# Probability and Intentional Action

Spencer Ericson
*Department of Psychology, University of Waterloo*

Stephanie Denison
*Department of Psychology, University of Waterloo*

John Turri
*Department of Philosophy and Cognitive Science Program, University of Waterloo*

Ori Friedman
*Department of Psychology, University of Waterloo*

Research was supported by grants from the Natural Sciences and Engineering Research Council of Canada awarded to OF and SD. Correspondence concerning this article should be addressed to Ori Friedman, Department of Psychology, University of Waterloo, 200 University Avenue W, Waterloo, Ontario, Canada N2L 3G1. E-mail: *friedman@uwaterloo.ca*

**Open access version of**

Ericson, S. R., Denison, S., Turri, J., & Friedman, O. (2023). Probability and intentional action. *Cognitive Psychology*, *141*, 101551. https://doi.org/10.1016/j.cogpsych.2023.101551

***This document may not fully correspond with the version published in the journal****

**Abstract**

How does probability affect attributions of intentionality? In five experiments (total N = 1410), we provide evidence for a probability raising account holding that people are more likely to see the outcome of an agent's action as intentional if the agent does something to increase the odds of that outcome. Experiment 1 found that high probability without probability raising does not suffice for strong attributions of intentionality. Participants were more likely to conclude a girl intentionally obtained a desired gumball from a single gumball machine when it offered favorable odds for getting that kind of gumball compared with when it offered poor odds, but their attributions of intentionality were lukewarm. Experiments 2 and 3 then found stronger attributions of intentionality when the girl raised her probability of success by choosing to use machines offering favorable odds over machines offering poor odds. Finally, Experiments 4 and 5 examined whether these effects of probability raising might reduce to consideration of agents' beliefs and expectations. We found that although these mental states do matter, probability raising matters too—people attribute intentional actions to agents who increase their odds of success, rather than to agents who merely become convinced that success is likely. We discuss the implications of these findings for claims that control and skill contribute to attributions of intentional action.

Keywords: intentional action; probability raising; choice; theory of mind; probability

## Probability and Intentional Action

"When Brown in an ordinary game of dice hopes to throw a six and does so, we do not say he threw six intentionally" (Butler, 1978). But those are poor odds. Suppose Brown instead hopes to throw an even number, or any number other than five. These goals bring higher odds of success, and so perhaps success with these goals would be deemed intentional. Or perhaps not. Purely probabilistic actions like rolling dice might be viewed as nonintentional even when the odds of success are high—regardless of the odds, the outcome does not depend on Brown's control or skill, and he does nothing to make success more likely.

In this paper, we ask how probability affects judgments of whether outcomes are produced intentionally. These judgments are central in people's reasoning about the social world. For example, whether someone produced an action intentionally can be the basis for deciding whether that person is morally responsible for it (e.g., Malle, 2021; Anderson et al., 2020). Much research has sought to uncover how people judge whether outcomes are intentional (e.g., Malle & Knobe, 1997; Quillien & German, 2021) and other work has highlighted the importance of probabilistic inference for many areas of cognition (e.g., Denison et al., 2013; Gershman et al., 2017; Xu & Tenenbaum, 2007). Nonetheless, no prior work has specifically examined whether, or how, probability affects attributions of intentionality. Investigating this may sharpen understanding of what people are doing when they judge whether outcomes are intentional.

### Attributing Intentionality

On first thought, the intentionality of an outcome might seem to depend only on the mental states of the person who produced it. Indeed, when people are asked to describe what it means to produce an action intentionally, they mostly mention mental states, including desires, beliefs, and awareness (Malle & Knobe, 1997). Their responses suggest, for example, that an outcome is more likely to be viewed as intentional if the agent desired it and believed their action would lead to the outcome. This cannot be the whole story, though, because desire and belief can be present when a person wins a lottery through sheer luck (Malle & Knobe, 1997; Knobe 2003b) and people do not think lotteries are won intentionally (Knobe, 2004; Nadelhoffer, 2005).

To address this issue, researchers identified further factors that contribute to attributions of intentionality. Besides depending on mental states, people also consider the agent's skill and control over the outcome. In general, outcomes are more likely to be viewed as intentional when they depend on skill, rather than on luck alone (e.g., Guglielmo & Malle, 2010; Malle & Knobe, 1997; Sousa & Holbrook, 2010). For instance, people see successfully shooting a bull's-eye as intentional if the shooter was experienced with rifles and aimed perfectly, but as unintentional if the shooter was an amateur whose hand slipped while firing (Knobe, 2003b).

Beyond these factors, people also consider morality and norms. When an agent produces an outcome that is especially bad or which violates norms, people often judge it was brought about intentionally even if the agent did not have one of the mental states usually linked with intentionality, or if the agent lacked skill and control (e.g., Knobe, 2003a, 2003b; Nadelhoffer, 2005; Quillien & German, 2021; Uttich & Lombrozo, 2010). For example, people judge that the chairman of a company intentionally harmed the environment even though he did not specifically want to harm it and was instead primarily motivated to increase profits (e.g., Knobe, 2003a). People likewise judge that a disgruntled employee intentionally caused a power plant to explode, even though he accomplished this by guessing at a computer code (Nadelhoffer, 2005; Quillien & German, 2021).

### Does Probability Matter?

We wondered if people might also consider probability in judging whether outcomes are intentional. Whereas most previously studied determinants of intentionality which center on agents (i.e., their mental states, skill, and control), the probability of an outcomes often depends on external facts—for instance, the odds of rolling a six when using a fair die. But probability might matter for intentionality because it affects judgments about many mental states, and social cognition more generally. For instance, adults and even young children consider probability when inferring preferences, happiness, and surprise (e.g., Buechel et al., 2014; Doan et al., 2018, 2020; Kushnir et al., 2010; Teigen & Keren, 2002) Although intentionality is not itself a mental state, its dependence on mental states suggests (impressionistically at least) it might likewise depend on probability.

The possibility that probability matters for intentionality may also fit with findings from studies investigating the effects of skill and control. In these studies, participants read scenarios where agents succeeded against unfavorable odds. These agents got lucky when rolling dice, tossing coins, guessing ten-digit numbers, and buying lottery tickets (e.g., Knobe, 2004; Mele & Cushman, 2004; Nadelhoffer, 2004, 2005, 2006). These agents also managed to hit targets when firing guns, playing pool, and throwing darts, despite potential impediments to success, such as slipping up and lacking skill (e.g., Guglielmo & Malle, 2010; Knobe, 2003b; Mele & Cushman, 2007; Nadelhoffer, 2004; Sousa & Holbrook, 2010). While the findings suggested that outcomes produced without skill and control are viewed as non-intentional, these results might also depend on probability, since the absence of skill and control usually mean low odds of success.[1]

How might people consider probability when judging whether an outcome was intentional? A straightforward possibility is that people simply view probable outcomes as more intentional than improbable ones. For instance, people might see Brown's success as intentional when he succeeds in rolling any number other than five, but not when he succeeds in rolling a six. One idea pointing towards this account is that almost no actions are guaranteed to bring desired outcomes. So variations in intentionality could come down to differences in the certainty actions provide. However, this account may raise doubts. When Brown succeeds in rolling a number other than five, his success has nothing to do with his actions—instead it just comes down to the distribution of potential outcomes (i.e., there were five ways Brown could succeed, and only one way he could fail). But people are sometimes reluctant to base social judgments on distributional information alone (e.g., Arkes et al., 2012; Niedermeier et al., 1999; Wells 1992). For instance, they deny the owner of a lottery ticket can know it will lose (i.e., before checking the results) even though the distribution of tickets overwhelmingly consists of losers (e.g., Ebert et al., 2018; Friedman & Turri, 2015; Turri & Friedman, 2014; Turri et al., 2017).

