Running head: CORRESPONDENCE BIAS

On the Robustness and Generality of the Correspondence Bias

Don A. Moore
University of California at Berkeley

Derek Schatz
Accenture, PLC

Zachariah Sharek
Highmark, Inc.

Samuel A. Swift
Bowery Farming

Francesca Gino
Harvard Business School

June 24, 2020

The authors gratefully acknowledge the support of National Science Foundation Grant SES-0718691, as well as the staff and facilities of the Center for Behavioral Decision Research at Carnegie Mellon University. Address correspondence to dm@berkeley.edu.

Abstract

People routinely assume correspondence between acts and dispositions. This “correspondence bias” leads people to neglect the power of the situation to influence behavior. However, most of the research documenting the correspondence bias uses paradigms in which the situation operates in the background. The three studies presented here employ novel paradigms in which people have full information about the situation and its effect on behavior. We find that the bias persists in the face of debiasing treatments and even in situations where avoiding it amounts to a simple subtraction problem. We also find that the correspondence bias generalizes to inanimate objects. In light of this new evidence, we reconsider the psychological origins of the correspondence bias.

Key words: correspondence bias; fundamental attribution error; judgment and decision making

On the Robustness and Generality of the Correspondence Bias

When we seek to explain another person's behavior, we usually rely on distinctive attributes of that person, such as the person's personality or abilities (Gilbert & Jones, 1986; Gilbert, Krull, & Pelham, 1988; Nisbett & Ross, 1991). It seems natural to assume that a person who performs poorly must lack talent; one who performs well must possess it (Block & Funder, 1986; Epley, Savitsky, & Gilovich, 2002). While this sort of inference makes intuitive sense, what it too often ignores is the powerful influence of situational forces on human behavior (Darley & Batson, 1973). Seeing someone give a speech supporting or opposing affirmative action leads others to infer that the person must believe what he or she said, even when they know that the speaker was compelled to take that position (Van Boven, Kamada, & Gilovich, 1999).

Jones (1990) referred to the assumed correspondence between acts and fundamental dispositions as the correspondence bias. The tendency is so pervasive that Ross (1977) declared it the "fundamental attribution error." In this paper, we explore the robustness of this tendency, and suggest that it may be even more fundamental than prior theories have supposed. We find that it persists beyond what most existing theories would predict. This persistence provides novel insights into and raises new questions about the ultimate causes of the correspondence bias.

Explanations for the Correspondence Bias

One way to think about the correspondence bias is as an instance of anchoring and insufficient adjustment: People anchor on observable behavior and fail to adjust sufficiently for the influence of the situation (Tversky & Kahneman, 1974). In Jones and Harris's classic (1967)

study, for example, participants read an essay, supposedly written by another participant, that either praised or criticized Cuban leader Fidel Castro. When participants believed that the essay's author had freely chosen what to write, it came as no surprise that they assumed the essay reflected the author's views. The interesting result was that participants still made this inference even when they believed the essay's author had been compelled to take that perspective.

Believing that the pro-Castro essay writer really loved Castro, even after learning the author was forced to express that view (as in Van Boven et al., 1999) could be described as anchoring on the essay and failing to adjust sufficiently for the assignment's directive.

One possible explanation for people's failure to adjust for situational influence is that dispositional inferences might be easier than situational adjustments. Trait inferences arise in the brain's medial prefrontal cortex (mPFC) more spontaneously or automatically than situational attributions, which require more intentional processing (Kestemont, Vandekerckhove, Ma, Van Hoeck, & Van Overwalle, 2012). An effortful process of correction in the dorsolateral prefrontal cortex (DLPFC) appears to be responsible for integrating situational information into attributions--as well as inhibiting automatic reactions such as dispositional trait inferences (Brosch, Schiller, Mojdehbakhsh, Uleman, & Phelps, 2013).

