Blame the Player and The Game

Drew Walker (dehoffma@ucsd.edu)
Department of Cognitive Science
University of California, San Diego

Edward Vul (evul@ucsd.edu)
Department of Psychology
University of California, San Diego

Abstract

How do people assign credit for others' actions? The Correspondence Bias — a classic bias in social psychology — purports that people are predisposed to attribute behaviors to dispositional, rather than situational, factors. However, recent work suggests that the pattern of data cited as evidence of a bias may be a natural consequence of attribution under uncertainty. Here we devise a novel "Bucket-Toss" task in which we can independently and parametrically manipulate and measure situation and disposition pressures to evaluate whether attribution to dispositions and situations are consistent with probabilistic inference. We find that as the strength of the situation or disposition is varied, attributions to the other (unobserved) cause follow roughly symmetric patterns of graded attribution. Together, these results confirm that social attribution appears to be largely consistent with unbiased inference under uncertainty.

Keywords: Social reasoning; probabilistic inference

According to the *Correspondence Bias (CB)*, one of the best known phenomena in psychology, humans have a bias to assign too much blame to internal qualities (e.g. attitudes, dispositions) because they have incorrect causal theories about behavior. We think situations exert less influence than dispositions, resulting in what has been called a *Fundamental Attribution Error*, in which we neglect the true influence of situations when making social judgments. In their classic study Jones and Harris (1967) asked observers to read an essay that either favored Fidel Castro, the then leader of Cuba, or an essay that took an anti-Castro stance. Although readers had been told that the essay author had been assigned their position and were just following instructions, observers still believed the author had a view of Castro that aligned with their essay position. Dozens of similar studies have replicated these general results, and have spawned countless theories attempting to explain why humans have this tendency to overestimate the influence of disposition on behavior, and underestimate the influence of situations (for review, see Gilbert & Malone, 1995; Gawronski, 2004).

According to models of human causal reasoning more generally, if we are considering the influence of one cause and learn about the presence of an additional cause the influence of the initial causes should be *discounted* – the presence of the second cause "casts doubt" on the first (e.g. Khemlani & Oppenheimer, 2011; Hall, Ali, Chater, & Oaksford, 2016). However, the amount of discounting that should occur will depend on the perceived influence of the other causal factor. For example, if your car won't start but you learn your gas tank is empty, there is no reason to think there is anything wrong with your engine. However, if internal qualities and situations affect behavior probabilistically rather than deterministically, then attributing behavior to disposition, even when it is somewhat constrained by situation, is not unreasonable. If an instruction to write a certain essay does not in fact guarantee that someone writes such an essay, than full discounting should not occur because a pro-Castro essay is still somewhat informative: the author was not so anti-Castro that they refused to write the essay in the face of instructions to do so.

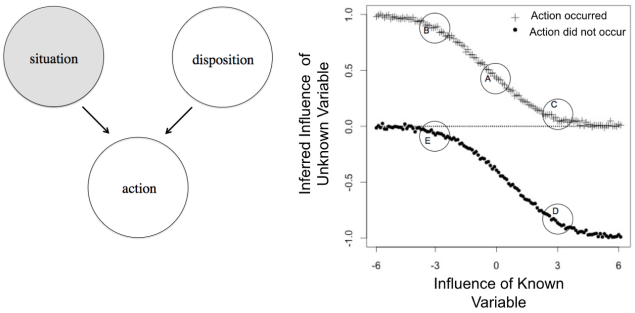

Figure 1: Probabilistic Social Attribution. (Left) Situation and disposition combine additively to determine the log-odds for an action. (Right) Conditioning on an observed action, cause is partly attributed to disposition, even if the situation favored that action. Stronger situational pressures yield more discounting of the role of disposition, but disposition is discounted to undetectable levels only under very extreme (log odds greater than 3) situational pressures.

A handful of theorists have formalized this intuition by showing that the judgments that observers make in these types of experiments are consistent with probabilistic inference under such non-deterministic causal influences. Indeed, judgments may well be optimal given the observers prior-beliefs about the strength of the situations (Trope, 1974; Ajzen & Fishbein, 1975; Morris & Larrick, 1995; Jennings, 2010; Walker, Smith, & Vul, 2015).