Another possibility is that attributions of intentionality depend on probability raising: People may be more likely to see the outcome of an agent's action as intentional if the agent does something to increase the odds of that outcome. To illustrate, when Brown rolls the die, this does not raise the odds of any particular outcome. For example, although he has good odds of rolling a number other than five, these odds are simply those expected by chance, and so his success might not be viewed as intentional. On this account, people will be more likely to see success as intentional if Brown does something to specifically increase the odds of his desired outcome.

[1] Although no studies investigated the effects of probability, one came close (Nadelhoffer, 2006). Participants were more likely to judge that a deer hunter shot a bystander intentionally when the hunter knew this would be the outcome of firing at a deer, compared with when he thought it was just a possibility. But rather than contrasting actual probabilities, this study contrasted levels of the hunter's certainty (i.e., subjective probability).

This may sound as though Brown needs to have skill in rolling, or control over the die. But as will see, increasing the odds of an outcome may only require these minimally.

This account was loosely inspired by probability raising accounts of how people attribute causality and responsibility when multiple factors contribute to an outcome. According to these accounts, factors are seen as more causal and responsible for an outcome to the extent that they increase its probability (e.g., Brewer, 1977; Spellman, 1997). Such accounts of causality (and others too) often struggle when multiple events or agents can each guarantee an outcome (Mandel, 2003; see also Gerstenberg & Lagnado, 2014 for a helpful discussion). For example, suppose Brown rouses his wife from sleep by saying "wake up" just seconds before her loud alarm clock goes off. Here Brown caused his wife to wake up on time, but arguably without raising the probability of this outcome – she was already guaranteed to wake on time because she'd set her alarm the night before. Such examples suggest that causality can be attributed in the absence of probability raising.

This challenge for probability raising accounts of causality may not arise for attributions of intentionality. Asking whether an agent produced an outcome intentionally makes sense if the agent had a causal role in producing it—if Brown had no causal role in waking his wife, this would rule him out as having woken her intentionally or unintentionally. But once it is decided that an agent did have a causal role, the assessment of probability raising doesn't need to look at other factors that contributed to the outcome, or that could have contributed. Instead, it involves checking whether the agent's actions raised the probability of the desired outcome in comparison with other actions the agent could have taken. As such, the account suggests that assessments of probability raising depend on counterfactual reasoning *about that agent*. So, when determining whether Brown woke his wife intentionally, people will not consider whether the alarm clock would have woken her up if he had not. Instead, they might conclude that by saying "wake up" (instead of remaining silent), Brown made it likely that she would wake up—that is, he increased the probability of her waking.

We explored how probability affects attributions of intentionality across five experiments. Because attributions of intentionality appear to be hair-triggered in morally charged situations, we limit our investigation to amoral situations where the outcome is not especially vile or virtuous. Specifically, we had participants consider scenarios where a girl uses a gumball machine in hopes of getting a yummy gumball. The machine also contained yucky gumballs, and so the girl was not guaranteed to succeed. Our goal was to determine whether participants might nonetheless sometimes view her success as intentional.

**General Methods**

Preregistrations, stimuli, code, and data for all experiments are available at https://osf.io/kux7f/. The Supplementary Materials also report two additional experiments conducted before those reported here. We have disclosed all measures, manipulations, and exclusions.

Participants were recruited and tested using CloudResearch (Litman et al., 2017). They were all residents of the US and had HIT approval rates greater than 95%. We tested different participants in each experiment and aimed to recruit enough participants for at least 100 per between-subjects condition after exclusions. We anticipated these sample sizes would suffice given findings from other papers with similar designs.

After completing the main task in each experiment, participants were asked comprehension questions to check they had paid attention. Participants who failed one or more

comprehension questions were excluded, as were participants who neglected to answer any test questions. Participants were also asked about their age and gender.

## Experiment 1

### Methods

**Participants.** We tested 355 participants (mean age = 40; 169 females, 185 males, and 1 indicating "other or prefer not to disclose"). An additional 46 participants were excluded for failing comprehension questions.

**Procedure.** Participants read a story about a girl who uses a gumball machine, hoping to get yummy pink gumballs rather than yucky gray ones; see Figure 1. The experiment used a 2X2X2 design. The girl used the gumball machines across four days. Each day corresponded with one cell yielded by crossing two within-subjects manipulations: whether the machine mostly contained yummy gumballs or yucky gumballs[2]; and whether the girl received a yummy or yucky gumball. Participants were asked about each use of the gumball machine, and the test question was manipulated between-subjects. Participants were either asked whether the outcome was intentional ("Did the girl get a [yucky/yummy] gumball intentionally?") or about whether it was intended ("Did the girl intend to get a [yucky/yummy] gumball?"). Participants responded to each test question using a 9-point scale ranging from "Definitely no" to "Definitely yes".

We were mainly interested in judgments about whether the girl received yummy gumballs intentionally. We included trials where she received a yucky gumball only because it would have been odd if she always obtained a yummy one. And we asked about what the girl had intended as a basis for comparison. We expected participants to agree the girl intended to get a yummy gumball regardless of whether her initial odds were high or low—after all, it's what she wanted to get. So, we thought these judgments could provide an upper benchmark for judgments about whether she got a yummy gumball intentionally.

### Results and Discussion

We analyzed ratings in a 2(judgment: intentional, intended) X 2(machine: mostly yummy, mostly yucky) X 2(received: yummy, yucky) generalized estimating equations model for linear data. All three main effects were significant: ratings were overall higher for judgments that the outcome was intended than intentional, $F(1) = 80.56$, $p < .001$, when the girl received a yummy gumball rather than a yucky one, $F(1) = 709.65$, $p < .001$, and when the machine mostly contained yummy gumballs rather than yucky ones, $F(1) = 30.55$, $p < .001$. However, these main effects were qualified by two 2-way interactions, received x judgment, $F(1) = 191.42$, $p < .001$, and machine x received, $F(1) = 81.98$, $p < .001$, and by the 3-way interaction $F(1) = 7.22$, $p = 0.007$. Only the 2-way interaction between machine and judgment was non-significant, $F(1) = 2.14$, $p = 0.144$.