Why would consideration of the situation require effortful correction? A number of explanations have noted that behavior is usually more obvious than situational influence (Gilbert, 1998; Gilbert & Malone, 1995). The essay author's love of Castro is evident for all to see in their written words. The situational pressure that compelled them to write those words is invisible. Indeed, making the situation more obvious tends to reduce the correspondence bias (Gawronski, 2004). Most of the research documenting the correspondence bias shares the key feature of the Jones and Harris study: The behavior (e.g., attitudes expressed in the essay) is

easier to identify than the situation (e.g., the situational pressures to write one essay or another). Cimpian and Salomon (2014) have emphasized this mechanism as also being a determinant of their inference heuristic, which they describe as operating through similar means as to the correspondence bias.

If the correspondence bias truly results from this sort of confusion, it should disappear when people understand both situation and behavior with equal clarity. The question we ask in this paper is whether the correspondence bias persists even when both the behavior and the situation are equally clear and presented in the same format. The four studies we present ask participants to infer dispositions when the behavior and the situation are both quantified in terms that allow people to easily apply Kurt Lewin's (1951) attributional equation. This equation, written in its simplest form, is: *Disposition = Behavior – Situation* (see Gilbert, 1995). When people know *Behavior* and *Situation* precisely, computing *Disposition* ought to come down to simple subtraction. Can the correspondence bias persist, even in this context?

The research paradigms we employ allow us to rule out several other explanations for the correspondence bias. For instance, Schwarz (Bless, Strack, & Schwarz, 1993; Schwarz, 2014) has noted the potential role of conversational norms in studies in which the researchers provide participants with better information about others' behavior than about the situation. Participants' tendency to overweight this information could be due to the assumption that the experimenter gave them specific information because it was relevant to the dispositional judgment they were asked to make. Gilbert and Jones (1986), for example, had their participants listen to verbal expressions of attitudes that they believed another person had been compelled to read aloud. The results showed that people made correspondent inferences, reporting that the other person must be liberal if the person made a majority of liberal statements. However, it was not obvious how

participants ought to have made use of the heartfelt delivery of ideology they had heard. Gilbert and Jones suggest that participants should have ignored it entirely. However, after presenting people with compelling and vivid (if misleading) information, it seems unfair to accuse them of behaving irrationally when they take it into account.

Sabini, Siepmann, and Stein (2001) argued that situational causes can be interpreted as dispositional causes. They theorize that Jones and Harris's (1967) participants were actually deciding if the essay writers were more influenced by their beliefs about Castro or their desire to please the instructor. If so, the correspondence bias Jones and Harris observed would be due to an overestimation of a particular internal cause (the writer's opinion of Castro) relative to another internal cause (the writer's desire to follow instructions) rather than a neglect of the situation.

Our experiments avoid the criticisms posed by Schwarz (2014), Sabini *et al.* (2001) and others by (1) providing participants with equally simple and clear information about both the behavior and the situation and (2) examining perceptions about objectively verifiable behaviors and situations. Our experimental paradigms extend the literature by examining the correspondence bias in a novel context, and testing the degree to which they apply to inanimate objects. Prior theories would have difficulty accounting for the persistence of the correspondence bias under these conditions. Instead, the persistence of the correspondence bias might require a different explanation. We consider the representativeness heuristic (Kahneman & Tversky, 1972) which leads people to assume a similarity between cause and effect. The world is more comprehensible and predictable when Castro opponents write anti-Castro essays and when smart people perform better on quizzes. As we will see, when the truth does not reflect appearance – such as when a person scores well on a quiz because the quiz is easy (a

situational influence) – and when the ‘adjustment’ from the anchor is made easy to compute and explicit, the representativeness heuristic can still lead to erroneous conclusions.

We report all conditions and all measures for the studies we present. Sample sizes were all determined prior to data analysis. Data and materials for all studies, as well as pre-registrations for Studies 1 and 2, are online:

https://osf.io/u6ze7/?view_only=e5fae8b6f6b74ae18c7e9b71fe60305b.

Study 1: Debiasing Anchoring

Study 1 employs a quiz-contestant paradigm and varies how explicitly we clarified the situational effects in the stimuli. In particular, we wanted to test how much of the effect was due to anchoring. We sought to do this by including a condition that ought to eliminate anchoring. We did this by instructing participants how much to adjust from the anchor. Note that anchoring is only a bias when adjustment from the anchor is insufficient. Recall, for instance, Epley and Gilovich’s (2001) demonstration of anchoring, in which participants adjusted from the year 1776 to estimate the year in which George Washington was elected President of the United States. If participants had been told, “George Washington was elected twelve years after the signing of the Declaration of Independence in 1776,” there can be little doubt that adjustment would have been sufficient and the anchor’s biasing effect would have been eliminated. Our first experiment tests whether the correspondence bias survives such an explicit adjustment hint.

