Folk Psychology and Proximal Intentions
Alfred Mele, Thomas Nadelhoffer, and Maria Khoudary

*This is a prepublication draft

Abstract:

There is a longstanding debate in philosophy concerning the relationship between intention and intentional action. According to the Single Phenomenon View, while one need not intend to A in order to A intentionally, one nevertheless needs to have an A-relevant intention. This view has recently come under criticism by those who think that one can A intentionally without any relevant intention at all. On this view, neither distal nor proximal intentions are necessary for intentional action. In this paper we present the results of two studies that explore folk ascriptions of proximal intentions and intentional actions in garden-variety, non-moral cases. Our findings suggest a very tight relationship between the two. We argue that the results from these two studies cohere with the Single Phenomenon View and give theorists who reject this view on conceptual grounds reason to worry.

Introduction

In everyday discourse, we often draw a distinction between actions that are performed intentionally (e.g., shutting a car door) and those that are not (e.g., shutting a car door on your finger). This is not a distinction without a difference. Whether an action was performed intentionally or not plays a crucial role in attributions of moral responsibility. Moreover, attributions of intentional action also help us navigate the world by helping us better understand

and predict people's behavior. The same is true of attributions of intentions. Intended actions are more indicative of our character and more predictive of our future behavior than actions that are not intended. However, while most philosophers agree that intentions and intentional actions play an important role in our folk psychology, there is wide-scale disagreement about the precise nature of this role. There is also disagreement about exactly how intentions and intentional actions are related to each other.

According to the Simple View of intentional action (SV), necessarily, agents intentionally *A* only if they have "an intention to *A*" (Bratman, 1987, p. 112). An alternative view, the Single Phenomenon View (SPV), is less demanding. Its thesis is that "to *A* intentionally I must intend to do something" (p. 113) but I do not need to intend to *A* (Bratman, 1987, p. 113). Some well-known studies by Joshua Knobe (2003, 2008) are widely regarded as showing that the majority of non-specialists make judgments that contradict SV, and, in our opinion, SV had already been falsified on theoretical grounds (see, for example, Michael Bratman's discussion of his video games scenario in 1987, ch. 8). SPV has not received nearly as much critical attention as SV, but some philosophers' claims about proximal intentions – in the simplest case, intentions to do something straightaway – do challenge SPV. One of us has examined philosophical arguments for these claims about proximal intentions elsewhere (XXXX 2019). Our concern here is to see how the claims at issue fare in the court of folk judgments. Our broader aim is to contribute to the growing body of evidence about folk conceptions of intention and intentional action by focusing on a kind of vignette that has received little attention in experimental philosophy. The vignettes featured here zoom in on cases in which the actions at issue are simple and require little to no thought. The vignettes are pertinent to some recent skeptical claims about proximal intentions.

1. Background

Bratman's characterization of SPV in the passage cited above is pretty casual. SPV can be formulated more rigorously and in a way that brings it into closer contact with his characterization of SV, which uses the word "intention." The following will do: "(i) necessarily, in any scenario in which an agent performs an intentional action, the agent has some relevant intention, and (ii) an agent who lacks an intention to *A* may *A* intentionally" (Mele, 2014, p. 38). Except for philosophers who deny the existence of states of mind, the move from "She intends to *A*" to "She has an intention to *A*" is smooth and easy, or so we assume in Study 1, reported below. This assumption is tested in Study 2.

There are two types of intention that are germane to our present discussion. Distal intentions are intentions to do something in the near or distant future. For instance, when one intends to go to the pub tomorrow for a pint, one has a distal intention to go to the pub (and a distal intention to have a pint). Proximal intentions, on the other hand, are focused on the present. Intentions to do something straightaway—e.g., to leave one's office for the pub *right now* – are proximal intentions. So are intentions to continue doing something that one is doing (for example, continue watching a movie after considering stopping) and intentions to start *A*-ing at once (for example, start swimming a lap right now) (Mele, 2009, p. 10).

We turn now to some restrictive claims about proximal intentions. Ann Bumpus conceives of "intentions as bound up with planning and means-end reasoning" (2001, p. 186). She asserts that if we follow her in this, "it would be unsurprising to find [that] the performance of many intentional actions" does not require an intention (p. 186). Her view is that although there are

proximal intentions, as when “Max’s distal intention to watch the *X-files* persists until Sunday night [and] issue[s] in a proximal intention to turn on the television now” (p. 189), there are not enough to go around. Bumpus highlights spontaneous intentional actions in this connection, contending that they are performed in the absence of any pertinent intention, including associated proximal intentions (pp. 184-190). She also claims, more generally, that although many distal intentions “lead to means-end reasoning and work to screen out incompatible options, most cases of intentional action do not involve this sort of reasoning” (p. 187). Her thesis is that the following proposition is false: “(IAP) If *S* A-s intentionally, then there is something *S* intends to do and it is the relation of *S*’s A-ing to a relevant intention that makes *S*’s A-ing intentional” (p. 177). IAP is a version of the Single Phenomenon View.

George Botterill seems to agree with Bumpus, and he may go even further than she does. He questions “whether folk psychology has any clear place for proximal occurrent intentions” (2010, p. 397), and he asserts that “the natural place for intentions within folk psychology is surely within some exercise of practical reasoning” (p. 397). We emphasize the expression “*the* natural place.” Should it be believed that intentions have just one natural place in folk psychology? We come back to this question in Section 2.

The thesis of a recent article by Marcela Herdova is that proximal intentions are not involved in a certain well-known style of neuroscience experiment (Herdova, 2016). In some experiments of this kind, participants are asked to perform an assigned task – for example, flex a wrist – at a time the participants arbitrarily select. In others, they arbitrarily select one of two buttons to press. (They are supposed to do these things many times over the course of an experiment. For representative experiments, see Fried, Mukamel, & Kreiman, 2011; Libet, 1985;

and Soon, Brass, Heinze, and Haynes, 2008.) The truth of the proposition that proximal intentions are not present in these studies is compatible with the truth of SPV (and Herdova does not reject SPV in Herdova, 2016). After all, the participants may intend to comply with the experimenter's instructions, and their intentions to do that may play a role in producing their intentional wrist flexes and button presses. But our concern in this connection is Herdova's thesis about proximal intentions in its own right and the way in which she defends it – namely, by arguing that the alleged proximal intentions in these neuroscience experiments do not have enough functions in common with paradigm proximal intentions to qualify as intentions. In a subsequent article, Herdova branches out to spontaneous actions not performed in laboratory settings and suggests that her "argument can form the basis for a stronger claim: that some intentional actions are not caused by intentions at all" (2018, p. 15).