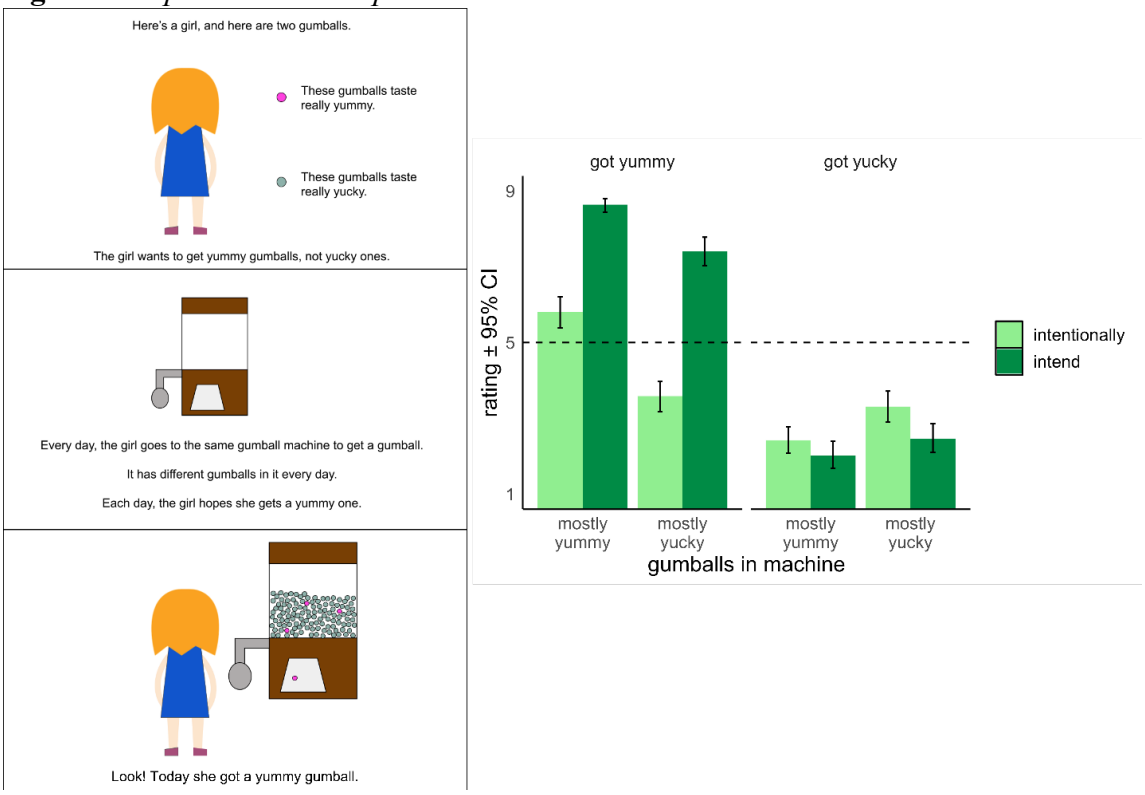
Given these interactions, we ran separate judgment x machine analyses for each outcome. When the girl received a yummy gumball, there were main effects of judgment, $F(1) = 240.25$, $p < .001$, and machine, $F(1) = 135.75$, $p < .001$, and a significant interaction between them, $F(1) = 11.17$, $p = .001$. Participants gave higher ratings when the machine contained mostly yummy gumballs rather than mostly yucky ones, but this difference was larger for judgments that the girl got a yummy gumball intentionally, $M_{difference} = 2.22$, $SE = 0.23$, $p < .001$, than judgments that she intended to get it, $M_{difference} = 1.23$, $SE = 0.18$, $p < .001$.

When the girl received a yucky gumball, there were main effects of judgment, $F(1) = 9.87$, $p = .002$, and machine, $F(1) = 14.39$, $p < .001$, but the interaction was not significant, $F(1) = 1.49$, $p = .222$. Opposite to the main effects from the omnibus analysis, ratings were higher

[2] In the Supplemental Materials, we detail the distribution of yummy and yucky gumballs for all stimuli.

when participants judged whether she got a yucky gumball intentionally than when judging if she intended to get it; and ratings were higher when the machine contained mostly yucky gumballs. Despite these significant effects, these judgments had very low ratings (see Figure 1).

**Figure 1.** *Experiment 1. Sample Scenario and Results*




*Note.* The left panel shows slides where the girl used a mostly yucky machine but nonetheless gets a yummy gumball. The right panel shows mean ratings.

In sum, participants generally agreed the girl got a yummy gumball intentionally when getting one was likely, but not when this was unlikely. The findings generally fit with the view that people see probable outcomes as more intentional than improbable ones. After all, ratings were clearly higher when the girl's success came from a machine offering good odds rather than bad ones. Nonetheless, there is reason for pause before settling on this account. Agreement that the girl got the yummy gumball intentionally was never very strong — even when she used the machine offering good odds, mean ratings barely exceeded the scale midpoint. By way of comparison, ratings were near ceiling when participants judged whether she had intended to get a yummy gumball.[3] So, although people may see probable outcomes as more intentional than improbable ones, something more is needed for strong agreement that an outcome was intentional.

We wondered if this missing ingredient might be probability raising. In this experiment, the girl had no way of increasing her probability of success. She was faced with one gumball machine at a time and could do nothing to affect how they operated. But if the girl was instead faced with a choice between machines, she would be able to raise the likelihood of success—that

[3] We also found similar results in Experiments S1 and S2, which are reported in the Supplemental Materials.

is, by choosing the machine offering better odds. We explored this idea in the next experiments. These experiments all used between-subjects designs. Rather than making judgments across multiple uses of gumball machines, each participant made a single judgment.

## Experiment 2

### Methods

**Participants.** We tested 301 participants (mean age = 40; 130 females, 168 males, and 3 indicating "other" or preferring not to disclose). An additional 18 participants were excluded for failing comprehension questions.[4]

**Procedure.** Participants again read a story about a girl who wants to get yummy pink gumballs rather than yucky gray ones; see Figure 2. The girl then chose between two gumball machines. The machine on the right mostly contained yucky gumballs, while the machine on the left varied across 3 between-subjects conditions. It either contained yummy gumballs only, mostly yummy gumballs, or mostly yucky ones (though with more yummy ones than the machine on the right). In each condition, the girl chose the machine on the left, and received a yummy gumball. Participants were then asked, "Did the girl get a yummy gumball intentionally?" (9-point scale).

Hence, by choosing the machine on the left, the girl always increased her odds of obtaining a yummy gumball. The conditions differed in how much her odds improved, and the "all yummy" condition was included to get a sense of the maximal score participants might provide.

