome. We found a significant indirect effect of condition on wrongness judgments through the manipulation check ($ab = -0.56$, 95% CI $[-1.03, -0.16]$) and no remaining direct effect, once the indirect effect was accounted for ($c' = -0.19$, $p = .469$, 95% CI $[-0.70, 0.32]$). This result is consistent with (though does not prove) the idea that the difference in wrongness judgments is psychologically mediated by the reduced harm to others, rather than by inferences of loneliness.⁷

⁷We also conducted analogous mediation analyses for the other nine manipulations, though they are less

TABLE 1: Descriptive and Inferential Statistics, Study 1.

Explanation	<i>M (SD)</i>		<i>df</i>	<i>t</i>	<i>p</i>	<i>d_{RM}</i>
	Control Condition	Experimental Condition				
Time Left	−5.05 (2.64)	−3.81 (3.49)	163	6.46	< .001	0.50
Social Harm	−5.08 (2.74)	−4.33 (3.24)	158	4.24	< .001	0.36
Soul Taint	−5.13 (2.69)	−3.87 (3.33)	151	6.19	< .001	0.50
Self-Harm	−5.65 (2.40)	−5.16 (2.66)	160	3.89	< .001	0.31
Severity of Circumstances	−5.01 (2.82)	−4.91 (2.84)	159	0.79	.432	0.06
Self-Euthanasia	−5.07 (2.64)	−2.07 (3.91)	165	11.05	< .001	0.86
Prior Exposure	−4.96 (2.74)	−4.89 (2.78)	161	0.47	.642	0.04
Attempted Suicide	−5.04 (2.70)	−4.84 (2.76)	161	1.55	.124	0.12
Impulsiveness	−5.58 (2.42)	−4.38 (3.21)	164	5.70	< .001	0.44
Moral-Religious Background	−5.02 (2.73)	−4.19 (3.21)	163	4.97	< .001	0.39

Notably, all condition means fell significantly below the scale midpoint of 0 (“Neither right nor wrong”), all p s < .001, indicating that, on average, participants considered suicide to be wrong, all things considered, in every single case. This is perhaps unsurprising, because these judgments are clearly multiply determined – seven out of ten manipulations had significant effects on participants’ normative judgments! It seems that each of these seven considerations exerts some influence on normative judgments of suicide, but most of these effect sizes are small-to-medium in magnitude. The exception is the case of self-euthanasia, which showed a large effect, $d_{RM} = 0.86$. Participants in this study seemed to view this as a qualitatively different kind of circumstance, where different principles apply. However, even in this special case, they still rated the act of suicide as somewhat wrong, overall.

We also examined how participants’ individual characteristics related to their judgments. Specifically, for all participants for whom we had complete demographic data ($n = 170$), we examined correlations between mean wrongness ratings across the 15 judgments (Cronbach’s $\alpha = .97$) and participant gender, a three-item measure of political liberalism (economic issues, social issues, and overall politics, $\alpha = .94$), and a two-item measure of religiosity and spirituality ($\alpha = .73$). We did not examine age in these analyses, because the undergraduate sample showed very little variation on this variable. Participant gender did not predict wrongness judgments ($r(168) = .05$, $p = .493$). The correlation between the composite politics measure and mean wrongness judgments was nearly significant ($r(168) = .15$, $p = .052$). Exploratory analyses showed that this relationship was driven by the question about social issues: social liberalism was significantly correlated with more per-

obviously necessary. The SPSS data file, syntax, and output are available at <https://osf.io/mjvqq>.

missive attitudes toward suicide (i.e., higher ratings on the wrong-to-right scale) ($r(168) = .17, p = .024$), whereas economic liberalism was not ($r(168) = .11, p = .156$), and “overall” liberalism was only marginally related ($r(168) = .14, p = .072$). The composite measure of religiosity/spirituality was not significantly correlated with wrongness judgments ($r(168) = -.06, p = .463$), though exploratory analyses showed that the religiosity item did correlate significantly with judgments ($r(168) = -.16, p = .038$), while the spirituality item did not ($r(168) = .02, p = .771$). Thus, participants with more liberal attitudes toward social issues considered suicide to be less wrong, while more religious participants considered it to be more wrong, across scenarios.