Probabilistic Social Attribution

The sorts of social inferences studied in the CB literature can be described as inferences in a three-node Bayes net (Walker et al., 2015) (see figure 1, left).

Situations and dispositions combine to probabilistically influence whether or not an action is taken, for example, whether a person writes a pro-Castro or anti-Castro essay. Whether an agent takes an action (a) follows a Bernoulli distribution with probability (q). It is assumed that the situational pressure (s) and dispositional pressure (d) influence the action probability additively in log odds, such that the log-odds of writing a pro-Castro essay is the sum of these two forces $\ln(q/(1-q)) = s+d$. Thus the situation and disposition pressures take on values from negative infinity (pressure toward an anti-Castro essay) to positive infinity (pressure toward and pro-Castro essay). A logistic transform on the sum of s and d yields the probability of writing a pro-Castro essay (q):

$$q = \frac{1}{1+e^{[-(s+d)]}}$$

.

A pro-Castro essay ($a = 1$) is written with probability (q), and an anti-Castro essay ($a = 0$) is written with probability (1-q):

$$P(a|q) = \begin{cases} q & \text{if } a = 1 \\ 1-q & \text{if } a = 0 \end{cases}$$

Therefore, whether a person writes a pro or anti-Castro essay will be influenced by the individual's unobservable internal quality (attitude about Castro) as well as the situation (pressure imposed to write a certain essay). A negative situation strength, say -2, amounts to a situational pressure to write an anti-Castro essay. A positive disposition strength (say +1) implies an internal predilection to write pro-Castro essays. The net pressure toward the positive action is given by the sum (-1). This net log-odds of -1 implies a probability of writing a pro-Castro essay of 27% (by converting log-odds into a probability via the logistic transform).

This model may not only be run forward — reasoning about the probability of an action given the pressure of the situation and the disposition — but may also be inverted, to estimate either the strength of an unknown situation or an unknown disposition given an observed action. For example, if we have observed an action (a pro or anti essay) and have an estimate of the strength of the situation, and have a sense of what proportion of people in the world are pro rather than anti Castro (prior distribution on disposition), we can use Bayes rule to find the posterior probability of disposition (this specific person's attitude about Castro).

$$P(d \mid a,s) = \frac{P(a|d,s)P(d)}{\sum_{d'} P(a|d',s)(P(d')}$$

A surprising result of this calculation is that the data reported by Jones and Harris (1967), and widely taken to be evidence of a bias, are actually consistent with optimal probabilistic inference given plausible situation strengths (Walker et al., 2015). For example, when there is little pressure to perform an action, such as when an author is allowed to choose their essay position and writes a pro-Castro essay, the ideal observer infers a strong influence of the internal quality, here a pro-Castro attitude (Figure 1, right panel, point A). However, when there is external motivation that is not interpreted to be deterministically strong, perhaps how observers view the influence of an essay assignment, the ideal observer infers a much weaker pro-Castro attitude, but this inference is still above zero (Figure 1, right panel, point C). These two points illustrate what in the past has been interpreted as the correspondence bias – an attitude inference despite a constrained situation – but it is a natural consequence of inference under uncertainty when the situation is not assumed to be totally compelling.