One of us has argued elsewhere (XXXX2019) that whether some proximal intention or other is at work in every case of intentional action is, in part, a question for neuroscience. Here, we investigate a different question. What place do proximal intentions have in folk thinking about intentional action? We had three related goals in running the studies to be discussed here. First, we wanted to shed light on the relationship between folk ascriptions of intention and intentional action in mundane, non-moral cases.¹ Second, we wanted to shed light on how laypersons think specifically about proximal intentions. Third, by shedding light on ascriptions of proximal intentions, we wanted to shed light on the debate about whether SPV is overly demanding. After all, if, for example, SPV is falsified by ordinary cases of spontaneous intentional action *on*

---

¹ This represents a departure from the experimental philosophy literature, which tends to focus on moral behavior. We note that two of our nine cases arguably have some moral significance.

conceptual grounds, as Bumpus suggests, one may expect a hefty percentage of people to reject the proposition that their agents intended to perform them. We might expect the same if, as Botterill suggests, folk psychology has no clear place for proximal intentions and if, as Herdova contends, the alleged proximal intentions at work in producing the spontaneous intentional actions featured in the neuroscience experiments we mentioned do not have enough functions in common with paradigm proximal intentions to qualify as intentions. As we will now show, this expectation was not borne out.

2. Study 1

2.1. Hypotheses and General Design

As we've explained, SPV is a theory about the relationship between intentions and intentional actions. According to some critics of SPV, proximal intentions need not play a role in the production of intentional actions and, in some cases, no distal intention is at work either. On our view, folk intuitions ought to inform one's theory of the relationship between intentions and intentional actions, especially if one is interested in ordinary concepts rather than technical, philosophical concepts.² Given this metaphilosophical assumption, figuring out the right analysis of an ordinary concept is a partly empirical affair; it involves gathering data on folk concepts.

---

² We are not suggesting that we ought not be interested in philosophical rather than ordinary concepts. Rather, we are merely pointing out that if and when one is interested in the latter, one is beholden to data about folk intuitions. For more on the relevance of folk intuitions to philosophy, see Nadelhoffer and Nahmias (2007); cf. Kauppinen (2007).

In designing the vignettes for this study, we wanted to focus on mundane, morally-neutral cases involving some clearly intentional actions as well as some clearly unintentional actions. We also wanted to ask questions that enabled us to get at whether participants judged that the agents in these cases intended to perform the actions under consideration at once or straightaway. Our prediction was that intuitions about intentions would correlate with intuitions about intentional actions. If this is how things turn out, philosophers who reject SPV on conceptual grounds have reason to worry that they have an unduly restrictive conception of intention.

To test our prediction, we used a within-subject experimental design whereby each participant was presented with a series of nine scenarios. Each scenario involved an agent who performed actions that did not require much, if any, deliberation. For each scenario, participants were presented with four statements. Two of the statements were about an action from the scenario that we thought could plausibly be described as intended and intentional and two of the statements were about an action that we thought could be plausibly described as not intended and not intentional. Our prediction was that participants would find the former to be both intended and intentional and the latter to be neither intended nor intentional—that is, we expected ascriptions of intention and intentional action to track our a priori evaluation of these actions. We took this tack to control for a potential worry about how discerning our participants may (or may not) be. In short, we wanted to be able to ensure that participants were adequately discriminating between actions that are intended and intentional and actions that are neither intended nor intentional. So, we included instances of both types of action as a safeguard. We figured that if our findings turned out to support our expectations—which they did—this would

show that (a) most participants discriminate between intended and intentional actions and actions that are neither intended nor intentional, and (b) most participants judge that intentions and intentional actions go together even in the kinds of cases that critics of SPV use to support their view.

This approach made our analysis simple and straightforward. Actions that elicited answers above the midpoint of a 7-point Likert scale were judged to be intended and intentional, respectively, and actions that elicited answers below the midpoint were judged not to be intended or intentional, respectively. So, below we report the mean scores and standard deviations for each question. From these scores, one can read off the average assessments of intention and intentionality. Our question is majoritarian: Do most participants judge that intentions and intentional actions go together or do they think the two come apart? As we will see, most responses about intentions and intentionality are of a piece, as SPV would predict regarding the vignettes we tested. When most participants judged that an action was intentional, they also judged that it was intended (and vice versa). When most participants judged that an action was not intentional, they also judged that it was not intended (and vice versa). This pattern held true across all eight scenarios that we discuss here.

2.2. Participants

This study was approved by the College of Charleston Institutional Review Board. The data set and Supplemental Materials (which contain all stimuli, measures, and figures) can be found on our OSF page at: <https://osf.io/vyjrf/>. Participants were recruited via Amazon Mechanical Turk

(MTurk).³ As a pre-requisite, all participants needed to have completed at least 500 surveys on MTurk in the past with a success rate of at least 97%. The other two requirements were that participants had to be at least 18 years old and located in the United States. Participants were paid \$1 for completing the survey—which took an average of 14 minutes and 4 seconds. 213 participants completed the survey. Once we eliminated participants who failed one of two attention checks, we ended up with 197 participants. As part of this study, we also collected basic demographic information: 71% of participants were white ($n = 140$), 53% were men ($n = 104$), the oldest participant was 74, and the youngest were 18 ($M = 36.3$, $SD = 11.5$).