In sum, seven of the ten possible considerations that we tested significantly affected normative judgments of suicide. We now turn to determining how complete this set of explanations is.

3 Study 2

Now that we have found several considerations that affect normative judgments of suicide, we turn in Study 2 to testing whether there are any further considerations that we have *not* yet examined. Specifically, in Study 2, we tested the combined effect of the seven manipulations that produced significant results from Study 1, compared to a control condition. If there are other considerations that matter to judgments of suicide that we have not examined, then it should still be considered somewhat wrong, even when all seven considerations from Study 1 push against this judgment. On the other hand, if opposition to suicide is wiped out when all of these considerations are combined, we can be reasonably confident that they exhaust the concerns that people consider when judging it. Moreover, in Study 2, we employ a between-subjects design, and study a more diverse sample than the undergraduates from Study 1.

3.1 Method

3.1.1 Participants

Participants were recruited through Amazon Mechanical Turk. Before beginning the study, participants were asked if they or anyone close to them had ever attempted suicide. Participants who responded in the affirmative were not permitted to complete the study, due to the potential for emotional distress, as in Study 1 (though they were still compensated for their participation). There were two between-subjects conditions in this study, so we aimed for approximately 100 total participants after exclusions, to have 50 participants per cell, as recommended by Simmons, Nelson, and Simonsohn (2013). The final sample consisted of $N = 97$ participants who completed the full study (64 male, $M_{Age} = 35.52, SD_{Age} = 8.67$).

3.1.2 Materials and Procedure

Participants were randomly assigned to either the control condition or the experimental condition. In both conditions, they read one scenario and responded to it. The control scenario read as follows: “Mr. K is a 40-year-old man who is very unhappy with his life, even though he is very healthy, physically. He was raised in a religious home and had always been taught that suicide was wrong, but he decides to take his own life. Without giving it much thought, Mr. K hangs himself, and, in a very slow and painful process, chokes to death.” In other words, this scenario combined the control scenarios of the seven significant comparisons from Study 1. The experimental scenario, on the other hand, combined the seven experimental scenarios: “Mr. K is a 90-year-old man who is very unhappy with his life because he is dealing with a chronic illness that causes very intense, frequent pain that cannot be treated. He has lived a life of criminality and deviance and has hurt many people over the years. He was raised in a religious home, and his religion teaches that suicide is a personal choice. He has no living family, and no close friends. Knowing that there is no one who would miss him if he was gone, he decides to take his own life to escape from the pain of his disease. After thinking hard about it for several months, Mr. K takes an entire bottle of sleeping pills and dies painlessly.”

After reading the scenario, participants responded to a single normative judgment question: “How right or wrong is what Mr. K did?”. Responses were made on the same scale as in Study 1. Following the response, participants completed a brief demographics survey (age, gender, political affiliation, religious affiliation, and religiosity). After completing the demographics survey, participants were debriefed, thanked, and paid.

3.2 Results and Discussion

Similar to Study 1, participants considered suicide to be more wrong in the control condition ($M = -1.69$, $SD = 4.55$) than in the experimental condition ($M = 0.40$, $SD = 5.08$). In other words, participants considered suicide to be less wrong when the actor had little time left to live, had no family or friends, had lived a life of criminality and deviance, had a painful chronic illness, and had thought hard about the decision to commit suicide, when the method of suicide was painless, and the actor’s own moral-religious background taught that suicide was a personal choice ($t(95) = 2.14$, $p = .035$, $d_{IG} = 0.43$).

Moreover, participants’ normative judgments were significantly below the scale midpoint (labeled “Neither right nor wrong”) in the control condition ($t(48) = 2.61$, $p = .012$), indicating that they considered the act to be at least somewhat wrong, though not as wrong as the undergraduates in Study 1 tended to. However, in the experimental condition, participants’ judgments did not differ significantly from the scale midpoint ($t(47) = 0.54$, $p = .592$), indicating that in this condition, participants did not consider suicide to be wrong, on average. Thus, there is no further opposition to be accounted for – when the seven moral considerations that we identified in Study 1 are taken off the table, normative opposition

to suicide disappears, at least when examining average opposition levels among online workers.