Method

Seeking 90% power, and based on the effect size of $\eta^2 = .05$ from a pilot study, our pre-registered research plan called for a sample size of 150. The experiment took place through an online survey through Amazon Mechanical Turk. We restricted participation to English speakers

18 years of age or older. Following the removal of three participants due to incomplete responses, the analyses included 147 participants. Participants received \$0.50 for participating, with a bonus of \$0.50 for each correct prediction of “high scorers” (similar to Study 1) resulting in total possible compensation of \$3.00. Each participant saw the performance of each of ten contestants on a 15-item geography quiz and then predicted which contestants would perform the best on another similar quiz.

Stimuli. Each contestant had actually each completed 30 geography quiz items, twenty of which participants saw: ten medium-difficulty geography questions, five very easy questions (which they all answered correctly) and five very hard questions (which they all answered incorrectly). Participants did not see contestants’ answers to the remaining ten geography test items, but instead tried to guess how these contestants would perform answering those ten questions.

We created “easy” and “hard” quizzes as stimuli by removing one set of the latter five questions from each contestant. For example, Contestant A’s “easy quiz” showed 15 of their responses – their ten responses to the medium difficulty questions, followed by their five correct responses to the easy questions. Conversely, Contestant A’s “hard quiz” showed their same ten medium difficulty responses, but now followed by their five incorrect responses to the hard questions. We created two sets of ten contestants, with each participant randomly rating only one set. Set A contained the easy quizzes of contestants 1-5 and the hard quizzes of contestants 6-10, and Set B contained the hard quizzes of contestants 1-5 and the easy quizzes of contestants 6-10. Participants saw not only the questions and contestant answers, but also a green check mark next to each correct answer and red X marking incorrect answers. At the bottom of the page was the total contestant score out of 15.

Manipulation. Within-subjects, we presented participants with one ‘set’ of contestants in random order. Instead of presenting numbered contestants, we assigned contestants letters of the alphabet A-J in random order. Our between-subjects manipulation entailed a control condition and an explicit condition. The control condition simply provided the task instructions to participants. The explicit condition additionally informed participants that there are actually two different quizzes, a hard quiz on which everyone answered the last five questions incorrectly, and an easy quiz on which everyone answered the last five questions correctly. The bottom of every contestant’s quiz contained the respective quiz average, in order to distinguish hard quizzes from easy quizzes. The average easy quiz score was 12 out of 15, and the average hard quiz score was 7 out of 15.

Measures. Participants rated the knowledgeability of each of the ten contestants they saw. In addition, participants bet on five of the ten contestants to be the top scorers on the ten-item geography test. Participants earned \$.50 for each contestant they correctly identified as being among the top half of scorers.

Results and Discussion

Knowledgeability Ratings. We computed two averages for ratings of knowledgeability: one for the five contestants whose easy quizzes participants saw, and another for the five contestants whose difficult quizzes participants saw. These averages were used in a 2 (quiz difficulty) x 2 (condition) mixed ANOVA with repeated measures on the first factor. The results show a significant within-subjects main effect of test difficulty on knowledgeability, $F(1, 146) = 116.60, p < .001, \eta^2 = .45$, and a significant interaction effect of difficulty and condition, $F(1, 146) = 25.53, p < .001, \eta^2 = .09$. See Figure 1.

Figure 1. Rated knowledgeability of contestants who had taken the easy or hard quiz, by experimental condition, Study 1. (Error bars indicate standard errors.)