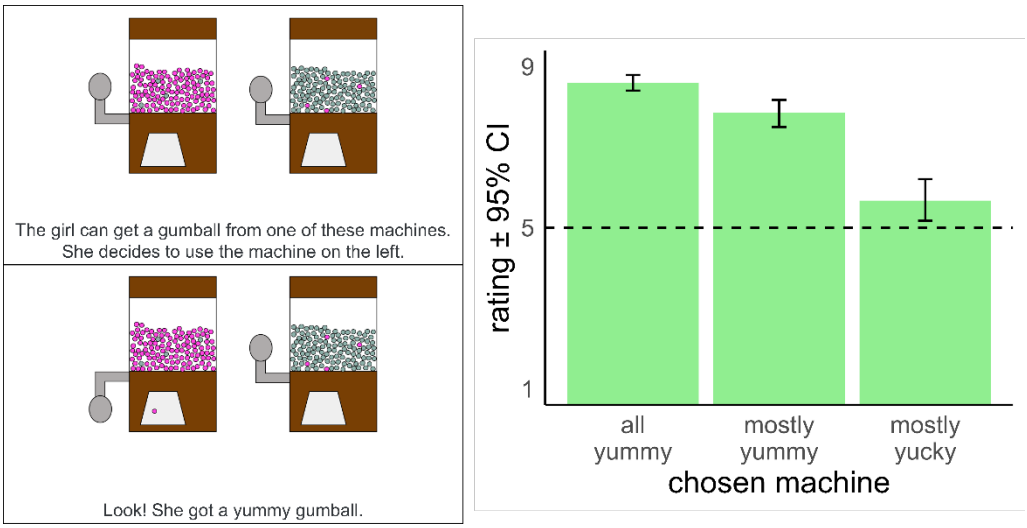
### Results and Discussion

Ratings of whether the girl got a yummy gumball intentionally depended on which gumball machine she chose (see Figure 2), $F(2.00, 176.52) = 50.92$, $p < .001$, $\eta^2 = 0.37$. Participants gave higher ratings when the chosen machine only contained yummy gumballs than when it mostly contained yummy gumballs, $p = .001$, though for both machines, participants strongly affirmed she got a yummy gumball intentionally. By contrast, participants denied the girl got a yummy gumball intentionally when she used a machine that mostly contained yucky gumballs, and ratings in this condition were significantly lower than in the other two, both $p$s < .001.

These findings again show that people view a desired outcome as more intentional if its odds of success were good rather than bad. Also, now people gave much more positive ratings—for example, when the girl ended up using a machine with mostly yummy gumballs, mean ratings were well above the scale midpoint. In fact, they were even above the scale midpoint when she used a machine offering poor odds. One key difference is that in contrast with the previous experiment, now the girl had a choice between gumball machines, and in all conditions her choice increased her odds of success. So, the increased ratings are consistent with the probability raising account. This said, the evidence is indirect. Clearer evidence would come from directly comparing performance from conditions with and without probability raising. We did this in the next experiment.

**Figure 2.** *Experiment 2. Sample Scenario (left) and Results (right)*

[4] One comprehension question in this experiment was poorly written and had two plausible answers in the condition where one machine had yummy gumballs only. Hence, we treated both answers as correct to avoid needlessly excluding participants.




*Note.* The left panel shows slides from the condition where the chosen machine mostly contained yummy gumballs. These slides were preceded by a slide introducing the girl and her preferences. The right panel shows mean ratings.

## Experiment 3

### Methods

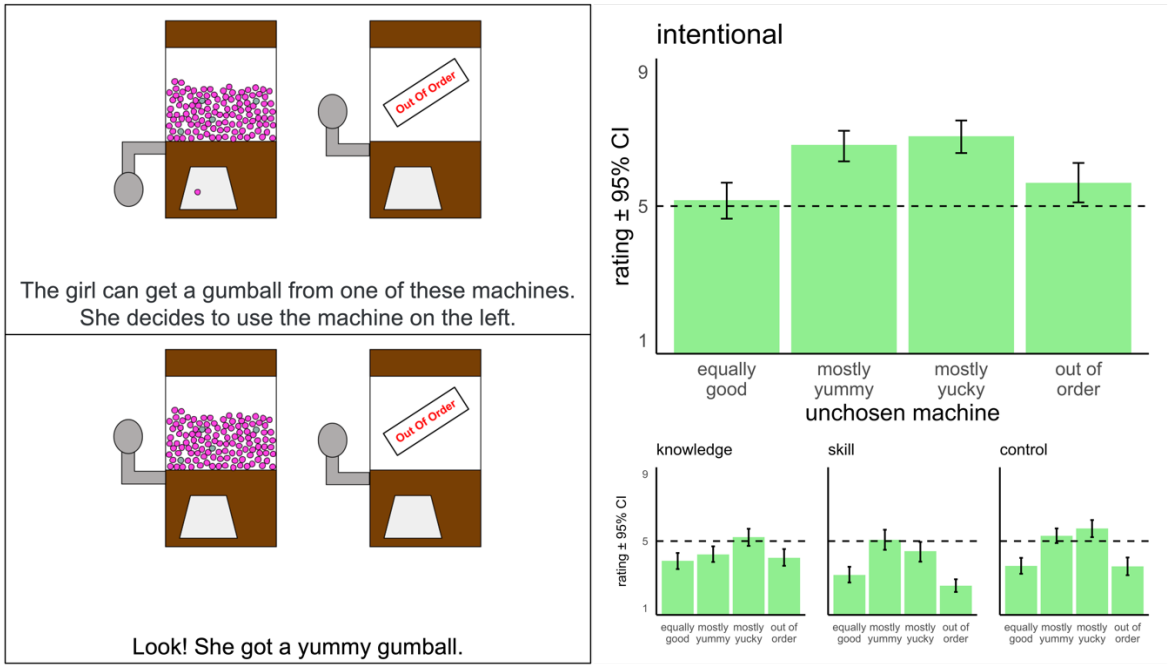
**Participants.** We tested 387 participants (mean age = 40; 186 females, 200 males, and 1 indicating "other or prefer not to disclose"). An additional 44 participants were excluded for failing comprehension questions.

**Procedure.** Participants were first introduced to the girl who wants yummy gumballs and not yucky ones. As in the previous experiment, the girl could choose between two gumball machines; see Figure 3. This time, the machine on the left mostly contained yummy gumballs, while the machine on the right varied across 4 between-subjects conditions. It either contained: the same distribution of gumballs as the machine on the left; a slightly worse distribution, but still with mostly yummy gumballs; mostly yucky gumballs; or it contained no gumballs and was out of order. The girl chose the machine on the left, and received a yummy gumball. Participants were then asked, "Did the girl get a yummy gumball intentionally?" (9-point scale).

After this, participants were asked three further exploratory questions (order randomized) about whether the girl knew she would get a yummy gumball, whether her skill helped her get a yummy gumball, and whether she had control over getting one. As with the main test question, participants responded to these questions on a 9-point scale from "Definitely No" to "Definitely Yes".

With this design, choosing the machine on the left raised the girl's probability of obtaining a yummy gumball in two conditions—when the unchosen machine had mostly yummy gumballs (but fewer than the chosen machine), and when it held mostly yucky gumballs. But this choice did not raise the odds when the unchosen machine had the same distribution as the chosen one, and when it was broken.

**Figure 3.** *Experiment 3. Sample Scenario and Results*




*Note.* The left panel shows slides from the condition where the nonchosen machine was broken. These slides were preceded by one introducing the girl and her preferences. The right panel shows mean ratings.

**Results and Discussion**

Ratings of whether the girl got a yummy gumball intentionally depended on the gumball machine she did *not* choose (see Figure 3), $F(3.00, 210.73) = 11.72$, $p < .001$, $\eta^2 = 0.14$. Participants affirmed the girl got a yummy gumball intentionally when the unchosen machine had a worse distribution of gumballs than the machine she chose—both when the rejected machine had mostly yummy gumballs (but fewer than the chosen machine) and when it mostly had yucky gumballs. Ratings did not differ between these two conditions, $p > .999$. But participants were less likely to affirm the girl got a yummy gumball intentionally when the unchosen machine had the same distribution as the chosen one, and when it was out of order. Ratings from these conditions also did not differ from one another, $p > .999$, but were lower than when the unchosen machine had mostly yummy gumballs, $p < .001$ and $p = .021$, and lower than when it had mostly yucky gumballs, $p < .001$ and $p = .003$.