Similar to Study 1, we also examined how individual differences relate to wrongness ratings. Specifically, we examined how participant sex, age, the three-item political liberalism measure from Study 1 ($\alpha = .96$) and the two-item religiosity-spirituality measure from Study 1 ($\alpha = .89$) predicted wrongness judgments. Neither sex ($r(95) = -.05$, $p = .626$) nor age ($r(95) = .11$, $p = .287$) nor political liberalism ($r(95) = .03$, $p = .761$) nor religiosity-spirituality ($r(95) = -.08$, $p = .443$) correlated significantly with wrongness judgments. Unlike in Study 1, none of the individual items from the composite measures was significantly correlated with wrongness judgments ($ps > .202$). Thus, the individual difference measures from Study 2 were not predictive of wrongness judgments. However, this is not to say that there was unanimous agreement among our participants – there was not. In the experimental condition, 15 participants still rated Mr. K's suicide below the scale midpoint (i.e., as “wrong” to some degree), 6 made ratings exactly at the midpoint (“neither right nor wrong”), and 27 rated Mr. K's suicide above the scale midpoint (i.e., as “right” to some degree). Thus, there was important variation in participants' responses, even when all seven considerations pushed against opposition to suicide, but this variation was not predicted by demographic variables.

4 General Discussion

In *The Myth of Sisyphus*, Camus (1991) posits that “there is but one truly serious philosophical problem, and that is suicide.” We disagree with Camus; we think that there are other serious problems in philosophy, and there are certainly other serious problems in empirical psychology. But we think Camus is correct that suicide presents a puzzle. It is widely lamented and viewed as *wrong* in some sense, but it is not yet clear why this is. In this research, we used an exploratory, “deep-dive” methodology (Landy et al., 2017) to study this question. We found that normative opposition to suicide is multiply determined, with seven factors contributing to it (Study 1): the amount of time the actor would have otherwise lived, emotional harm to the actor's close conspecifics, whether the actor lived a moral life, the amount of pain caused to the actor by the act, whether the actor is healthy or in poor health, whether the actor thought hard about the decision to commit suicide, and whether the actor's own moral-religious background permits suicide or not. Moreover, when all of these considerations push against opposition to suicide, it becomes, on average, normatively neutral (Study 2). This strongly suggests that these seven considerations essentially exhaust the psychological causes of opposition to suicide, at least as we have operationalized it.

4.1 Limitations

There are several important limitations to this research. Perhaps the most obvious is the lack of detail and realism in our scenarios. The scenarios were very brief, and participants knew very little about the characters in them, who were not even given full names (e.g., “Mr. K”). Depersonalized scenarios of this sort have rightly drawn criticism for omitting details that could affect normative judgments in the real world (Hester & Gray, 2020). However, we made this design choice deliberately, to minimize the emotionality of the scenarios, to avoid causing unnecessary distress to participants. Nonetheless, we acknowledge that real-world normative judgments of suicide probably depend on details of the identity of the person who commits it. Related to this, all of our scenarios focused on male characters. Future research could investigate whether normative judgments of suicide vary by the actor’s gender, though, for our part, we do not see any clear reason why they would.

Our samples were also drawn only from American populations. We did sample from two different populations (i.e., undergraduates and online workers), but neither of these can be said to accurately represent the United States population as a whole, let alone people from other countries and cultures. Cross-cultural research on normative judgments of suicide is an important task for further research, especially because prior research suggests that attitudes toward suicide can differ between different cultures (Gard et al., 2005).