Looking closer at the interaction, we subset the data by condition and ran one-way repeated measures ANOVAs for each between-subjects condition, using quiz difficulty as the independent variable to predict ratings of knowledgeability. In the control condition, difficulty significantly affects ratings of knowledgeability, $F(1,71) = 79.84, p < .001, \eta^2 = .52$. Surprisingly, difficulty also significantly affects knowledgeability in the explicit condition, $F(1,76) = 39.89, p < .001, \eta^2 = .41$, though its effect is smaller than in the control condition. These results show that participants rate easy quiz contestants as more knowledgeable than hard quiz contestants, but that the explicit condition only partially attenuated the effect, as we saw from the interaction effect of test difficulty and condition on knowledgeability reported above, $F(1,146) = 25.53, p < .001, \eta^2 = .09$.

Predictions of Higher Scorers. We conducted a one-way ANOVA predicting number of easy contestants chosen as higher scorers by condition. There was a significant effect of

condition on number of easy contestants chosen, $F(1, 146) = 18.59, p < .001, \eta^2 = .06$. One-sample t-tests reveal that both those in the control condition ($M = 3.85$ of easy quiz contestants chosen as future high scorers, out of five), $t(141) = 13.51, p < .001$, and the explicit condition ($M = 3.28$ out of five) $t(141) = 8.91, p < .001$, chose easy quiz contestants at a rate significantly greater than 50%.

Making the required adjustment of quiz score explicit to participants did attenuate the effect of correspondence bias on knowledgeability ratings. Nevertheless, participants still rated easy quiz contestants as more knowledge than those who took hard quizzes. Further, even those in the explicit condition favored easy quiz contestants to be high scorers on a new quiz. Considered together, these results suggest that anchoring and insufficient adjustment partially explains the correspondence bias we observe, but it does not explain the effect entirely and room remains for another mechanism – perhaps the representativeness heuristic.

Study 2: Seeking Representativeness

Study 2 seeks to replicate results from the first study and, in addition, provide a more direct test for the role of representativeness as an explanation for the correspondence bias. Study 2 employs an even more heavy-handed debiasing intervention designed to eliminate the effect of anchoring. It also includes new questions devoted to measuring reliance on the representativeness heuristic. The representativeness heuristic invites the inference that high quiz scores are the result of knowledgeable quiz-takers. Variation in scores is therefore attributable to variation in individual knowledgeability. At the extreme, someone in the grip of the representativeness heuristic would ignore situational contributions to performance, such as quiz difficulty and attribute all variation in performance to variation the abilities of test-takers.

If the representativeness heuristic leads participants to attribute differences in scores to knowledgeability of the individual test-takers, we should observe its effects on how they apportion variation in scores. High scores on a test are ‘representative’ of an intelligent person. We expected the representativeness heuristic should lead participants to attribute too much of the variance in quiz scores to individual differences and too little variance to differences between quizzes.

Method

Study 2 was a close replication of Study 1. Like Study 1, Study 2 varied (between-subjects) the explicitness with which the effect of quiz difficulty was explained. In this study, before participants made their predictions for ‘top-half’ scorers, we provided them with a table showing participant scores. In the explicit condition, this table also had the quiz mean (either 7/15 or 12/15) next to each contestant score. We also provided a blank version of the ‘new’ quiz for which participants predicted contestants’ performance.

Participants. Our pre-registered research plan called for 150 participants recruited from MTurk. Participants each received \$1 for participating, with an additional incentive for prediction accuracy in the form of tickets for a lottery draw of \$30. Following the consent form, participants read a short description of the task before them, noting that each quiz had 15 questions and that contestants’ correct answers would be marked with a green check mark while their incorrect answers would be marked with a red X. The attention check on the following page tested participants’ comprehension of these instructions with two questions: “How many questions are there on each quiz?” and “What are wrong answers marked with on each quiz?”

Measures. As in the Study 1, participants rated the knowledgeability of each of the ten contestants they saw. In this study, participants provided a percentage chance prediction for each

of the ten contestants to score in the top-half on a new ten-item geography test. To motivate accuracy, we scored participants using the quadratic scoring rule (Selten, 1998). More accurate estimates improved participants' chances of winning the \$30 prize.