**Follow-up questions.** Participants also rated whether the girl knew she would get a yummy gumball and whether she had skill and control over this outcome. As shown in Figure 3 and detailed in the Supplemental Materials, responses differed across the four conditions for all three questions. Responses for the skill and control questions received higher ratings in the two conditions with probability raising than in the two conditions without this, and the effect for control was similar in magnitude to the effect for intentionality. However, in contrast with intentionality, skill and control were never rated much higher than the scale midpoint. We discuss these ratings further in the General Discussion, including limits stemming from how the questions were worded.

These findings support the probability raising account. The girl's success in getting a yummy gumball was viewed as more intentional in the two conditions where her action increased its likelihood. We can see that the raw likelihood of success did not contribute much

because it was the same in all conditions—in each, the girl used the same "mostly yummy" machine. However, we cannot be sure that participants' judgments depended on probabilities and odds per se. Instead, participants could have considered the girl's mental states. The girl knew the odds offered by each machine, so perhaps probability mattered only insofar as it affected her beliefs and expectations — maybe it only mattered because it increased her *certainty* about getting a yummy gumball.

Some findings contradict the idea that strong belief of success (or increase in this belief) is necessary for actions to be viewed as intentional. For example, participants in one study affirmed that a skilled bowler intentionally knocked down bowling pins, even though he doubted he would succeed—he wrongly believed the pins weighed 200 pounds each and would be difficult to knock down (Mele & Cushman, 2007). As reviewed, participants also attributed intentionality to an employee who caused a power plant to explode by guessing at a 10-digit code (Nadelhoffer, 2005; also see Quillien & German, 2021), even though it's hard to imagine the employee believed he would succeed. But this earlier work asked about vignettes where the agent either had skill or produced morally charged outcomes, and beliefs about success could matter more outside these contexts. Hence, to address the concern that the findings hinged on the girl's beliefs and expectations, we next asked about a situation where she was unaware of the success and had no basis for believing she would succeed.

## Experiment 4

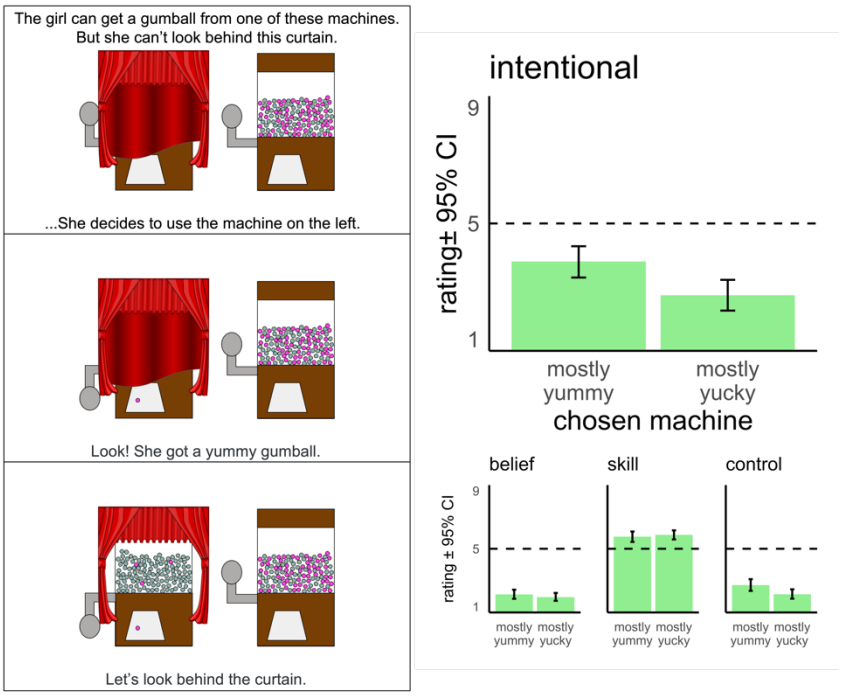
### Methods

**Participants.** We tested 195 participants (mean age = 41; 97 females, 95 males, and 3 indicating "other or prefer not to disclose"). An additional 20 participants were excluded for failing comprehension questions.

**Procedure.** Participants were introduced to the girl and her preferences for yummy over yucky gumballs, and then saw two gumball machines she could choose between; see Figure 4. The machine on the left was covered by a curtain so its gumballs could not be seen; the machine on the right had a 50/50 split of yummy and yucky gumballs. The girl then chose to get a gumball from the occluded machine and received a yummy one. The curtain was then removed to reveal the gumballs in the machine. The contents varied across between-subjects conditions: The machine mostly contained either yummy gumballs or yucky ones. Participants judged whether the girl got a yummy gumball intentionally (9 point scale) and were then asked 3 exploratory questions (order randomized, 9 point scales) about whether she believed she would get a yummy gumball, whether her skill helped her get a yummy one, and whether she had control over getting a yummy one.

With this design, the girl has no basis for expecting success. So, if probability matters only insofar as it affects beliefs of success, judgments should be identical regardless of whether the girl's choice increased or decreased the odds of winning. However, if attributions of intentionality also depend on the actual odds, ratings should be higher when the chosen machine mostly contains yummy gumballs rather than yucky ones.

**Figure 4.** *Experiment 4. Sample Scenario and Results*




*Note.* The left panel shows slides from the condition where the chosen machine mostly contains yucky gumballs. These slides were preceded by one introducing the girl and her preferences. The right panel shows mean ratings.

**Results and Discussion**

Participants gave higher ratings when the chosen machine mostly had yummy gumballs than when it mostly had yucky ones, Welch t-test, $t(192.79) = 3.02$, $p = .002$, Cohen's d = 0.43; see Figure 4. These findings provide some support for the idea that participants are sensitive to the actual odds offered by each location (i.e., rather than merely being sensitive to the girl's beliefs). However, this evidence is weak and unsatisfying because ratings were low and participants mostly denied the girl got the gumball intentionally.

Participants were also asked follow-up questions about the girl's belief, skill, and control. As shown in Figure 4 and detailed in the Supplemental materials, ratings of belief and skill did not vary across the conditions. Ratings of control were higher in the mostly yummy than mostly yucky condition, but near floor in both conditions. As in the previous experiment, though, the effects for control and intentionality were similar in magnitude.

Given these findings, we used a different approach to test whether probability itself matters, or whether it instead matters only insofar as it affects expectations. In our final experiment, the girl always had grounds to believe that choosing one location would increase her chances of getting yummy gumballs. However, we manipulated whether this belief was true or false.