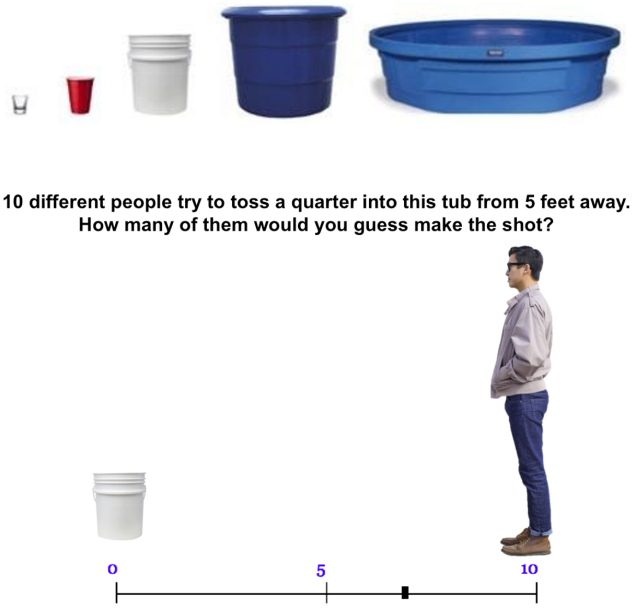

Figure 2: Estimates of situation strength. (Top) Observers estimated how many people out of ten would be able to toss a quarter into each of five containers from 5 feet away. (Bottom) Each trial showed an image of a person and one container. People moved a slider to indicate how many out of ten people would succeed in the toss.

Empirical Evaluation of Probabilistic Social Attribution

Previously, we used the empirical observations of attribution from the classic CB literature to show that these results are *consistent* with PSA (Walker et al., 2015) and therefore a systematic reasoning bias is not *required* to account for them. However, the data reported in the classic literature did not allow us to fully constrain the model – we had to guess at the correct situation strengths believed by subjects 50 years ago. Although we used strengths that seemed plausible and were endorsed by the literature (Sherman, 1980), constraining the

model with ground-truth situation perceptions is critical before advocating for the idea that humans use something like optimal probabilistic inference to reason socially. This is because the PSA model predicts the same qualitative effects as the CB, and differs only in the quantitative magnitude of these effects as situation pressure changes. If observers indeed believed the essay assignment was a totally constraining situation, an attitude inference would indeed reflect a systematic error in reasoning. Since the same results may be interpreted as a profound bias, or as the optimal conclusion, depending on the perceived strength of the situation, in order to advocate for a PSA model of social reasoning, it is critical to constrain the model with observer's actual perception of the situation.

Given how dependent attitude attribution should be on the perceived situation pressure, surprisingly little work has been done to modulate perceptions of situation strength, and characterize resulting shifts in attitude attribution. Many studies in the social psychology literature have expanded on Jones and Harris (1967) paradigm by adding manipulations of such things as perceived action strength and base-rate assumptions (Jones, Worchel, Goethais, & Grumet, 1971). But because they maintain the binary and ill-defined manipulation of lower versus higher situation constraint, the conclusions we could draw from our reinterpreted of the original literature using the PSA framework were limited.

Recent researchers taking a probabilistic approach to understand attitude attribution have attempted to measure observers' perception of these binary situations, in an effort to test internal consistency with a Bayesian probabilistic inference framework. Morris and Larrick (1995) asked subjects to report subjective probabilities of the "sufficiency" of instructions to compel a pro-Castro essay. Likewise, Jennings (2010) replicated the Jones and Harris (1967) study but using the death penalty as the controversial topic; they asked participants to rate how much overall choice they thought the participants had to write the assigned essay. Both studies found attributions consistent with Bayesian inference. That said, it is important to note the distinction between the idea of the coherence of a judgment, and correspondence or a judgment (Hammond, 1996), while coherence refers to internal consistency of an attitude given a person's own assumptions about the strength of the situation, correspondence refers to an observer's judgment aligning with ground-truth. Past probabilistic accounts have been interested entirely in the coherence of observers judgments (internal consistency). Although we do not have ground truth in the studies that follow, we do attempt to approximate ground truth by obtaining the judgment of situational pressure by a large independent group of observers, rather than using individuals own subjective probabilities as done in previous work.

In the current work we are able to show that human judgments compare to an ideal observer along a much wider range of perceived situation strengths (including situations that are in fact perceived to be near deterministic) which were quantified by asking a large sample of independent judges. We go beyond measuring subjective perceptions of binary situation strengths, using instead parametric manipulations in a novel experimental paradigm to test the key qualitative predictions of the Probabilistic Social Attribution (PSA) account. Specifically, if observers' social reasoning is consistent with PSA, we expect: (1) Attribution to internal qualities should decrease as situational pressure is perceived to increase. (2) Judgments that have classically been interpreted as an CB -– observers should infer that the unobserved feature (either internal quality or situational influence) is consistent with the observed behavior (e.g. a pro-Castro essay should indicate a pro-Castro attitude), for all but the strongest situation strengths. (3) When situation is perceived to totally compel the observed behavior, then no attributions should be made regarding disposition: this would be evidence of a genuine CB. Finally, (4) All of these results should also hold for attribution to unknown situations in the presence of observed dispositions. This reversal has been observed in the CB literature (Quattrone, 1984) and is a critical prediction of PSA. Symmetric patterns when judging situations and dispositions would demonstrating that the CB pattern is a general feature of imperfect discounting, rather than an asymmetric preference for attributing to disposition.