4.2 Notable Findings and Future Directions

Overall, we think our results represent a thorough examination of the understudied, but widely expressed, normative judgment that suicide is wrong in some way. However, we think some of our findings warrant further discussion. First, most of the effect sizes in Study 1 were small-to-medium in size, indicating that laypeople do not view suicide as wrong for one or two very important reasons, but rather for several smaller, additive reasons. The exception was the effect size for the Self-Euthanasia manipulation, which was large in size. We speculated above that people may see cases like this as something fundamentally different from typical cases of “suicide”, and therefore as subject to a different set of normative standards. As noted, there is some philosophical precedent for this, though there are also philosophers, such as Aquinas and Kant, who would not treat such cases as any different from other types of suicide (see Amundsen, 1989; Beauchamp, 1989). Philosophical disputes aside, our participants seemed quite convinced that ending one’s life to escape a chronic, painful illness is far less wrong than doing so for other reasons. It is interesting to note that this belief did not extend to severe financial difficulties, as there was no effect of the “Severity of Circumstances” manipulation in Study 1. Serious health problems seem to be viewed as qualitatively different than problems of other sorts when it comes to normative evaluations of suicide.

This raises an interesting question: how would people evaluate a suicide enacted in response to severe *mental* health problems? It is known that suicide is often precipitated by

clinical depression or other mental health issues (Angst et al., 1999; Harwitz & Ravizza, 2000; Westefeld & Furr, 1987). Does people's permissive attitude toward self-euthanasia extend to such cases? We do not have any evidence that speaks to this question, but we speculate that the answer is no, chiefly because mental health issues like depression are treatable, and in some cases episodic, and therefore differ importantly from the chronic, untreatable condition described in our self-euthanasia scenario. It is even possible that normative judgments of suicide may differ between different mental health conditions (e.g., depression versus schizophrenia) in interesting and predictable ways. Future research could examine this question directly.

Related to this point, we did not examine the effects of biological or genetic risk factors for suicide on normative judgments. Suicide, like depression and other mental illnesses, has a genetic component, and normative judgments of the act might vary depending on a person's predisposition toward it (for discussion of how genetic testing for suicide risk might affect stigmatization of *persons*, see Docherty et al., 2021; Kious et al., 2021). Of course, genetics do not directly cause suicide – the genetic cause must still be mediated through a psychological cause, and, ultimately, a deliberate, intentional act. The intervening psychological mechanism is very likely to be a mental illness like depression, so perhaps future research on normative judgments of suicide in response to mental health problems could also examine genetic predisposition as a moderating variable.

Another noteworthy result is that our manipulation of “soul taint” failed (according to our manipulation check), yet it still exerted a significant effect on moral judgments. There are several possible reasons for this. First, it could simply be the case that the manipulation check did not function as intended. Implicit in the design of this manipulation check is an assumption that people believe that there is some upper limit to “how tainted” a soul can be, and/or that soul taint is sub-additive, i.e., each successive bad act taints the soul less than the prior one. However, we do not have direct evidence of this. Perhaps taint to souls is actually *super*-additive, and each successive bad act is *more* tainting than the last. This could plausibly explain why suicide was rated as *more* tainting to the soul of deviant criminal Mr. N than to the soul of Mr. K. Another possibility is the manipulation of soul taint was successful, but only at an implicit level, so participants did not explicitly report that the actor's soul was tainted, despite implicitly believing this to be true. This explanation, of course, is not incompatible with the previous one. Rottman et al. (2014), for their part, seem to believe that perceptions of soul taint are mostly implicit. It is also possible that the effect of the “soul taint” manipulation was mediated through some inference other than a tainted soul. For instance, our participants may have believed that the immoral Mr. N “deserved” his untimely end, or that the world was “better off without him”, i.e., that his death was just and/or beneficial. Another possibility is that our participants viewed suicide as the only way for this evil individual to rid himself of his immoral desires, a position that Plato stakes out in the *Laws* (see Cooper, 1989). Or, perhaps they believed that suicide can be a form of repentance for one's bad deeds, a position with some precedent in Jewish

ethical theory (see Brody, 1989). At this time, we cannot adjudicate between these possible explanations for this result.

We also found that suicide is considered more wrong when it is done impulsively, rather than after much deliberation. Similar results have been found for other kinds of acts that are considered wrong (Pizarro et al., 2003; Starmans & Bloom, 2016). As noted above, the reason for this seems to be that quick decisions indicate a lack of decision conflict. However, in the case of suicide, we think it is more likely that a quick, impulsive decision to end one's life reflects a lack of value placed on one's own existence. In other words, by making an impulsive decision to kill oneself, one treats one's own life as "cheap", a thing to be disposed of with little consideration (see Boyle, 1989). Inferring that one does not value one's own life is a rather different inference from inferring that one experienced little conflict about one's decision, though the former might plausibly imply the latter. It could also be the case that people think that impulsive decisions are more error-prone, and therefore less normatively defensible. We leave investigation of the mechanism underlying this result to future work.