## Experiment 5

**Methods**

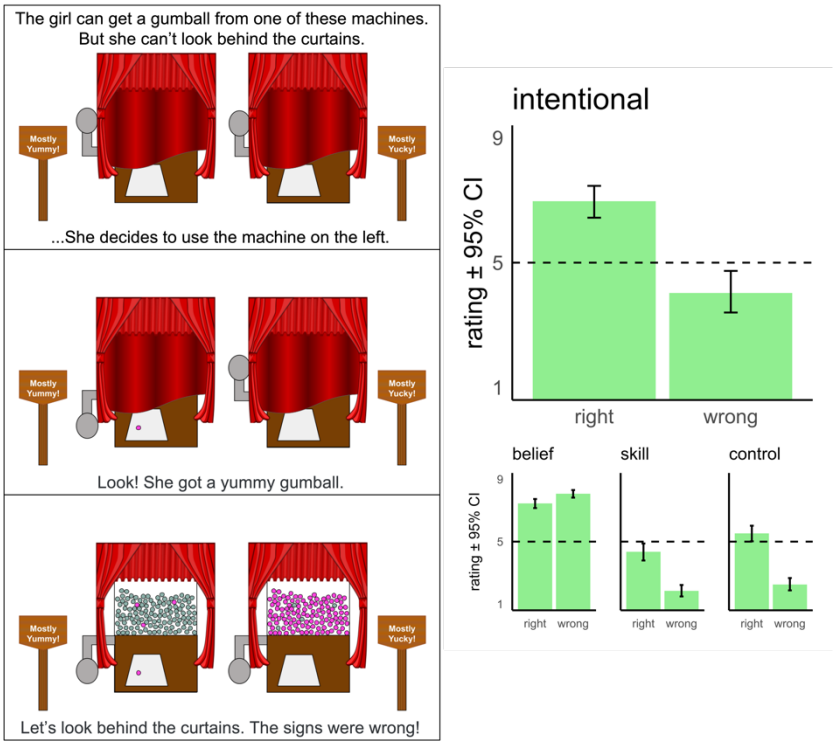
**Participants.** We tested 172 participants (mean age = 41; 91 females, 80 males, and 1 indicating "other or prefer not to disclose"). An additional 43 participants were excluded for failing comprehension questions.

**Procedure.** Participants were introduced to the girl and her preferences for yummy over yucky gumballs, and then saw two gumball machines she could choose between; see Figure 5.

This time, both machines were covered, but each had a sign advertising its contents: one said "Mostly Yummy!" and the other said "Mostly Yucky!". The girl used the machine advertised as having mostly yummy gumballs and received a yummy gumball. The curtains were then lifted to reveal the contents of both machines. In one between-subjects condition, the signs were accurate about the contents of the machines. But in the other condition, the signs were wrong. The machine advertised as mostly having yummy gumballs mostly contained yucky ones, and vice versa. Participants were then asked whether the girl got a yummy gumball intentionally (9 point scale). They were then asked 3 further exploratory questions (order randomized, 9 point scales) about whether she believed she would get a yummy gumball, whether her skill helped her get a yummy one, and whether she had control over getting a yummy one.

In this experiment, the girl always had grounds for believing she would get a yummy gumball. Hence, on an expectations-only account, participants should judge she got a yummy gumball intentionally regardless of the actual odds offered by the machine she chose. But if ratings are higher when her choice did increase the odds, this will show that probability itself matters.

**Figure 5.** *Experiment 5. Sample Scenario and Results*




*Note.* The left panel shows slides from the condition where the signs are wrong. These slides were preceded by a slide introducing the girl and her preferences. The right panel shows mean ratings.

**Results and Discussion**

Ratings were higher when the signs were correct rather than incorrect, Welch t-test, $t(145.36) = 6.68$, $p < .001$, Cohen's $d = 1.05$; see Figure 5. These findings clearly show that attributions of intentionality do not just depend on expectations. The girl had reason to expect

she would get a yummy gumball in both conditions. But participants only agreed she got one intentionally when the machine she chose provided good odds for obtaining a yummy gumball.

In responding to the follow-up questions, participants gave higher ratings for skill and control when the signs were correct than when they were incorrect, and these effects were similar in magnitude to the effect for intentionality (see Figure 5 and the Supplemental Materials). In contrast with intentionality though, ratings for control and skill were not substantially above the scale midpoint. We revisit these findings in the General Discussion along with responses to the follow-up question from Experiments 3 and 4.

## General Discussion

We investigated how probability affects people's judgments about whether an agent produces outcomes intentionally. To this end, we had participants consider scenarios where a girl received a yummy gumball from a gumball machine, rather than a yucky one. Participants saw the girl's success as more intentional when her initial odds of success were high rather than low. However, participants only strongly endorsed her success as intentional when she chose between gumball machines in ways that made success more likely—for example, by choosing to use a machine mostly containing yummy gumballs, instead of one mostly containing yucky ones.

We also investigated whether participants' judgments of intentionality depended on the agent's beliefs and expectations about the likelihood of success. These mental states did matter. Participants denied the girl's success was intentional when she had no basis for believing she would succeed and was unaware of the odds offered by the gumball machine she chose. Nonetheless, the effects of choice do not reduce to this sensitivity to mental states. When the girl was unaware of the odds provided by the machine she chose, participants still gave higher ratings when it provided good odds rather than bad ones. Also, the final experiment found that when the girl *believed* she had chosen a gumball machine offering good odds, judgments of whether success was intentional hinged on whether the machine really did provide good odds.

Together, these findings support a probability raising account of intentional action. They suggest that in judging whether an agent's success was intentional, people consider whether the agent did something to specifically increase its odds—even if just choosing between probabilistic options. Other work has found that adults and children also consider probability when attributing mental states, such as surprise and happiness (e.g., Kushnir et al., 2010; Ong et al., 2015; Teigen & Keren, 2002). Most of this work did not examine probability raising itself. However, the findings from one study do suggest that young children consider it when attributing preferences (Doan et al., 2022), and other work suggests it matters in people's attributions of causation and responsibility (Brewer, 1977; Spellman, 1997).

Besides assessing intentionality, participants in the three last experiments also assessed the agent's skill and control over the outcomes. These ratings were generally low and did not substantially exceed the scale midpoint, even when participants strongly agreed that actions were intentional. At the same time, control and intentionality appeared to be related—they were similarly affected by manipulations of probability raising. That is, when we compared how intentionality and control were affected by condition the analyses never yielded significant interactions with condition (see Supplemental Materials).

What do participants' ratings of control and skill tell us? One implication is that probability raising may provide a context where people attribute intentionality to agents who do

not have much skill or control.[5] This separation of intentionality from skill and control is precedented by earlier work showing that people see negative outcomes as intentional even when these depend on sheer luck (e.g., Nadelhoffer, 2005; Quillien & German, 2021; but see Guglielmo & Malle, 2010 for an alternative perspective on these findings).

Even so, the correspondence between control and intentionality could be viewed as consistent with a control-based account—people might see an agent whose actions raise the probability of an outcome as having *some* degree of control (i.e., even if minimal), which may in turn support seeing the agent's success as intentional (probability raising → control → intentionality). Alternatively, control could be epiphenomenal. People might see probability raising as supporting inferences of both control and intentionality (control ← probability raising → intentionality) without linking control with intentionality. This latter view raises an ambiguity in previous work suggesting that control matters for intentionality. Specifically, effects of control, and perhaps skill too, might reduce to probability raising—control and skill might primarily matter for intentionality insofar as they are viewed as impacting the probability of success.