Bucket Toss

To more rigorously test the predictions of PSA we need a scenario in which it is easy to fairly unambiguously manipulate situation strength parametrically, and where judgments of the relationship between internal qualities and behaviors can be made quantitatively, and fairly consistently across observers. To achieve these ends, we designed a scenario that involves a made-up carnival game called "Bucket Toss." The objective of "Bucket Toss" is to throw a quarter into a container (the observed action) from a fixed distance. Although the classic CB studies we've discussed have asked observers to make attributions about attitude, judgments of skill have also been elicited to study social attribution (Ross, Amabile, & Steinmetz, 1977; Lau & Russell, 1980). In "Bucket Toss" the size of the container can change (situational influence) and —- as is usually the case in the world — different individuals are differently skilled at tossing coins. We first obtain objective measure of the perceived strength of the situation (e.g. how many people would make a quarter into a given container from 5 feet away), and then in two experiments we varied the container size and known skill, respectively, to compare human inferences to an ideal observer making inferences about either an unknown internal quality or an unknown situation with identical information.

Experiment 1: Estimates of Situation Strength

The goal of the first study was to simply determine how strong situations are perceived to be, which was operationalized as the proportion of people able to make a quarter from five feet away. This experiment was run online with 143 participants from Amazon Mechanical Turk.

Methods

For each participant, each of the five containers appeared once in a random order along with a profile image of a person standing 5 feet away from it (Figure 2). Under the PSA account, the strength of the situation is defined as the proportion of people who would be compelled to perform a certain action, so questions were phrased to elicit responses appropriate to this definition. On each trial participants were asked "if ten different people tried to toss a quarter into this container from five feet how many of these ten would you guess make the shot?" Participants used a mouse to move a toggle on a sliding scale. At the beginning of each trial the toggle started at the center (which corresponded to five people) and participants could indicate they thought more than five people would make the shot by moving the toggle to the right, or less than five people by moving the toggle to the left. The number of people the slider position mapped to appeared below the scale, and this changed dynamically as the participant moved the slider.

Results

The median responses revealed the expected pattern of situation strengths: observers believed that only .06% of people could toss the quarter into a shot glass, 23% of people could make it into the red cup, 59% of people could make the 5-gallon bucket, 79% of people could make the ice-tub, and 99% of people would make the shot into the kiddie pool. Importantly, as the bucket size increased, judgments about the proportion of people who could make the shot increased monotonically and were largely non-overlapping (Figure 3). These features are important as they validate that these container sizes effectively communicate unambiguous situational pressures across a very wide range.

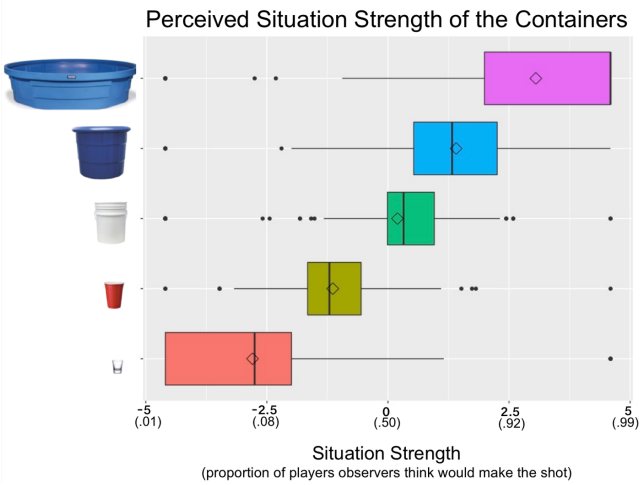

Figure 3: Estimated situation strengths differed for the five containers. We transformed estimated proportions into log-odds (x-axis) to put them on the same scale as situation strengths as the model (proportions shown in parentheses).