One surprising finding of Study 2 is that even participants in the control condition rated suicide as substantially less wrong ($M = -1.69$) than participants in *any* condition in Study 1 (M s -5.58 to -2.07). This could be due to the differing samples: perhaps South Florida undergraduates are generally more opposed to suicide than are online workers located throughout the United States. Another possible explanation is that, by giving more detail in Study 2, we made the focal character seem more real to participants, making his actions more understandable. We are not aware of any research showing that more detail in hypothetical scenarios results in more lenient normative judgments, but it is possible. Future research could investigate this possibility directly.

As discussed above, we did not examine non-prototypical forms of self-killing such as martyrdom, altruistic suicide, accidental death, and "suicide by cop". We chose to exclude these from our study for two reasons. First, our interest was in people's understanding of the normative status of "typical" cases of suicide, and the circumstances that would lead to these kinds of deaths seem to us to be anything but typical. Second, we suspected that – with the possible exception of suicide by cop – these actions may not be normatively opposed in the first place. There is arguably something praiseworthy about martyrdom and altruistic suicide, and Westerners, at least, do not typically condemn accidents (see, e.g., Cushman, 2008; McNamara et al., 2019). Indeed, accidentally killing oneself may not even be considered a form of suicide at all. As Kollareth et al. (2019, p. 554) put it, "suicide is implicitly intentional, and killing oneself accidentally is not called suicide." This is not to say that people do not engage in risky behaviors (e.g., drug use) that can lead to accidental deaths (e.g., by overdose) for many of the same reasons why they might commit suicide. Certainly, they do. But cases like this seem to lack the true intentionality that characterizes prototypical acts of suicide. To put it simply, *not wanting to live* seems to us to be a different psychological state than actively *wanting to die*. Future research could explore what specific

sorts of risky or harmful behaviors are categorized by laypeople as suicide, and which are not, and what implications this has for normative judgments of these acts.

In addition to contributing to our understanding of normative judgment, the present research may have real-world applications. Understanding what makes suicide seem so wrong may be valuable in designing new interventions to divert people away from committing it. For instance, knowing that one reason that suicide is opposed is that it is seen as causing harm to one's family and friends suggests that reminding suicidal individuals of this harm could help to dissuade them from attempting to kill themselves. At first glance, suicide prevention efforts like this, which focus on what is "wrong" about the act, may seem to be at odds with efforts to destigmatize suicide. However, we think that this is not the case. Specifically, it is valuable to draw a distinction between stigmatizing the *person* who attempts or commits suicide, versus opposing the act itself. Our reading is that efforts aimed at reducing stigma are very much directed towards the former. We are not aware of any serious attempts to reduce opposition to the *act* of suicide (i.e., to convince people that "suicide is okay"), except in the specific context of euthanasia cases. So, we do not see suicide prevention efforts that build upon our findings as contrary to efforts at destigmatization – we can try to prevent the act while still accepting and loving those who have attempted it. We leave applying this research to design real-world suicide-prevention interventions as an important task for future work.

4.3 A Tentative Theoretical Framework

The present research was exploratory in nature. We did not commit to a particular theoretical framework at the outset, but rather simultaneously tested multiple different explanations for normative opposition to suicide. This allowed us to uncover seven objections to suicide that all contribute to laypeople's opposition to it. That said, we now sketch a framework to make sense of our results, in the hopes that future work can develop this tentative, post-hoc idea into a fully formed theory of normative opposition to suicide.