2.3. Materials and Results

Our aim was to explore how people ascribe intentions and intentional actions in some pedestrian cases that bear on the philosophical claims about proximal intentions reported in section 1. Participants were randomly assigned to one of two conditions. In one condition, they read the following cover story and instructions:

Imagine that Gunnar is a college student from Sweden who has moved to the United States with the goal of becoming a perfect speaker of English. We have agreed to help him achieve his goal, and we're asking for your assistance with this project. Your task has to do with how the words "intended" and "intentionally" are used. Other groups will deal with other words. With Gunnar's situation in mind, please read the following nine scenarios. Then, please tell us whether you disagree or agree with the statements that

---

³ MTurk is an online survey service that enables researchers to recruit and pay for participants for completing surveys of studies. For findings concerning the benefits of using MTurk—including the quality of the data and the improved diversity of the participant pool—see Burhmester, Kwang, & Gosling (2011); Paolacci, Chandler, & Ipeirotis (2010); Rand (2012)..

follow. A neutral response is available in case you neither agree nor disagree. The other options are: strongly agree, agree, somewhat agree, somewhat disagree, disagree, strongly disagree. For each scenario, you will have the opportunity to scroll back up and read the scenario while you're responding to the statements. Once you've finished responding to the statements for a particular scenario, you will be able to advance to the next scenario.

In the second condition, participants instead read the following cover story and instructions:

We are interested in how the words “intended” and “intentionally” are used in ordinary English. We are asking for your assistance with this. What we want to know is how you yourself use these words. Keep our goal in mind and then please read the following nine scenarios. Then, please tell us whether you disagree or agree with the statements that follow. A neutral response is available in case you neither agree nor disagree. The other options are: strongly agree, agree, somewhat agree, somewhat disagree, disagree, strongly disagree. For each scenario, you will have the opportunity to scroll back up and read the scenario while you're responding to the statements. Once you've finished responding to the statements for a particular scenario, you will be able to advance to the next scenario.

These two cover stories represented our attempt to make it clear to participants that we were interested in understanding how they normally use the words “intended” and “intentionally.” We were not sure which cover story would be the most effective; so, we tried them both. In each of the two conditions, participants received the same set of scenarios and the same statements. The only difference between the two conditions was the cover story. Because a linear mixed

effects model (with random intercepts for participant and scenario) returned no difference in rating between conditions ($\beta = 0.02$, $SE = 0.06$, $t(6295) = 0.33$, $p = 0.74$), we collapsed the two conditions for the purposes of analysis.⁴

Below are (a) eight of the nine scenarios we used,⁵ (b) four of the statements we used about them,⁶ and (c) the means and standard deviations for each statement for all 197 participants. For each scenario, there are two pairs of statements and for each pair one statement has the form, “S intended to A,” and the other has the form “S intentionally A-ed.” Our prediction was that most participants would agree that the action described in the first pair was both intentional and intended and that they would disagree that the action described in the second pair was either intentional or intended. The order of presentation of the scenarios was

---

⁴ All statistical analyses for Studies 1 and 2 were computed in R Version 4.0 (R Core Team, 2020). With the lme4 package (Bates et al., 2015; version 1.1-23), we used linear mixed effects models to analyze main effects and interactions among our variables. All reported models had random intercepts for participants and scenario, meaning that trends in response patterns were free to vary in their origin on a per subject and per scenario basis. Including random intercepts for these dimensions of our study controls for any particular idiosyncrasies of our sample or stimuli, thus increasing our ability to generalize our findings beyond these specific experiments. All models were fit with restricted maximum likelihood estimation (REML) and t-tests of the estimates use Satterthwaite’s method. Degrees of freedom were calculated using the Kenward-Roger method. Plots were made with the ggplot2 package (Wickham, 2016; version 3.3.0) and mixed model summary tables with the sjPlot package (Lüdtke, 2020; version 2.8.3.1). Code for the analysis and plots can be found at <https://github.com/maria-khoudary/folk-psych-proximal-intentions>.

⁵ Originally, we envisioned a broader project than the one we eventually decided to undertake here. So, our study involved an additional case that focused on act individuation—e.g., whether people intend each of the individual steps they take while walking. While we do not report those findings here, they, too, were perfectly consistent with SPV. The data for that study can be found on our aforementioned OSF page.

⁶ For some scenarios, we included more than the four statements. We ultimately decided to exclude some of those statements from our analysis in order to streamline our presentation of the results. Sometimes this was because we didn’t pair a particular statement about “intended to x” with a statement about “intentionally x-ing” and sometimes it was because we included two (or more) items that we expected participants to clearly judge both to be intended and intentional in the same scenario—e.g., in Scenario 1, AI’s unlocking his door, turning on his computer, and checking his email. In these cases, we just focused on one of the relevant items so that for our analysis, each scenario had two questions about intention and two questions about intentionality that we expected would pull people’s intuitions in opposite directions based on the actions in question—e.g., AI turning on his computer vs. AI making his computer crash. This post hoc strategy corrected two design flaws on our part that we failed to identify until we had already collected the data. All of the data, stimuli, and dependent variables we used can be found in the Supplemental Materials on our aforementioned OSF page.

randomized as was the order of the statements within each scenario. Responses were provided on a 7-point Likert scale ranging from 1 = strongly disagree to 7 = strongly agree. As such, scores between 5-7 indicate that participants judged that the action was intentional and/or intended whereas scores between 1-3 indicate that participants judged that the action was not intentional and/or not intended. Here are the materials we used plus the basic findings:

Scenario 1: Today, like many days, Al drives to work, walks up the stairs to his office, unlocks the door, turns on his computer, and then starts checking his e-mail. When he starts checking e-mail, his computer crashes.

1. When he turned on his computer, Al intended to do that ($M = 6.45$, $SD = 1.21$).
2. Al intentionally turned on his computer ($M = 6.61$, $SD = 0.90$).
3. When he made his computer crash, Al intended to do that ($M = 1.40$, $SD = 0.98$).
4. Al intentionally made his computer crash ($M = 1.45$, $SD = 1.12$).

Scenario 2: Beth is driving to work when a dog darts in front of her car. She slams on the brakes to avoid hitting the dog, and in the process her lunch falls off the front seat onto the floor.

1. When she slammed on her brakes, Beth intended to do that ($M = 5.94$, $SD = 1.57$).
2. Beth intentionally slammed on her brakes ($M = 6.32$, $SD = 1.25$).
3. When she made her lunch fall onto the floor, Beth intended to do that ($M = 1.50$, $SD = 1.14$).
4. Beth intentionally made her lunch fall onto the floor ($M = 1.60$, $SD = 1.21$).