Similar questions can be asked about other factors that matter for attributions of intentionality. Consider causation: Quillien and German (2021) recently provided evidence that people attribute intentionality by assessing whether an outcome causally depends on the agent having a positive attitude towards it. Applied to our experiments, considerations of probability raising might have been one link in a process wherein participants judged whether the girl got a yummy gumball intentionally by asking if her desire for a yummy gumball caused her to get one (attitude → probability raising → causal role → intentionality). Future work could examine how probability raising relates to causal attributions. Following Quillien and German (2021), participants could be asked to rate whether the outcome was caused by the agent's desires. Or studies could explore how they attribute intentionality in situations where probability raising and causality come apart (e.g., situations where an agent increases the odds of success, but without this increase causing the success).

Related to this, future work could also look at when beliefs matter for intentionality, and the even thornier question of how people integrate beliefs with probability raising. In contrast with the idea that intentionality requires an agent to believe their action will produce the desired outcome (e.g., Malle & Knobe, 1997), Quillien and German (2021) suggested – and showed – that even weak beliefs will sometimes do. In many cases, the agent needs only to believe the action had *some* (i.e., non-zero) possibility of causing it. At the same time, our findings indicate that strong belief is sometimes necessary, as participants in our fourth experiments denied intentionality when the girl was unaware of the odds of the chosen machine, even though she had some basis for believing it might provide a yummy gumball.[6] One possibility is that stronger belief are necessary for actions and outcomes that are trivial and harmless (e.g., using a gumball

[5] This said, it's possible that control and skill ratings would be stronger with other question wordings. For instance, it might have been clearer to probe judgments about skill by referencing the girl's *ability* to get a yummy gumball rather than directly referencing skill.

[6] This is not evidence against Quillien and German's (2021) account, as their account does not deny that strong belief could sometimes be necessary, and they convincingly showed that in some contexts weak belief is enough. We are more curious about the fit between their account and the pattern in Experiment 5 where the girl read signs about the contents of each gumball machine. Participants affirmed intentionality when the signs were accurate and the girl used a mostly yummy machine, but denied it when the signs were inaccurate and she inadvertently used a mostly yucky one. We're curious whether corresponding patterns would be seen in ratings of whether the girl's desire caused her to get a yummy gumball.

machine to getting yummy gumball) in comparison with outcomes that are harmful or otherwise questionable to pursue (e.g., using a rifle to shoot someone; see Quillien & German, 2021). If so, the relation between probability raising and intentionality may be quite complicated because (as we consider next) assessments of probability raising may themselves depend on whether they involve actions that are questionable to pursue.

Another area for future research is to more carefully examine the process by which people assess probability raising. We suggested that people assess whether the agent's actions raised the probability of the outcome in comparison with other actions the agent could have taken. But more can be learned about how people select the relevant alternative actions. To illustrate why, let's return to the scenario from the first experiment where the girl is faced with just one machine at a time. The scenario indicated that the girl wanted to get yummy gumballs and it was taken as a given that she would use gumball machines every chance she had. With this background, the girl does not affect the odds of getting a yummy gumball when she uses the only machine available. Again, using the machine doesn't increase the odds of getting a yummy gumball because it's assumed she will use the machine. To affect the odds, the girl has to do something more, like choose between machines offering different odds. But suppose instead that there are compelling reasons not to use the machine—maybe getting certain gumballs will bring negative consequences for others, or maybe the machine costs a lot of money to use. With this change, it is no longer a given that the girl will decide to use the machine—although she wants yummy gumballs, this desire doesn't outweigh the negative consequences or costs that come with using the machine. Hence people might now compare the girl's actual use of the machine with a counterfactual alternative where she does not use it, and her action might be viewed as increasing the odds of all its possible outcomes.

This account of how people assess probability raising is relevant to findings showing that people often judge that negative outcomes were intentional even when the agent producing them lacked desire, or control and skill. Although these findings led some researchers to posit that people have multiple notions of intentionality, with different notions triggered across contexts (e.g., Cova et al., 2012; Nichols & Ulatowski, 2007; also see Laurent et al., 2019, 2021), other researchers maintain that judgments of intentionality can be explained by people possessing a single understanding of intentionality. For example, judgments that an agent produced an action intentionally may depend on assessments of whether the agent's attitude towards the action caused it (Quillien & German, 2021) or whether the agent's actions (interpreted in relation to expected actions) suggest they were willing or reluctant to produce the outcome (Knobe, 2022). The account outlined here relates to these "single understanding" accounts, and to Machery's (2008) proposal that attributions of intentional action increase when actions have tradeoffs (i.e., when there are reasons why a person might not take the action). The main addition is suggesting that negative outcomes (or other compelling reasons to not act) may affect judgments by changing how people assess probability raising.

Finally, in further investigating how probability raising affects intentionality, it may be useful to explore judgments about other kinds of probabilistic systems (i.e., besides gumball machines). This will be important for assessing the generalizability of the effects of probability raising. For example, participants could be shown scenarios where an agent wants to roll a certain number on a dice, or wants to land a roulette wheel on a certain color. We focused on gumball machines because they offer so little control over the outcome—with gumballs, there is no equivalent to using more or less force when rolling a dice and spinning a roulette wheel. Regardless, it may also be worth exploring whether people reason similarly in situations that

involve probabilities, but which might not be viewed as purely probabilistic (e.g., when an agent successfully produces a chemical using a process that works some percentage of the time).

**References**

Anderson, R. A., Crockett, M. J., & Pizarro, D. A. (2020). A theory of moral praise. *Trends in Cognitive Sciences*, *24*(9), 694-703. https://doi.org/10.1016/j.tics.2020.06.008

Arkes, H. R., Shoots-Reinhard, B., & Mayes, R. S. (2012). Disjunction between probability and verdict in juror decision making. *Journal of Behavioral Decision Making*, *25*, 276–294. https://doi.org/10.1002/bdm.734

Buechel, E. C., Zhang, J., Morewedge, C. K., & Vosgerau, J. (2014). More intense experiences, less intense forecasts: Why people overweight probability specifications in affective forecasts. *Journal of Personality and Social Psychology, 106*(1), 20–36. https://doi.org/10.1037/a0034478

Brewer, M. B. (1977). An information-processing approach to attribution of responsibility. *Journal of Experimental Social Psychology*, 13, 58–69. https://doi.org/10.1016/0022-1031(77)90013-0

Butler, R. J. (1978). Report on analysis "problem" no. 16. *Analysis*, *38*, 113-114.

Cova, F., Dupoux, E., & Jacob, P. (2012). On doing things intentionally. *Mind & Language*, *27*, 378-409. https://doi.org/10.1111/j.1468-0017.2012.01449.x

Denison, S., Bonawitz, E., Gopnik, A., & Griffiths, T. L. (2013). Rational variability in children's causal inferences: The sampling hypothesis. *Cognition*, *126*(2), 285-300. https://doi.org/10.1016/j.cognition.2012.10.010

Doan, T., Friedman, O., & Denison, S. (2018). Beyond belief: The probability-based notion of surprise in children. Emotion, 18(8), 1163–1173. https://doi.org/10.1037/emo0000394

Doan, T., Friedman, O., & Denison, S. (2020). Young children use probability to infer happiness and the quality of outcomes. *Psychological Science*, 31(2), 149-159. https://doi.org/10.1177/0956797619895282

Doan, T., Stonehouse, E. E., Denison, S., & Friedman, O. (2022). The odds tell children what people favor. *Developmental Psychology*.