Humans are social creatures, and the functional purpose of normative beliefs is thought to be to allow for cooperative group living (Curry et al., 2019, Haidt & Kesebir, 2010). If this is right, then perhaps a person's perceived value to the group drives how judgments of how wrong it would be for them to take their own life. This could explain several of our findings: a 90-year-old presents less potential for future cooperation than a 40-year-old, a person with no social ties presents little such opportunity, someone who has lived a life of criminality and deviance is clearly an undesirable cooperative partner, and someone with a chronic illness may be incapacitated, and therefore less able to contribute to cooperative group endeavors. It is less clear how this framework would explain our other results, however: someone who dies in a painful manner is not obviously a less desirable partner than someone who dies painlessly, nor is someone who commits suicide impulsively obviously a less desirable partner than someone who does so after much thought, nor is it clear how a person's own moral-religious background affects their value as a member of a

cooperative group. Nonetheless, we think at least some of our findings can be interpreted thusly: when evaluating the wrongness of a suicide, people attend to the value that the deceased could have provided to their group. High-value cooperative partners represent a greater loss to the group than low-value partners, making their deaths more wrong.

4.4 Conclusion

Despite the prevalence and costs of suicide, little prior research has investigated normative judgments of this act. We took an exploratory approach to this problem, testing ten possible explanations for why suicide is opposed. We found that seven of these explanations contribute to moral judgments of suicide, and that, taken together, they seem to fully explain these judgments, at least on average. In addition to broadening our understanding of lay normative judgments, this research can also contribute to new interventions to deter people from this irreversible act.

References

- Amundsen, D. W. (1989). Suicide and early Christian values. In B. A. Brody (Ed.), *Suicide and euthanasia* (pp. 77–153). Kluwer Academic Publishers.
- Angst, J., Angst, F., & Stassen, H. (1999). Suicide risk in patients with major depressive disorder. *Journal of Clinical Psychiatry*, *60*(2), 57–62.
- Beauchamp, T. L. (1989). Suicide in the age of reason. In B. A. Brody (Ed.), *Suicide and euthanasia* (pp. 183–219). Kluwer Academic Publishers.
- Boyle, J. (1989). Sanctity of life and suicide: Tensions and developments within common morality. In B. A. Brody (Ed.), *Suicide and euthanasia* (pp. 221–250). Kluwer Academic Publishers.
- Brody, B. A. (1989). A historical introduction to Jewish casuistry on suicide and euthanasia. In B. A. Brody (Ed.), *Suicide and euthanasia* (pp. 39–75). Kluwer Academic Publishers.
- Camus, A. (1991). *The myth of Sisyphus and other essays*. (O'Brien, J., Trans.). Alfred A. Knopf, Inc. (1942).
- Chakroff, A., Dungan, J., & Young, L. (2013). Harming ourselves and defiling others: What determines a moral domain? *PLoS ONE*, *8*, e74434.
- Coleman, L. (1987). *Suicide clusters*. Faber and Faber.
- Cooper, J. M. (1989). Greek philosophers on euthanasia and suicide. In B. A. Brody (Ed.), *Suicide and euthanasia* (pp. 9–38). Kluwer Academic Publishers.
- Cropper, M. L., Aydede, S. K., & Portney, P. R. (1994). Preferences for life saving programs: How the public discounts time and age. *Journal of Risk and Uncertainty*, *8*, 243–265.
- Curry, O. S., Mullins, D. A., & Whitehouse, H. (2019). Is it good to cooperate? Testing the theory of morality-as-cooperation in 60 societies. *Current Anthropology*, *60*, 47–69.