Scenario 3: Carlos is a college student participating in a neuroscience experiment. There are two buttons in front of him, a blue one and an orange one. His task is to press whichever button he feels like pressing while scientists measure his brain waves. Carlos doesn't know any more about brain waves than you do. This time, he presses the orange button and he produces a brain wave of type Q.

1. When he pressed the orange button, Carlos intended to do that ($M = 6.27, SD = 1.06$).
2. Carlos intentionally pressed the orange button ($M = 6.28, SD = 1.15$).
3. When he pressed the orange button, Carlos intended to produce the brain wave he produced ($M = 2.08, SD = 1.58$).
4. Carlos intentionally produced the brain wave he produced ($M = 2.18, SD = 1.58$).

Scenario 4: Debby is a college student participating in a neuroscience experiment. Her task is to watch a fast clock, press the Z key on a keyboard whenever she feels like it while scientists measure her brain waves, and then tell the scientists where the hand was on the clock when she pressed the key. Debby doesn't know any more about brain waves than you do. Although Debby knows she has to press the Z key sooner or later, it is up to her when to do it. When she pressed the Z key, she produced a brain wave of type Q.

1. When Debby pressed the Z key, she intended to do that ($M = 6.35, SD = 1.10$).
2. Debby intentionally pressed the Z key ($M = 6.45, SD = 0.98$).
3. When Debby pressed the Z key, she intended to produce the brain wave she produced ($M = 2.19, SD = 1.60$).
4. Debby intentionally produced the brain wave she produced ($M = 2.16, SD = 1.59$).

Scenario 5: Ezra walked into his house and turned on his kitchen light by flipping a switch. Ezra didn't know that a burglar was in his house. When the burglar saw the light go on, the burglar was frightened and left through a back window.

1. When he flipped the switch, Ezra intended to do that ($M = 6.57, SD = 0.86$).
2. Ezra intentionally flipped the switch ($M = 6.58, SD = 0.86$).
3. When he frightened the burglar, Ezra intended to do that ($M = 1.54, SD = 1.07$).
4. Ezra intentionally frightened the burglar ($M = 1.55, SD = 1.07$).

Scenario 6: Filomena is stopped at a red light in an old car that she has been driving for many years. When the light turns green, she presses down on her gas pedal. She doesn't know it, but there's a problem with her engine. By pressing on the gas pedal, Filomena releases a lot of exhaust fumes into the air and irritates the cyclist behind her.

1. When she pressed down on her gas pedal, Filomena intended to do that ($M = 6.45, SD = 1.04$).
2. Filomena intentionally pressed down on her gas pedal ($M = 6.53, SD = 0.93$).
3. When she released a lot of exhaust fumes, Filomena intended to do that ($M = 1.72, SD = 1.34$).
4. Filomena intentionally released a lot of exhaust fumes ($M = 1.71, SD = 1.38$).

Scenario 7: When Greg drives to work, he always turns left at James St., unless the road is blocked. Today, the road isn't blocked. He turns left at James St., as usual, and, out of habit, he uses his turn signal. He also drives over a small, empty plastic bag that he notices but doesn't care about.

1. When he used his turn signal, Greg intended to do that ($M = 6.07, SD = 1.32$).
2. Greg intentionally used his turn signal ($M = 5.97, SD = 1.49$).
3. When he drove over the plastic bag, Greg intended to do that ($M = 4.2, SD = 2.08$).
4. Greg intentionally drove over the plastic bag ($M = 4.38, SD = 2.08$).

Scenario 8: As Ivan was walking to the library, a skateboarder almost knocked him over. Ivan immediately shouted an obscenity at the skateboarder. A professor who was walking behind Ivan was upset by the obscenity.

1. Ivan intended to shout at the skateboarder ($M = 5.55, SD = 1.72$).
2. Ivan intentionally shouted at the skateboarder ($M = 5.98, SD = 1.44$).
3. Ivan intended to upset the professor ($M = 1.63, SD = 1.23$).
4. Ivan intentionally upset the professor ($M = 1.53, SD = 1.02$).

To examine a potential Language (intended vs. intentionally) x Presumed Intentionality (intentional vs. not intentional) interaction, we ran a linear mixed effects model with random intercepts for participant and scenario. We found a large main effect of Presumed Intentionality, such that those scenarios we expected could not plausibly be identified as intended or intentional were indeed less likely to be identified as such ($\beta = -4.17, SE = 0.05, t(6097) = -85.0, p < .001$). We found a smaller main effect of Language ($\beta = 0.13, SE = 0.05, t(6097) = 2.74, p = .006$), indicating that actions described as being performed intentionally were rated as slightly more intentional than actions described as intended. There was no Language x Presumed Intentionality interaction (see Supplementary Materials for model formula and summary). To further probe participants' judgments about actions described as intentional vs. those described as intended, we computed

a Pearson's product-moment correlation between each participant's mean rating for intended actions and mean rating for intentional actions per scenario. The correlations per scenario ranged from 0.30 to 0.76, and the overall correlation was strong and significant (Pearson's $r = 0.63$, $t(1574) = 32.75$, $p < .001$; see Supplementary Materials for a table with correlations for each scenario). Further, there were no scenarios wherein participants, on average, described an action as intentional but not intended (see Figure 1).

Figure 1: Scenario means for Study 1. Average response across participants for each condition in each scenario. Responses to “Intended” items are plotted in gray, and responses to “Intentionally” items are plotted in white. Error bars are the standard error. The dashed line at 4 is the midpoint, for which the intentionality judgment was neutral.

2.4. Discussion

Taken together, these results indicate that participants agreed that actions we presumed to be intentional and intended were both intentional and intended, and they agreed that actions we presumed to be neither intentional nor intended were neither intentional nor intended. For our purposes, however, the most important finding was that actions that were judged to be intentional were almost invariably judged to be intended since this speaks to the criticisms of SPV we mentioned earlier. While it is true that there was variation in how strongly participants agreed

(or disagreed) that the actions in question were intentional and/or intended, average responses were always on the same side of the midpoint for all cases (see Supplemental Materials). Moreover, only a very small percentage of responses (between 1% and 6%, depending on the case) are such that actions judged to be intentional were not judged to be intended (or vice versa). In short, judgments about intentional actions almost always align with judgments about intentions—which is precisely what a proponent of SPV would predict.