To measure participants' reliance on the representativeness heuristic, we then elicited participants' beliefs about the variance in scores attributable to quiz difficulty. We did this by asking, *"Thinking about the 10 contestants you just evaluated, what percentage of the variance in test scores you saw was due to variation in quiz difficulty? In other words - how much (in %) did quiz difficulty affect their test score? (from 0-100%)."* We elicited beliefs about variance in scores attributable to individual factors such as ability or luck by asking, *"Thinking about the 10 contestants you just evaluated, what percentage of the variance in test scores you saw was due to variation among test-takers (such as knowledge, ability, luck)? In other words - how much (in %) did the individual characteristics of each contestant affect their test score? (from 0-100%)."* Because these two ought to sum to 100%, we summed each participants' responses and then divided their answers by the total, thereby normalizing them to proportions totaling 100%.

Results and Discussion

Consistent with our pre-registered analysis plan, we first excluded the participants who failed a comprehension check, leaving us with a total sample size of 110.

We conducted a 2 (quiz difficulty) x 2 (explicitness) mixed ANOVA with repeated measures on the first factor using reported knowledgeability as the dependent measure. As in Study 1, the results reveal a significant main effect of test difficulty, $F(1, 214) = 69.09, p < .001$, $\eta^2 = .39$, and a significant interaction effect between difficulty and explicitness, $F(1, 214) = 14.31, p < .001, \eta^2 = .07$. The main effect of explicitness is not significant, $F(1, 214) = .85, p = .36, \eta^2 = .0$.

When we conduct the same mixed ANOVA on performance prediction probabilities, the results again show a significant main effect of test difficulty, $F(1, 214) = 34.89, p < .001, \eta^2 = .10$, no main effect of explicitness, $F(1, 214) = .50, p = .48, \eta^2 = .0$; and a significant difficulty by explicitness interaction, $F(1, 214) = 11.71, p < .001, \eta^2 = .05$. These interactions are consistent with the anchoring explanation: the size of the correspondence bias shrinks substantially in the explicit condition. These results, illustrated in Figure 2, replicate our primary findings from the first study. Nevertheless, the effect of test difficulty remains substantial even in the explicit condition, implying that anchoring cannot account for the entirety of the effect.

Figure 2. Rated knowledgeability of contestants who had taken the easy or hard quiz, by experimental condition, Study 2. (Error bars indicate standard errors.)

Reported variance. Consistent with the representativeness heuristic, participants attributed 50.1% of the variance to individual variation when in fact it only accounted for 24.24%, a significant difference by one-sample t-test, $t(109) = 14.67, p < .001$. This over-

attribution to individual differences is remarkable because it even includes variance rightly attributable to chance or luck, which we made explicit to participants.

We tested for the effect of the experimental manipulation of explicitness on participants' attribution to individual variation. There was no significant main effect, $t(1, 109) = .16, p = 0.87$, on percentage of variance attributed to individual.

Reported variance predicting knowledgeability and predictions. Our theory holds that over-attribution to ability should contribute to the correspondence bias. To compute this, we subtracted participants' reported variance due to ability from the actual variance due to ability. We then correlated this difference with the difference in knowledgeability between easy quiz contestants and hard quiz contestants for each participant, as a measure of correspondence bias. We also correlate this variance difference with the difference in performance probability between easy quiz contestants and hard quiz contestants for each participant. The results show marginal results for both, knowledgeability $r = .14, t(149) = 1.67, p = .10$, and prediction probability $r = .14, t(149) = 1.73, p = .08$. These results imply that the degree to which participants over-attribute scores to individual attributes (thus showing use of the representativeness heuristic) is only marginally related to participants' display of the correspondence bias.

This study replicated our prior results while adding further direct tests of representativeness as a mechanism for the correspondence bias. These results suggest that anchoring and insufficient adjustment again partially explains the correspondence bias we observe, since correcting for insufficient adjustment does not eliminate the correspondence bias. While representativeness is a plausible explanation for the remaining bias, it only marginally correlates with the dispositional errors participants commit.

Study 3: M&Ms

As with most of the research on the correspondence bias, our focus thus far has been on attributions about people. But we wonder whether the correspondence bias might be more basic and more general than that. Indeed, the natural sciences used to subscribe to what can only be called as a correspondence bias toward inanimate objects (Gilbert & Malone, 1995). In Aristotle's (350 B.C.) account, the behavior of physical objects was attributable to their individual dispositions: Light things had levity and heavy things had gravity, which explained why they rose and fell. Fortunately Galileo (1632) came along and helped scientists appreciate the power of the situation: An apple falls through air but rises through water. Objects move not only because of their natures but because of the external forces acting upon them.