- Cushman, F. A. (2008). Crime and punishment: Distinguishing the roles of causal and intentional analyses in moral judgment. *Cognition*, *108*, 353–380.
- DeScioli, P., Gilbert, S. S., & Kurzban, R. (2012). Indelible victims and persistent punishers in moral cognition. *Psychological Inquiry*, *23*, 143–149.
- Docherty, A., Kious, B., Brown, T., Francis, L., Stark, L., Keeshin, B., ... & Coon, H. (2021). Ethical concerns relating to genetic risk scores for suicide. *American Journal of Medical Genetics Part B: Neuropsychiatric Genetics*, *186*, 433–444.
- Dolan, P., Shaw, R., Tsuchiya, A., & Williams, A. (2005). QALY maximisation and people's preferences: A methodological review of the literature. *Health Economics*, *14*, 197–208.
- Durkheim, E. (1979/1897). *Suicide: A study in sociology*. The Free Press.
- Fairbairn, G. (1995). *Contemplating suicide: The language and ethics of self-harm*. Routledge.
- Ferngren, G. B. (1989). The ethics of suicide in the Renaissance and Reformation. In B. A. Brody (Ed.), *Suicide and euthanasia* (pp. 155–181). Kluwer Academic Publishers.
- Gard, G., Cavlak, U., Sundén, B. T., & Ozdinler, A. R. (2005). Life-views and ethical viewpoints among physiotherapy students in Sweden and Turkey – A comparative study. *Advances in Physiotherapy*, *7*, 20–31.
- Goodwin, G. P., & Darley, J. M. (2008). The psychology of meta-ethics: Exploring objectivism. *Cognition*, *106*, 1339–1366.
- Goodwin, G. P. & Landy, J. F. (2014). Valuing different human lives. *Journal of Experimental Psychology: General*, *143*, 778–803.
- Gray, K. (2014). Harm concerns predict moral judgments of suicide: Comment on Rottman, Kelemen and Young (2014). *Cognition*, *133*, 329–331.
- Haidt, J. & Kesebir, S. (2010). Morality. In S. T. Fiske, D. T. Gilbert, & G. Lindzey (Eds.), *Handbook of social psychology* (pp.797–832). Hoboken, NJ: Wiley.
- Hecht, J. M. (2013). *Stay: A history of suicide and the arguments against it*. Yale University Press.
- Hester, N. & Gray, K. (2020). The moral psychology of raceless, genderless strangers. *Perspectives on Psychological Science*, *15*, 216–230.
- Harwitz, D., & Ravizza, L. (2000). Suicide and depression. *Emergency Medicine Clinics of North America*, *18*(2), 263–271.
- Johannesson, M., & Johannesson, P.-O. (1997). Is the valuation of a QALY gained independent of age? Some empirical evidence. *Journal of Health Economics*, *16*, 589–599.
- Kious, B. M., Docherty, A. R., Botkin, J. R., Brown, T. R., Francis, L. P., Gray, D. D., ... & Coon, H. (2021). Ethical and public health implications of genetic testing for suicide risk: family and survivor perspectives. *Genetics in medicine*, *23*, 289–297.
- Kollareth, D., Allam, A., & Russell, J. A. (2019). On judging the immorality of someone having taken his or her own life. *Social Cognition*, *37*, 547–570.
- Kollareth, D., & Russell, J. A. (2018). On an observer's reaction to hearing someone harm him or herself. *Psychological Studies*, *63*, 298–314.