Two possible lines of objection to Study 1 merit attention. First, it might be claimed that we did not make it sufficiently clear that the intentions at issue are *proximal* intentions. We used statements of the form “When *S* *A*-ed, *S* intended to do that,” and one might worry that participants simply ignore the word “when.” One might also worry that participants who register the “when” might not assume that the agent intended to perform the action straightaway. Second, it may be argued that asking whether an agent intended to *A* is an inadequate way of discovering whether respondents believe that an had an intention to *A* and that a better strategy is simply to ask whether the agent had that intention. In the latter connection, we observe that some philosophers are skeptical of the claim that token states of mind, such as intentions and beliefs, exist (Child, 1994; Hornsby, 1993; Steward, 1997). These philosophers are happy to say that people intend to do things but not that people have intentions to do things. Such skepticism is not expressed in the work by Botterill, Bumpus, and Herdova that we have discussed, and, of course, it is skepticism about the existence of intentions in general (as states of mind), not just about proximal intentions. The skeptics about token states of mind might nevertheless think that our move from “*S* intends to *A*” to “*S* has an intention to *A*” is unwarranted. In response to these two worries, we ran a second study.

3. Study 2

3.1. Hypotheses and General Design

Study 2 only differed in a few ways from Study 1. For starters, we wanted to make an effort to address a recent worry that has arisen when it comes to data collected on Mturk—namely, that bots are being used to illicitly earn money. These bots may have the ability to pass some common attention checks and they corrupt surveys with bad data. So, we took preemptive steps to try to prevent bots from participating. At best, a bot would have a 1/36 chance of randomly getting both of the two items we used correct (see below for details). So, while we certainly didn't eliminate the possibility of bots, we took active steps to make bot participation much less likely. Only three participants failed one of the two bot checks—and, as a result, they were unable to participate in the survey.⁷ Second, because we found no differences between the two main conditions in Study 1, we only used the second cover story for Study 2. Like Study 1, Study 2 also had a within-subject design. All participants received the same eight scenarios we reported in Study 1. Finally, rather than framing the statements about intention in terms of what *S* “intended,” we framed the statements as “*S* had an intention to *A* right then.” In this way, we respond to both aforementioned worries.

3.2. Participants

This study was approved by the College of Charleston Institutional Review Board. Here again, the data set and supplemental materials (which contain all stimuli and measures) can be found on

---

⁷ For recent discussions of the potential problem with bots on MTurk, see: *New Scientist* (2018): <https://www.newscientist.com/article/2176436-bots-on-amazons-mechanical-turk-are-ruining-psychology-studies/>
Wired (2018): <https://www.wired.com/story/amazon-mechanical-turk-bot-panic/>

our OSF page at: <https://osf.io/vyjrf/>. Participants were recruited via MTurk. As a pre-requisite, all participants needed to have completed at least 500 surveys on MTurk in the past with a success rate of at least 97%. The other two requirements were that participants had to be at least 18 years old and located in the United States. Participants were paid \$1 for completing the survey—which took an average of 7 minutes and 27 seconds. 238 participants completed the survey. Once we eliminated participants who failed one of two attention checks, we ended up with 231 participants.⁸ As part of this study, we also collected basic demographic information: 71% of participants were white ($n = 165$), 63% of participants were men ($n = 165$), the oldest participant was 69 and the youngest were 18 ($M = 35.02$, $SD = 10.22$).

3.3. Materials and Results

Before participants could participate in the survey, they had to pass the following bot check:

Instructions:

Before you take the survey, we would like you to answer a few questions to help ensure you are a real person! Pay attention and answer carefully. If you respond incorrectly, you will be taken to the end of the survey and you will not be compensated.

1. Monday is the first day of the week. What is the third day of the week? Answer Choices:
Tuesday through Sunday.
2. January is the first month of the year. What is the fifth month of the year? Answer Choices: February through July.

---

⁸ Presumably, fewer participants failed the attention checks in Study 2 than in Study 1 because we told them up front this time that attention checks would be included and that failing them would mean they wouldn't be compensated (see below for instructions).

Any participants who missed either item were kicked out of the study prior to participation. We also used the same two attention checks that we used in Study 1 (which were once again randomly dispersed throughout the survey). Once participants passed the bot check, they were presented with a slightly modified form of the instructions that we used in the second condition of Study1—with the main goal of emphasizing that we were implementing attention checks:

Instructions:

We are interested in how the words “intention” and “intentionally” are used in ordinary English. We are asking for your assistance with this. What we want to know is how you yourself use these words. Keep our goal in mind and then please read the following eight scenarios. Then, please tell us whether you disagree or agree with the statements that follow. For each scenario, you will have the opportunity to scroll back up and read the scenario while you're responding to the statements. Once you've finished responding to the statements for a particular scenario, you will be able to advance to the next scenario. The eight scenarios we used in Study 2 were worded precisely the same as they were in Study 1, so we won't repeat them here. The primary difference between Studies 1 and 2 is the difference between one of the two dependent variables we mentioned earlier—namely, our statements about what was intended were replaced with statements of the following form: “When *S* *A*-ed, *S* had an intention to *A* right then” (except in Scenario 8, for a reason we provide below). The order of presentation of the scenarios was once again randomized as was the order of the statements within each scenario. Responses were provided on a 7-point Likert scale ranging from 1 = strongly disagree to 7 = strongly agree. Here are the materials we used plus the basic findings. Here again, scores between 5-7 indicate that participants judged that the action was intentional and/or

intended whereas scores between 1-3 indicate that participants judged that the action was not intentional and/or not intended:

Scenario 1: Al

1. When he turned on his computer, Al had an intention to do that right then ($M = 6.09$, $SD = 1.40$).
2. Al intentionally turned on his computer ($M = 6.21$, $SD = 1.24$).
3. When he made his computer crash, Al had an intention to do that right then ($M = 2.30$, $SD = 1.90$).
4. Al intentionally made his computer crash ($M = 2.08$, $SD = 1.74$).

Scenario 2: Beth

1. When she slammed on her brakes, Beth had an intention to do that right then ($M = 5.76$, $SD = 1.55$).
2. Beth intentionally slammed on her brakes ($M = 6.07$, $SD = 1.32$).
3. When she made her lunch fall onto the floor, Beth had an intention to do that right then ($M = 2.40$, $SD = 1.83$).
4. Beth intentionally made her lunch fall onto the floor ($M = 2.20$, $SD = 1.72$).