Experiment 2: Inferring internal qualities from actions and situations

In this experiment we parametrically varied the strength of the situation, and elicited judgments of internal quality. Observers made inferences about a player's "Bucket Toss" ability when the player made (or missed) a shot, as a function of container size. The pattern of judgments across these conditions will reveal whether inferences about disposition track the predictions of PSA.

Methods

This experiment was conducted online with 299 participants from Amazon Mechanical Turk, each making ten judgments. Participants were told the rules of "Bucket Toss": that it is a game in which players try to toss a quarter into various containers from a fixed distance. These initial instructions were presented along with an image of the five containers that would be seen during the experiment. Participants were also told that they would see a series of players and that they would need to guess how skilled each person was at the game.

On each trial, participants saw a profile image of a person standing a fixed distance away from a container (like Figure 2, bottom), and were told the person either made or missed the shot. Observers were asked to report how skilled they thought that individual was at the game, and they made this judgment by moving a slider to indicate a percentile from "very unskilled" on the left and "very skilled" on the right, with "average" in the center. As the participant moved the slider the words below the scale changed dynamically to reflect what that setting meant in terms of what percent of people were better and worse than the target player. On each of the ten trials the container presented was selected randomly with the constraint that each would appear twice per participant. Whether the player was shown to make or miss the shot was determined randomly on each trial. The name and image of the player shown was selected randomly from a bank of 12 (6 males and 6 females) with the constraint that each person/name would only appear once per participant.

Model and Results

We can now compare how humans attribute actions to traits of an individual and the attributions of an ideal observer. The ideal observer uses estimates of situation strength (container difficulty) obtained in Experiment 1, to estimate the skill level of a "Bucket-Toss" player who either made, or missed, a shot into a bucket of a particular size (Figure 4, panel A). Likewise, we obtain human observers' ratings of skill level on those same types of trials (Figure 4, panel B). Human judgments track the critical predictions of the PSA model: (1) For a particular action (made or miss) attribution to internal qualities decreases as perceived situational pressure increases. (2) The unobserved variable (skill) is estimated to be in a direction consistent with the observed behavior (above average skill is inferred when the person makes the shot, below average skill is inferred when the person misses the shot). (3)

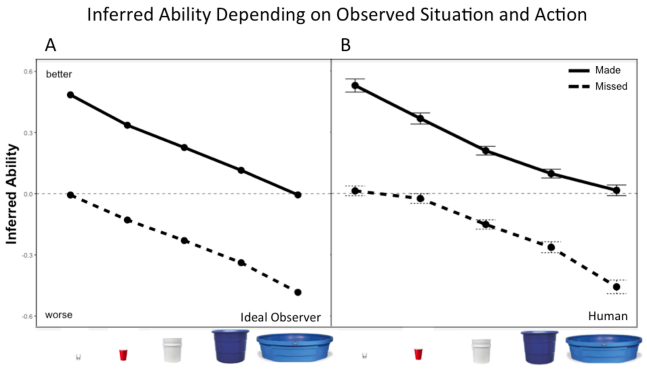

Figure 4: Inferred ability at "Bucket Toss" as a function of container size, and whether the player made or missed the shot. (A) Probabilistic Social Attribution predictions. When a shot is made (solid line) the player is deemed to have above-average ability, and this inference is more extreme for smaller containers. The converse is true when a shot is missed (dashed line). No ability is inferred under for nearly deterministic situations: missing the shot glass or making the kiddie pool. (B) Human inferences are consistent with Probabilistic Social Attribution. When action is held constant, attribution to ability decreases systematically as situational pressure is increases. Behavior-consistent inferences are made so long as the situation is not so strong as to be deterministic.

However, no attributions are made when the situation is perceived to be totally compelling – when the player makes the quarter into the kiddie pool or misses the shot glass.