- Landy, J. F. (2016). Representations of moral violations: Category members and associated features. *Judgment and Decision Making, 11*, 496–508.
- Landy, J. F., Walco, D. K., & Bartels, D. M. (2017). What's wrong with using steroids? Exploring whether and why people oppose the use of performance enhancing drugs. *Journal of Personality and Social Psychology, 113*, 377–392.
- Lewis, P. A., & Charny, M. (1989). Which of two individuals do you treat when only their ages are different and you can't treat both? *Journal of Medical Ethics, 15*, 28–34.
- Li, M., Vietri, J., Galvani, A. P., & Chapman, G. B. (2010). How do people value life? *Psychological Science, 21*, 163–167.
- Martin, J.W. & Cushman, F.A. (2015). To punish or to leave: Distinct cognitive processes underlie partner control and partner choice behaviors. *PLoS ONE, 10*, 1–14.
- Martin, J.W. & Cushman, F.A. (2016). Why we forgive what can't be controlled. *Cognition, 147*, 133–143.
- McNamara, R. A., Willard, A. K., Norenzayan, A. & Henrich, J. (2019). Weighing outcome vs. intent across societies: How cultural models of mind shape moral reasoning. *Cognition, 182*, 95–108.
- Montoya, A. K. & Hayes, A. F. (2017). Two-condition within-participant statistical mediation analysis: A path-analytic framework. *Psychological Methods, 22*, 6–27.
- Morris, S. B. & DeShon, R. P. (2002). Combining effect size estimates in meta-analysis with repeated measures and independent-groups designs. *Psychological Methods, 7*, 105–125.
- Mulgan, T. (2020). *Utilitarianism*. Cambridge University Press.
- NASMHPD. (2020). National Association of State Mental Health Program Directors Weekly Update. *National Association of State Mental Health, 6*(6).
- National Institute of Mental Health. (2021, May). *Suicide*. Mental Health Information. <https://www.nimh.nih.gov/health/statistics/suicide>
- Nucci, L., (1982). Conceptual development in the moral and conventional domains: Implications for values education. *Review of Educational Research, 52*, 93–122.
- Nucci, L., Guerra, N., & Lee, J. (1991). Adolescent judgments of the personal, prudential, and normative aspects of drug usage. *Developmental Psychology, 27*, 841–848.
- Nucci, L., & Turiel, E. (1993). God's word, religious rules, and their relation to Christian and Jewish children's concepts of morality. *Child Development, 64*, 1475–1491.
- Pizarro, D., Uhlmann, E., & Salovey, P. (2003). Asymmetry in judgments of moral blame and praise: The role of perceived metadesires. *Psychological Science, 14*, 267–272.
- Platt, J. R. (1964). Strong inference. *Science, 146*, 347–353.
- Ramírez Sierra, D. F. (2020). *Las razones detrás de la condena moral del suicidio. Una investigación empírica*. [Unpublished master's thesis]. Universidad Nacional de Colombia.
- Ratcliffe, J. (2000). Public preferences for the allocation of donor liver grafts for transplantation. *Health Economics, 9*, 137–148.

- Rodríguez, E., & Pinto, J. L. (2000). The social value of health programmes: Is age a relevant factor? *Health Economics*, *9*, 611–621.
- Rottman, J., Kelemen, D., & Young, L. (2014a). Tainting the soul: Purity concerns predict moral judgments of suicide. *Cognition*, *130*, 217–226.
- Rottman, J., Kelemen, D., & Young, L. (2014b). Purity matters more than harm in moral judgments of suicide: Response to Gray (2014). *Cognition*, *133*, 332–334.
- Sheehan, L. L., Corrigan, P. W., Al-Khouja, M. A., & Stigma of Suicide Research Team. (2017). Stakeholder perspectives on the stigma of suicide attempt survivors. *Crisis*, *38*, 73–81.
- Simmons, J., Nelson, L. D., & Simonsohn, U. (2013). Life after p-hacking. In S. Botti & Labroo, A. (Eds.), *Advances in consumer research, volume 41* (p. 775).
- Starmans, C. & Bloom, P. (2016). When the spirit is willing, but the flesh is weak: Developmental differences in judgments about inner moral conflict. *Psychological Science*, *27*, 1498–1506.
- Tisak, M., & Turiel, E. (1984). Children's conceptions of moral and prudential rules. *Child Development*, *55*, 1030–1039.
- Tsuchiya, A., Dolan, P., & Shaw, R. (2003). Measuring people's preferences regarding ageism in health: Some methodological issues and some fresh evidence. *Social Science & Medicine*, *57*, 687–696.
- Turiel, E. (1983). *The development of social knowledge: Morality and convention*. Cambridge, UK: Cambridge University Press.
- United Health Foundation. (2020). *America's Health Rankings analysis of CDC WONDER Online Database, Underlying Cause of Death, Multiple Cause of Death files, United Health Foundation*. America's Health Rankings. <https://www.americashealthrankings.org/explore/annual/measure/Suicide/state/ALL>
- Westefeld, J. S., & Furr, S. R. (1987). Suicide and depression among college students. *Professional Psychology: Research and Practice*, *18*(2), 119–123.
- Wilcox, H. C., Kuramoto, S. J., Lichtenstein, P., Långström, N., Brent, D. A., & Runeson, B. (2010). Psychiatric morbidity, violent crime, and suicide among children and adolescents exposed to parental death. *Journal of the American Academy of Child and Adolescent Psychiatry*, *49*, 514–523.
- Wreen, M. (1988). The definition of suicide. *Social Theory and Practice*, *14*, 1–23.