Scenario 3: Carlos

1. When he pressed the orange button, Carlos had an intention to do that right then ($M = 5.84$, $SD = 1.41$).
2. Carlos intentionally pressed the orange button ($M = 5.96$, $SD = 1.29$).

3. When he pressed the orange button, Carlos had an intention to produce the brain wave he produced right then ($M = 2.85, SD = 1.93$).
4. Carlos intentionally produced the brain wave he produced ($M = 2.91, SD = 1.91$).

Scenario 4: Debby

1. When Debby pressed the Z key, she had an intention to do that right then ($M = 5.94, SD = 1.32$).
2. Debby intentionally pressed the Z key ($M = 6.16, SD = 1.14$).
3. When Debby pressed the Z key, she had an intention to produce the brain wave she produced right then ($M = 3.02, SD = 1.89$).
4. Debby intentionally produced the brain wave she produced ($M = 2.90, SD = 1.97$).

Scenario 5: Ezra

1. When he flipped the switch, Ezra had an intention to do that right then ($M = 5.99, SD = 1.42$).
2. Ezra intentionally flipped the switch ($M = 6.14, SD = 1.27$).
3. When he frightened the burglar, Ezra had an intention to do that right then ($M = 2.42, SD = 1.90$).
4. Ezra intentionally frightened the burglar ($M = 2.23, SD = 1.79$).

Scenario 6: Filomena

1. When she pressed down on her gas pedal, Filomena had an intention to do that right then ($M = 5.9, SD = 1.52$).

2. Filomena intentionally pressed down on her gas pedal ($M = 6.07, SD = 1.38$).
3. When she released a lot of exhaust fumes, Filomena had an intention to do that right then ($M = 2.40, SD = 1.91$).
4. Filomena intentionally released a lot of exhaust fumes ($M = 2.32, SD = 1.82$).

Scenario 7: Greg

1. When he used his turn signal, Greg had an intention to do that right then ($M = 5.62, SD = 1.45$).
2. Greg intentionally used his turn signal ($M = 5.60, SD = 1.45$).
3. When he drove over the plastic bag, Greg had an intention to do that right then ($M = 4.18, SD = 2.02$).
4. Greg intentionally drove over the plastic bag ($M = 4.33, SD = 2.05$).

Scenario 8: Ivan

1. Ivan had an intention to shout at the skateboarder ($M = 5.60, SD = 1.68$).
2. Ivan intentionally shouted at the skateboarder ($M = 5.79, SD = 1.56$).
3. Ivan had an intention to upset the professor ($M = 2.26, SD = 1.83$).
4. Ivan intentionally upset the professor ($M = 2.34, SD = 1.87$).

We ran another linear mixed effects model to investigate any potential Language x Presumed Intentionality interactions. We found a large main effect of Presumed Intentionality, such that the actions we expected could not plausibly be identified as intentional or as actions the agents intended to perform were more likely to be rated as such ($\beta = -3.11, SE = 0.05, t(7151) = -59.54, p < .001$). We again found a small effect of Language, such that actions described as intentional

were judged to be slightly more intentional than those described as having been performed by agents who had the intention to perform them right then ($\beta = 0.16$, $SE = 0.05$, $t(7151) = 3.01$, $p = .002$). Additionally, we found a Language x Presumed Intentionality interaction, such that presumably unintended actions were judged as *less* intentional than presumably unintended actions that were described as having been performed by agents who had the intention to perform them right then ($\beta = -0.23$, $SE = 0.07$, $t(7151) = -3.07$, $p = .002$; see Supplementary Materials for a model formula and summary). Again, to probe perceptions of “intentionally” vs. “intention”, we computed a Pearson’s product-moment correlation between each participant’s mean rating for these items per scenario. The correlations per scenario ranged from 0.57 to 0.82, and the overall correlation was strong and significant (Pearson’s $r = 0.68$, $t(1846) = 39.92$, $p < .001$; see Supplementary Materials for a table with correlations per scenario). Further, in no scenario did participants judge, on average, that an agent performed an action intentionally but did not have an intention to perform it or vice versa (see Figure 2).

Figure 2: Scenario means for Study 2. Average response across participants for each condition in each scenario. Responses to “Had intention” items are plotted in gray, and responses to “Intentionally” items are plotted in white. Error bars are the standard error. The dashed line at 4 is the midpoint, for which the intentionality judgment was neutral.

Thus, even with our new wording in Study 2—that is, “When *S* *A*-ed, *S* had an intention to *A* right then”—actions deemed to be intentional were almost always deemed to be actions the agents had intentions to perform, which here again speaks to the criticisms of SPV we mentioned earlier.

3.4. Discussion

For starters, the results from Study 2 are in line with the results from Study 1. Most importantly, actions that were judged to be intentional were also usually judged to be intended (and vice versa). While there was once again variation in how strongly participants agreed (or disagreed) that the actions in question were intentional and/or intended, average responses were always on the same side of the midpoint for all cases. Moreover, as before, only a very small percentage of responses (between 1% and 5%, depending on the case) are such that actions judged to be intentional were not judged to be intended or vice versa (see Supplemental Materials). Taken together, the results from Studies 1 and 2 provide evidence that judgments about whether an action was intended and whether it was intentional almost always go together in the eyes of the folk. To determine whether judgments about what agents have intentions to do (as in Study 2) are out of line with judgments about what agents intend to do (as in Study 1), we combined the data from both studies and ran a linear mixed effects model with an interaction between Scenario and Language as the fixed effect, and a random intercept for participant as the random effect. Because we were interested in this comparison specifically, we did not include judgments of actions described as being performed intentionally. There was no main effect of Language ($\beta = 0.27$, $SE = 0.17$, $t(6470) = 1.60$ $p = 0.11$). But there was a Language x Scenario interaction for the case of Greg, such that actions of his described as intended were judged as slightly more intentional than actions of his described as having had an intention ($\beta = 0.50$, $SE = 0.24$, $t(6425)$

= 2.13, $p = 0.033$). However, the estimated difference in judgments is quite small (0.5 response points), and the statistical significance of this scenario is an outlier relative to the other scenarios (see Supplementary Materials for model formula and summary).