These data also show the expected "Correspondence Bias" style behavior: when situations weakly motivate the behavior (missing the solo cup v. making the ice bucket), inferred ability differs between making and missing a toss ($t(603) = 10.78$, $p.001$). This inference is weaker than in the least constraining situation (five-gallon bucket; $t(597) = 11.51$, $p.001$), but it vanishes entirely in the very influential situations (shot glass vs kiddie pool; $t(580) = .075$, $p = .94$. Together, these results show that when we parametrically manipulate situation strength, and measure the disposition strengths people estimate for these situations, human behavior is qualitatively consistent with probabilistic attribution.

Experiment 3: Inference When Situational Pressure is Unknown

In the previous experiment the unknown variable was "Bucket Toss" ability, which was inferred given knowledge about the success of the shot and the container size. This is analogous to most attribution experiments——participants are asked to infer an unknown internal quality given an observed situation and action. However, the Bayesian social attribution framework says that the pattern of inferences in these studies arises due to inferring unknown variables under uncertainty; patterns of inference should be similar regardless of whether the unknown feature is the situation or the disposition. So, in this experiment we were interested in the inferences people would make about the situation (container size) when the internal quality (skill at "Bucket Toss") is known, and how this compares to the inferences of an ideal observer operating with the same information.

Methods

Participants were shown a profile picture of a player, and were told whether the player made a shot, and the players' skill-level in terms of a percentile (e.g., "Jenifer is better than 99% of people at Bucket Toss"). There was a large question mark in the location where a container would be, and participants had to estimate what container was the likely target.

We wanted to convey to participants that the container was chosen independently of the player's skill (i.e. it was not a "level" they had achieved). Therefore, we described "Bucket Toss" as a carnival game of "skill and chance", in which the first part was "the luck of the spin", wherein a player spun a pointer which determined which of the five container they would then try to toss the quarter into.

This experiment was conducted online with 305 participants from Amazon Mechanical Turk, each completed eight trials. On each trial the skill percentile (1st, 25th, 50th, 75th, or 99th) and whether or not the person made or missed the shot was selected at random. The person-image and name were selected at random with the constraint that they could only appear once per subject.

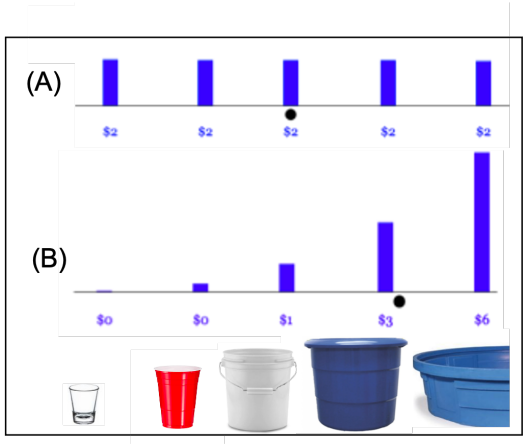

Figure 5: Participants estimated which container was the target, given that a person made/missed the shot. Responses were made by a slider that distributed bets over containers in the situation estimation experiment. (A) At the midpoint, bets were distributed uniformly over the containers. As the slider moved away from the midpoint, bets were skewed accordingly (e.g., toward large containers in B).

Response Scale Participants estimated which container they thought the player had been aiming for in the form of a bet. We changed the scale in this way to provide participants with a midpoint that corresponded to uncertainty about the target container [1]. Participants were told to distribute a $10 bet over multiple containers. They distributed these bets

by dragging a uni-dimensional slider. At the center of the scale the bets were distributed equally over the 5 containers ($2 on each; Figure 5A). As the participants dragged the toggle to the right this skewed the distribution of bets toward larger containers (Figure 5B); and vice versa when moving the slider to the left. A bar chart corresponding to the proportion of money bet on each container served as a visual cue that changed dynamically as participants moved the slider.

Participants were first given a few examples to illustrate the meaning of the response scale. For example, they were told that if betting about which container "a man had used to give his dog a bath" then the slider might be moved to the right, allocating more bets to larger containers (like the kiddie pool) than smaller containers (like the shot glass). If the bet pertained to "a container a man used to hold liquid", a slider position at the midpoint, which distributed bets uniformly over all containers, would be sensible.