Thanks in part to Kurt Lewin (1931) psychologists have become more Galilean in their theorizing about human behavior. Few psychologists now question the basic notion that human behavior is an interaction between the person and the situation (see Ross & Nisbett, 1991). Intuitive human judgment routinely neglects situational influence and instead focuses on the individual actor as the cause of behavior (Krull et al., 1999), though one should be cautious of generalizing this tendency too broadly (Miyamoto & Kitayama, 2002; Morris & Peng, 1994). However, if human intuition, unaided by the insights of Galileo's physics, is prone to show the correspondence bias for inanimate objects, is it really a property of interpersonal perception or is it a more general property of human inference? Study 3 asks this question.

This experiment put participants in the role of a shopper who has to choose from a selection of products that vary in both packaging weight and net weight of the product. This judgment was complicated by variance in the weight of the product packaging. This is a task which should put participants on guard against the ulterior motives and potential manipulation of

vendors for whom it is less costly to provide packaging than product (see Fein, 1996; Fein, Hilton, & Miller, 1990). We used as stimulus materials a set of similarly sized pint glasses of different weights, covered with colored construction paper and lids, and filled with various quantities of M&Ms. The M&Ms were the focal product, while the pint glasses of varying weights provided a situational influence. As in Studies 1 and 2, removing the biasing influence of situational forces came down to a simple subtraction problem.

Method

Experimental manipulation. The product packaging manipulation had three levels: the containers were light (112g), medium (288g), or heavy (419g). Into these, we measured contents in nine randomly chosen quantities between 50 and 300 grams. Each participant encountered nine “products,” three of each weight container with M&Ms inside as well as a fourth empty container of each type, useful for comparison. The containers were each labeled with a symbol (@, (), !, #, \$, %, ^, &, or *) to facilitate identification.

Three sets of stimuli allowed non-instrumental factors to be counterbalanced. Each set of nine cups contained the same randomly chosen nine quantities of M&Ms, but in different randomly chosen orders. The average content weight in each container condition across sets was 185g. The colors of construction paper used to identify the different container weights were also counterbalanced across sets. (See Appendix 2 for full specification of container and content weights.)

Participants. Participants were 62 volunteers recruited from a busy street in Pittsburgh, Pennsylvania. Compensation for the 15-minute experiment was \$3 in cash and a bag of M&Ms.

Dependent measures. Participants first ranked the nine “products” by sorting them from heaviest to lightest. Participants could pick up and feel the weight of each focal container as well

as the corresponding empty reference container. They could not shake the containers or look inside them.

Participants estimated the weight of each product individually. To reduce noise in the estimates, participants were truthfully told that the range of contents was between 50 and 300 grams, and that each M&M weighed approximately one gram.

Results and Discussion

Participants each made nine estimates of the net weight of the contents of the nine products they saw. Two participants failed to complete the estimation task and were removed from the analysis. The remaining 60 participants generated a total of 540 estimates, on which we conducted a regression with *container weight* and *content weight* predicting participants' estimates of content weight with random effects modeled on the participant level. Both *container weight* and *content weight* were highly significant. A one gram increase in the weight of the container led to a .208 gram increase in the estimated weight of its contents, $t = 7.79, p < .001$, and a one gram increase in the weight of the cup's contents led to an increase of .469 grams in its estimated weight, $t = 9.35, p < .001$. See Figure 2. Recall that an accurate estimate of the contents' weight requires that one completely factor out the weight of the container. The highly significant predictive value of *container weight* suggests that participants failed to factor out the weight of the container.

Figure 2. *Best-fit linear regression lines for the estimated weight of the contents given actual weight and the weight of the container.*

General Discussion

Gilbert and Malone's (1995) analysis of the correspondence bias suggested four contributing causes: (a) lack of awareness of the situation, (b) unrealistic expectations for behavior, (c) improper categorizations of behavior, and (d) incomplete corrections or insufficient adjustment. None of the first three offer plausible explanations for the data we present in this paper. While each of these may undoubtedly contribute to the correspondence bias in some contexts, none of them appear to be necessary conditions for the presence of the effect.