Thus, in general, we find no difference between “*S* intended to *A*” and “*S* had an intention to *A*”—which quashes the worry that, among lay folk, judgments about what agents have intentions to do will be out of line with judgments about what agents intend to do. It appears that, for the overwhelming majority of our participants, there is a smooth and easy transition between *S*’s intending to *A* and *S*’s having an intention to *A* after all. In line with our predictions, most participants agreed with our assessment concerning whether or not particular actions are intended and intentional (e.g., turning on one’s computer vs. making one’s computer crash) and they also apparently attribute proximal intentions to agents who perform mundane, spontaneous, and morally neutral actions. Whether we worded the statements in terms of the agents’ intending to *A* when they *A*-ed or in terms of their having an intention to *A* right then, whenever most participants judged that an agent *A*-ed intentionally, they judged that he or she had a proximal intention to *A*.

4. General Discussion

So, with these findings in hand, it is time to apply our results to the aforementioned claims by Botterill, Bumpus, and Herdova. We start with Botterill, who questions “whether folk psychology has any clear place for proximal occurrent intentions” (2010, p. 397) and asserts that “the natural place for intentions within folk psychology is surely within some exercise of practical reasoning”

(p. 397).⁹ Now, of course, “proximal intention” is a technical term, and we are not interested in how non-specialists might use this term. Our way of getting at proximal intentions was simple. We did not ask respondents to react to such statements as “Al intended to unlock his door” or “Carlos intended to press the orange button.” Agreement with those statements is consistent with thinking that these agents did not have proximal intentions to do these things. For example, someone may think that, a minute or so before he pressed the orange button, Carlos had an intention to press that button the next time he felt like pressing a button (a minute later, as it turned out) but that he did not have an intention to press the orange button *when* he pressed it. This is why we presented respondents with such statements as “When he pressed the orange button, Carlos intended to do that.”¹⁰ And, in response to the worry that such statements do not make it sufficiently clear that *proximal* intentions are at issue, we switched to “had an intention to do that right then” statements in Study 2 – with results that reinforce our findings in Study 1.

If our studies are any indication, there is a clear place for proximal intentions in folk psychology. We take it that ordinary usage of terms such as “intended” and “intention” provides evidence about folk psychology, and our respondents were happy to attribute proximal intentions to agents in all eight vignettes. Also, concerning Botterill’s claim about “the natural place for intentions within folk psychology,” notice that there is no clear place for practical

---

⁹ Since Botterill is reacting to Mele (2009), we take it for granted that by “occurrent” in this context he means what Mele means by it. There Mele endorses a proposal that he made in an earlier publication: “there are two ways for an intention to A to be an *occurrent* intention at a time. One way is for it to be ‘suitably at work at that time in producing relevant intentional actions or in producing items appropriate for the production of relevant intentional actions’; and the other is, roughly, for it to be a conscious intention at that time, provided that the intention ‘is not wholly constituted by a disposition to have occurrent intentions to A’” (2009, p. 4). Under normal conditions, all proximal intentions are suitably at work.

¹⁰ There is one exception. Because Ivan *immediately* shouted an obscenity at the skateboarder who almost knocked him over, there is no room for confusion about the timing of a relevant intention attributed to him.

reasoning in our vignettes. (If someone wishes to claim that there is very rapid practical reasoning in the vignette in which Beth hits the brakes, we do not object. But this might not have been the sort of thing Botterill had in mind when making his claim about practical reasoning.)

Bumpus, as we mentioned, highlights spontaneous actions in her critique of a version of SPV, and the critique is made on conceptual grounds. Carlos and Ivan perform spontaneous actions, and Debby seemingly spontaneously picks a time to press the Z key. Even so, our respondents report that, when these agents performed these actions, they intended (or had an intention) to do so. Of course, this does not prove that Bumpus is wrong about SPV. But it does contribute to a case for the judgment that she sets an uncommonly high bar for the presence of intentions (on this, see XXXX 2019).

As we reported, the thesis of Herdova (2016) is that proximal intentions are not involved in a certain kind of neuroscience experiment, and she defends this thesis by arguing that the alleged proximal intentions at work in these experiments do not have enough functions in common with paradigm proximal intentions to qualify as intentions (for a reply to this argument, see XXXX 2019). Our Carlos and Debby vignettes were designed to get evidence about lay reactions to such scenarios. The responses to “When he pressed the orange button, Carlos intended to do that” and “When Debby pressed the Z key, she intended to do that” were strongly positive, with means of 6.27 and 6.35, respectively. And in Study 2 the results were similar, with “intention” means of 5.84 for Carlos and 5.94 for Debby. Obviously, this does not prove that Herdova’s claim about experiments of this kind is false. In this connection, we observe that one of us has argued elsewhere that whether proximal intentions are present in these scenarios is an open empirical question requiring neuroscientific investigation (XXXX 2019), and all three of us

are open to the possibility that Herdova happens to be right about the absence of proximal intentions in these experiments (though not for the conceptual reasons she offered; see XXXX 2019). Again, our goal here is to contribute to the growing body of evidence about lay conceptions of intention and intentional action by focusing on scenarios pertinent to some recent skeptical claims about proximal intentions, scenarios of a kind that has received little attention in the experimental philosophy literature. Even so, we point out that, as with Bumpus, our findings contribute to a case for the verdict that Herdova sets an unusually high bar for the presence of intentions (see XXXX 2019).

One might worry that even such statements as “When he turned on his computer, Al had an intention to do that right then” might not be read by participants as we intended. Someone might suggest that a substantial number of participants might read our statement as reporting that Al had an intention right then to turn on his computer at a later time – that is, at some time later than the time at which he actually turned it on – rather than attributing to him a proximal intention to turn it on. We find this unlikely. First, the worrisome reading is not the most natural reading of the statement. Second, if readers were thinking that Al turned his computer on earlier than he intended to, why would the mean ratings for “intentionally” be so high? Things that we do earlier than we intend to do them are often good candidates for being unintentional actions. Some examples: “I intended to get off the interstate at exit 99; but without changing my mind, I got off at exit 89.” “It was my intention to send the signal at 2:00; but without changing my mind, I sent it at 1:00.” Third, in two of the scenarios (2 and 8), it would make no sense for the agent to intend to perform the action later than he or she performed it, and our findings for those scenarios are in line with our findings for the others. This is some reason to believe that

participants interpreted the statements of intention in, for example, Al's scenario as they did in these scenarios.