Model and Results

Again, we can compare human attributions to the attributions of an ideal observer. The ideal observer uses internal quality (skill percentile) to estimate the situation strength (container size) when the player either made, or missed a shot (Figure 6A). Likewise, we obtain human observers' judgments of the most probable container size on those same types of trials (Figure 6B).

Again, we see the graded pattern of inferences predicted by probabilistic attribution: as we increase the ostensible skill of the player (disposition strength), observers' estimate a missed shot to be ever more diagnostic of container size, and made shots become ever less diagnostic. However, in these data, human judgments are not consistent with all of the predictions of the PSA model: container size is only estimated to be in a direction consistent with the observed behavior when the player makes the shot and is at or above the 50th percentile, or misses the shot and is at or below the 50th percentile. The pattern for missed shots might make sense as a consequence of the scale we used — there is no difference in judgments between the 25th and 1st percentile conditions, perhaps because people are encouraged to distribute bets homogeneously. However, the counter-intuitive pattern for made shots is more puzzling: people estimate that a 99th percentile player who made a shot is likely to be throwing into a particularly small container; this peculiar pattern might be a consequence of people not fully believing that container size is allocated randomly to players, and thus believe that an especially skillful player is likely to be matched up with an especially challenging target. This might imply a fascinating direction for future research: people do not believe situations and dispositions to be independent —- dispositions likely influence what situations we find ourselves in.

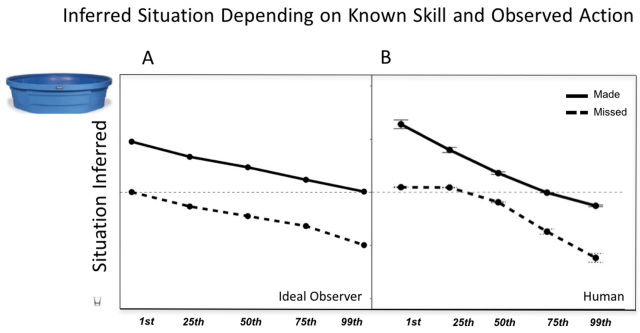

Figure 6: Experiment 3 results. Model (A) and human (B) estimates of the most likely situation strength, given that a player made (solid) or missed (dashed) a shot, as a function of the skill of the player (x-axis, as a percentile). In both cases, made shots are assumed to have reflected a larger container, and vice versa for missed shots. In both cases, inference strengths diminishes as player skill becomes more extreme to explain it away. Surprisingly, people make no inferences about container size when a player in the 75th percentile makes the shot, or a player in the 25th percentile or below misses the shot.

Discussion

Here we evaluated how people make social attributions. The dominant narrative in social psychology asserts that people make these attributions with systematic bias: they tend to over-attribute to dispositions, and under-attribute to situations. In the classic CB literature, observers are asked to judge disposition after witnessing a behavior, usually in either a neutral situation or one designed to compel a certain action. Observers tend to judge that disposition is in a direction consistent with the observed action, even when the situation was motivating the action, and this has been interpreted as a error in social reasoning and named the Correspondence Bias (CB). Although a persistent thread of research has argued that such inferences may simply reflect probabilistic inference (Trope, 1974; Ajzen & Fishbein, 1975; Morris & Larrick, 1995; Jennings, 2010; Walker et al., 2015), much of this research has relied on flexible assumptions about observers' priors and expectations to accommodate the data. For example, in previous work, we argued that a reasoning error is not required to explain this empirical result – if observers don't actually believe the situation is deterministic, an action-consistent inference about disposition would be predicted by optimal Bayesian inference, and show that when an ideal observer assumes situation strength is less than deterministic, action-consistent disposition are still inferred (Walker et al., 2015). When plausible situation strengths are assumed, optimal inference aligns well with human judgments. However,