Ebert, P. A., Smith, M., & Durbach, I. (2018). Lottery judgments: A philosophical and experimental study. *Philosophical Psychology*, *31*, 110-138. https://doi.org/10.1080/09515089.2017.1367767

Friedman, O., & Turri, J. (2015). Is probabilistic evidence a source of knowledge? *Cognitive Science*, *39*, 1062-1080. https://doi.org/10.1111/cogs.12182

Gershman, S. J., Pouncy, H. T., & Gweon, H. (2017). Learning the structure of social influence. *Cognitive Science*, *41*, 545-575. https://doi.org/10.1111/cogs.12480

Gerstenberg, T., & Lagnado, D. A. (2014). Attributing responsibility: Actual and counterfactual worlds. In J. Knobe, T. Lombrozo, and S. Nichols (Eds.), *Oxford studies in experimental philosophy, Volume 1*. (pp. 91-130). *Oxford University Press*. https://doi.org/10.1093/acprof:oso/9780198718765.003.0005

Guglielmo, S., & Malle, B. F. (2010). Enough skill to kill: Intentionality judgments and the moral valence of action. *Cognition*, *117*, 139-150. https://doi.org/10.1016/j.cognition.2010.08.002

Knobe, J. (2003a). Intentional action and side effects in ordinary language. *Analysis*, *63*(3), 190-194.

Knobe, J. (2003b). Intentional action in folk psychology: An experimental investigation. *Philosophical Psychology*, *16*, 309-324. https://doi.org/10.1080/09515080307771

Knobe, J. (2004). Folk psychology and folk morality: Response to critics. *Journal of Theoretical and Philosophical Psychology, 24*(2), 270–279. https://doi.org/10.1037/h0091248

Knobe (2022). Morality and possibility. In M. Vargas & J. M. Doris (Eds.), *The Oxford handbook of moral psychology* (pp. 310-332). Oxford University Press.

Kushnir, T., Xu, F., & Wellman, H. M. (2010). Young children use statistical sampling to infer the preferences of other people. *Psychological Science*, *21*(8), 1134-1140. https://doi.org/10.1177/0956797610376652

Laurent, S. M., Reich, B. J., & Skorinko, J. L. M. (2019). Reconstructing the side-effect effect: A new way of understanding how moral considerations drive intentionality asymmetries. *Journal of Experimental Psychology: General, 148*(10), 1747–1766. https://doi.org/10.1037/xge0000554

Laurent, S. M., Reich, B. J., & Skorinko, J. L. (2021). Understanding side-effect intentionality asymmetries: Meaning, morality, or attitudes and defaults? *Personality and Social Psychology Bulletin*, *47*(3), 410-425. https://doi.org/10.1177/0146167220928237

Machery, E. (2008). The folk concept of intentional action: Philosophical and experimental issues. *Mind & Language*, *23*(2), 165-189. https://doi.org/10.1111/j.1468-0017.2007.00336.x

Malle, B. F. (2021). Moral judgments. *Annual Review of Psychology*, *72*(1), 293-318. https://doi.org/10.1146/annurev-psych-072220-104358

Malle, B. F., & Knobe, J. (1997). The folk concept of intentionality. *Journal of Experimental Social Psychology*, *33*, 101-121. https://doi.org/10.1006/jesp.1996.1314

Mandel, D. R. (2003). Judgment dissociation theory: An analysis of differences in causal, counterfactual and covariational reasoning. *Journal of Experimental Psychology: General*, *132*, 419–434. https://doi.org/10.1037/0096-3445.132.3.419

Mele, A. R., & Cushman, F. (2007). Intentional action, folk judgments, and stories: Sorting things out. *Midwest Studies in Philosophy*, *31*, 184-201. https://doi.org/10.1111/j.1475-4975.2007.00147.x

Nadelhoffer, T. (2004). The Butler problem revisited. *Analysis*, *64*, 277-284.

Nadelhoffer, T. (2005). Skill, luck, control, and intentional action. *Philosophical Psychology*, *18*, 341-352. https://doi.org/10.1080/09515080500177309

Nadelhoffer, T. (2006). Desire, foresight, intentions, and intentional actions: Probing folk intuitions. *Journal of Cognition and Culture*, *6*, 133-157. https://doi.org/10.1163/156853706776931259

Nichols, S., & Ulatowski, J. (2007). Intuitions and individual differences: The Knobe effect revisited. *Mind & Language*, *22*(4), 346-365. https://doi.org/10.1111/j.1468-0017.2007.00312.x

Niedermeier, K. E., Kerr, N. L., & Messé, L. A. (1999). Jurors' use of naked statistical evidence: Exploring bases and implications of the Wells effect. *Journal of Personality and Social Psychology, 76*, 533-542. https://doi.org/10.1037/0022-3514.76.4.533

Ong, D. C., Zaki, J., & Goodman, N. D. (2015). Affective cognition: Exploring lay theories of emotion. *Cognition*, *143*, 141-162. https://doi.org/10.1016/j.cognition.2015.06.010

Quillien, T., & German, T. C. (2021). A simple definition of 'intentionally'. *Cognition*, *214*, 104806. https://doi.org/10.1016/j.cognition.2021.104806

Sousa, P., & Holbrook, C. (2010). Folk concepts of intentional action in the contexts of amoral and immoral luck. *Review of Philosophy and Psychology*, *1*, 351-370. https://doi.org/10.1007/s13164-010-0028-x

Spellman, B.A. (1997). Crediting causality. *Journal of Experimental Psychology: General*, *126*, 323-348. https://doi.org/10.1037/0096-3445.126.4.323

Teigen, K. H., & Keren, G. (2002). When are successes more surprising than failures? *Cognition & Emotion*, *16*(2), 245-268. https://doi.org/10.1080/02699930143000275

Turri, J. & Friedman, O. (2014). Winners and losers in the folk epistemology of lotteries. In J. Beebe (Ed.), *Advances in experimental epistemology* (pp. 45-69). Continuum.

Turri, J., Friedman, O., & Keefner, A. (2017). Knowledge central: a central role for knowledge attributions in social evaluations. *Quarterly Journal of Experimental Psychology*, *70*, 504-515. https://doi.org/10.1080/17470218.2015.1136339

Wells, G. L. (1992). Naked statistical evidence of liability: Is subjective probability enough? *Journal of Personality and Social Psychology*, *62*, 739–752. https://doi.org/10.1037/0022-3514.62.5.739

Uttich, K., & Lombrozo, T. (2010). Norms inform mental state ascriptions: A rational explanation for the side-effect effect. *Cognition*, *116*(1), 87-100. https://doi.org/10.1016/j.cognition.2010.04.003

Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, *114*(2), 245–272. https://doi.org/10.1037/0033-295X.114.2.245