An insightful referee identified a pair of limitations of our studies. First, because participants were responding to so many related vignettes, they might at some point, as the referee put it, "start providing their answers at least in part based on recognizing certain patterns (rather than judging each case independently." Second, our simple stories about Carlos and Debby do not fully bring out the spontaneity that researchers seek to elicit in neuroscience experiments of the kind at issue. These issues may be taken up in future work. One way to deal with the first issue is to present participants with only one or two cases. The results can be compared with our results. In response to the second issue, one may test beefed-up versions of the Carlos and Debby scenarios that emphasize the spontaneity of the agents. In the case of both issues, the results can be compared with our results.

We should acknowledge another limitation of our present studies before closing. While participants recruited through MTurk tend to be more diverse and representative than usual convenience samples drawn from universities—which tend to be skewed especially when it comes to age—we nevertheless don't want to draw any hasty generalizations. After all, our online participants were drawn from a country that is Western, Educated, Industrialized, Rich, and Democratic (WEIRD: Henrich, Heine, & Norenzayan, 2010). Given that most of the work on folk intuitions about intentions and intentional action has been done using WEIRD participants—and given that we were trying to respond to and build upon this work—it made sense for us to follow suit in this regard. While there are a limited number of cross cultural studies that have explored the folk concept of intentional action (e.g., Knobe & Burra, 2006; Robbins, Shepard, &

Rochat, 2017), much work on this front remains to be done. By our lights, this is a ripe area for future research. But trying to address this issue would take us beyond the limited scope of our present investigation.

This limitation notwithstanding, we believe that our findings show that proximal intentions certainly seem to be alive and well in folk psychology (at least in our target population). Moreover, our results turn up no problems for SPV. Whether there are good reasons to deny that proximal intentions are actually at work in some or all of the vignettes we tested (whatever lay folk may think about this) or good reasons to reject SPV are topics for other occasions.¹¹¹²

References

- Botterill, G. (2010). Review of *Effective intentions: The power of conscious will*. *Analysis*, 70, 395–398
- Bratman, M. (1987). *Intention, plans, and practical reason*. Cambridge, MA: Harvard University Press.
- Bumpus, A. (2001). Actors without intentions: The Double Phenomena View. *Philosophical Studies*, 103(2), 177–199.
- Burhmester, M., Kwang, T., & Gosling, S. (2011). Amazon’s Mechanical Turk: A new source of inexpensive, yet high-quality data? *Perspectives on Psychological Science*, 6(1), 3–5.
- Child, W. (1994). *Causality, interpretation, and the mind*. New York: Oxford University Press.
- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48.
- Fried, I., Mukamel, R., & Kreiman, G. (2011). Internally generated preactivation of single neurons in human medial frontal cortex predicts volition. *Neuron*, 69(3), 548–562.
- Henrich, J., Heine, S.J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33, 61–83. <https://doi.org/10.1017/S0140525X0999152X>

---

¹¹ For one of those occasions, see XXXX 2019.

¹² We would like to thank Adam Feltz and an anonymous referee for very helpful feedback on an earlier draft of this paper. Working through their suggestions greatly improved our paper.

- Herdova, M. (2016). Are intentions in tension with timing experiments? *Philosophical Studies*, 173(3), 573–587.
- Herdova, M. (2018). Trigger warning: No proximal intentions required for intentional action. *Philosophical Explorations*, 21(3), 364–383.
- Hornsby, J. (1993). Agency and causal explanation. In J. Heil & A. Mele (Eds.), *Mental causation* (pp. 161–185). Oxford: Clarendon Press.
- Kauppinen, A. (2007). The rise and fall of experimental philosophy. *Philosophical Explorations*, 10(2), 95–118.
- Knobe, J. (2003). Intentional action and side-effects in ordinary language. *Analysis*, 63, 190–94.
- Knobe, J. (2008). The concept of intentional action: A case study in the uses of folk psychology. In J. Knobe & S. Nichols (Eds.), *Experimental philosophy* (pp. 129–148). New York: Oxford University Press.
- Knobe, J., & Burra, A. (2006). The folk concept of intention and intentional action: A cross-cultural study. *Journal of Cognition and Culture*, 6(1-2), 113–132.
- Libet, B. (1985). Unconscious cerebral initiative and the role of conscious will in voluntary action. *Behavioral and Brain Sciences*, 8, 529–66.
- Lüdecke D (2020). *sjPlot: Data visualization for statistics in social science*. R package version 2.8.4, <https://CRAN.R-project.org/package=sjPlot>.
- Mele, A. (2009). *Effective intentions*. New York: Oxford University Press.
- Mele, A. (2014). The Single Phenomenon View and experimental philosophy. In M. Vargas & G. Yaffe (Eds.), *Rational and social agency: The philosophy of Michael Bratman* (pp. 38–55). Oxford: Oxford University Press.
- Nadelhoffer, T. & Nahmias, E. (2007). The past and future of experimental philosophy. *Philosophical Explorations*, 10(2), 123–149.
- Paolacci, G., Chandler, J., & Ipeirotis, P. (2010). Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making*, 5, 411–419.
- R Core Team (2020). R: A language and environment for statistical computing. R Foundation for statistical computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Rand, G. (2012). The promise of Mechanical Turk: How online labor markets can help theorists run behavioral experiments. *Journal of Theoretical Biology*, 299, 172–179.
- Robbins, E., Shepard, J., & RoCHAT, P. (2017). Variations in judgments of intentional action and moral evaluations across eight cultures. *Cognition*, 164, 22–30.

Soon, C. S., Brass, M., Heinze, H.-J., & Haynes, J. D. (2008). Unconscious determinants of free decisions in the human brain. *Nature Neuroscience*, 11, 543–545.

Steward, H. (1997). *The ontology of mind: Events, processes, and states*. London: Clarendon Press.

Wickham, H. (2016). *ggplot2: Elegant graphics for data analysis*. New York: Springer-Verlag.

XXXX. 2019.