[1]Simply choosing one of the containers would not yield meaningful behaviors in the middle of the scale. When indicating skill percentile (Experiment 2), the midpoint is a sensible choice when an observer has little or no information about a person's skill (for example, in Experiment 1 when the player made a shot that about 50% of people could make) it makes sense to guess this player is about average — its an intuitive guess when you have no other information. However, when choosing containers, if a participant has little or no useful information about container size there is no meaningful way to express this by choosing one of the containers, since the center of the scale (the 5-gallon bucket) is no more likely to have occurred than any other container.

since perception about situation strength were not reported in the original literature, we needed to make assumptions about them to accommodate the existing data. One critical assumption is that observes don't in-fact assume that these strong situations are deterministic – if they did, drawing a corresponding inference would also violate optimal Bayesian inference, and calling it a Correspondence Bias would be justified.

In the current work we aimed to clarify these positions by collecting independent observers' judgments of situation strength to constrain the PSA model, and parametrically varied both situation and disposition over a wide range – from non-constraining to nearly deterministic. We find that whether people are judging situations or dispositions people largely track the predictions of Probabilistic Social Attribution. Inconsistent with the classic CB argument, when situations are interpreted to be close to deterministic, no action-consistent inferences about disposition are made. Additionally, we find that the pattern of inference that arise from judging an unknown disposition when situation is known is similar to the pattern of inference that emerges when observers judge unknown situation when disposition is known (though with some notable oddities). Overall, there seems to be nothing biased in the way we account for the causal effects of disposition, but instead these inferences seem to track well what we would expect from inference about any unknown variable under uncertainty. Our intention here is not to suggest that our specific causal model underlies all social attribution, but instead to provide further evidence that behavior that has been interpreted as a bias in the classic CB literature is consistent with an unbiased probabilistic reasoning framework, even with objectively measured situation strengths.

References

Ajzen, I., & Fishbein, M. (1975). A bayesian analysis of attribution processes. *Psychological bulletin*, *82*, 261.

Gawronski, B. (2004). Theory-based bias correction in dispositional inference: The fundamental attribution error is dead, long live the correspondence bias. *European review of social psychology*, *15*, 183-217.

Gilbert, D., & Malone, P. (1995). The correspondence bias. *Psychological Bulletin*, *117*, 21.

Hall, S., Ali, N., Chater, N., & Oaksford, M. (2016). Discounting and augmentation in causal conditional reasoning: causal models or shallow encoding? *PloS one*, *11*, e0167741.

Hammond, K. (1996). *Human judgment and social policy: Irreducible uncertainty, inevitable error, un-available injustic*. New York: Oxford University Press.

Jennings, K. (2010). Determining the internal consistency of attitude attributions. In *Proceedings of the cognitive science society.*

Jones, E., & Harris, V. (1967). The attribution of attitudes. *Journal of Experimental Social Psychology*, *3*, 1-24.

Jones, E., Worchel, S., Goethais, G., & Grumet, J. (1971). Prior expectancy and behavioral extremity as determinants of attitude attribution. *Journal of Experimental Social Psychology*, *7*, 59-80.

Khemlani, S., & Oppenheimer, D. (2011). When one model casts doubt on another: A levels-of-analysis approach to causal discounting. *Psychological bulletin*, *137*, 195.

Lau, R., & Russell, D. (1980). Attributions in the sports pages. *Journal of personality and social psychology*, *39*, 29.

Morris, M., & Larrick, R. (1995). When one cause casts doubt on another. a normative analysis of discounting in causal attribution. *Psychological Review*, *102*, 331.

Quattrone, G. (1984). Overattribution and unit formation: When behavior engulfs the person. *Journal of personality and social psychology*, *42*, 593.

Ross, L., Amabile, T., & Steinmetz, J. (1977). Social roles, social control, and biases in social-perception processes. *Journal of personality and social psychology*, *35*, 485.

Sherman, S. (1980). On the self-erasing nature of errors of prediction. *Journal of personality and Social Psychology*, *39*, 211.

Trope, Y. (1974). Inferential processes in the forced compliance situation: A bayesian analysis. *Journal of Experimental Social Psychology*, *10*, 1-16.

Walker, D., Smith, K., & Vul, E. (2015). The fundamental attribution error is rational in an uncertain world. In *Proceedings of the cognitive science society.*