Could our results be due to incomplete corrections? Yes, incomplete correction may be an accurate description of the results we report. But as Studies 1 and 2 show, making the adjustment easy for participants to compute does not fully attenuate the correspondence bias. Had our participants been applying Lewin's attributional equation directly by adjusting from an

anchor, their task would have been reduced to simple subtraction. Our results suggest that this does not tell the full story.

Kahneman and Tversky (1972) described a more rudimentary error, one that they attributed to a “representativeness heuristic” in human judgment. The representativeness heuristic leads us to want the cause to look like—to be representative of—the effect. The world is more orderly when heavy things feel heavy and when smart people perform better on quizzes. Assuming that a smart person will perform well on quizzes is reasonable, but the representativeness bias can lead people to mistakenly assume the converse: that someone who performs well on a quiz is smart. This does not follow, because easy quizzes allow dunces to perform well. Nevertheless, the logic fits Reeder's schematic model of dispositional inference (Gawronski, 2003; Reeder, 1993; Reeder & Brewer, 1979) and provides an explanation for why participants neglected even the easily quantified situational effects in our experiments: the representativeness heuristic serves as an implicit schema that leads our participants to (incorrectly) underweight the situation because they believe the observed behavior can be (mostly) accounted for by a dispositional attribution.

Study 2 attempted to measure the degree to which individuals saw high scores as representative of contestant knowledgeability rather than quiz difficulty: we asked them to apportion the variance in quiz scores to individual differences and quiz difficulty. Variation in their responses are only weakly predictive of the degree to which they rated easy quiz-takers as more knowledgeable. While this weak result might appear to be a problem for the representativeness explanation, we would point to a much stronger result—namely, participants’ massive overestimate of the variance attributable to individual differences. Participants estimated that half the variance in scores was attribute to individual differences, when it was

actually only a quarter. This sizable error is entirely compatible with a representativeness heuristic that takes quiz scores as directly indicative of contestant knowledgeability and neglects the power of situational effects. If our participants made this error consistently, then we would not be able to observe much variation in their situational attributions and so would not observe a strong correlation between this question and their ratings of contestant knowledgeability. If everybody does it, there simply is not much variability between people.

In a classic demonstration of the representativeness heuristic, Tversky and Kahneman (1983) gave their participants the following brief profile: *“Linda is 31 years old, single, outspoken and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.”* Then participants had to rank the likelihood of various possible paths for Linda’s life, including: 1) *Linda is active in the feminist movement.* 2) *Linda is a bank teller.* 3) *Linda is a bank teller and is active in the feminist movement.*

It is impossible for the conjunction of two events to be more likely than either of those events alone, and most of Tversky and Kahneman’s respondents claimed to know it. But 85% of them nevertheless reported believing that Linda was more likely to be a feminist bank teller than a bank teller. Just as those people couldn’t shake the feeling that Linda sounded like a feminist, so our participants apparently couldn’t shake the feeling that those who scored well on the quiz ought to be the more knowledgeable contestants. Remarkably, this feeling persisted even among those who had just taken both a hard and an easy quiz and observed their differential scores. In truth, of course, causes are not always representative of their effects. Sometimes people who appear to have performed well in fact just had an easier task. Decisions that fall victim to the correspondence bias will wind up favoring those people. For instance, evidence suggests that

alumni from institutions with lenient grading are more likely to gain admission to graduate school, in part because their high grades are mistaken for evidence of high achievement (Moore, Swift, Sharek, & Gino, 2010). A truly meritocratic society must be able to assess merit independent of the circumstances into which an individual is born. Our results do not leave us optimistic regarding the ease with which human judgment can achieve this ideal.

Our results do, however, underscore the robustness and generality of the correspondence bias. We found that the bias generalized to inanimate objects and persisted even in situations where avoiding it amounted to a simple subtraction problem. It persisted in the face of debiasing manipulations designed to eradicate it. We are prompted to remember that Jones proposed the correspondence bias as “a candidate for the most robust and repeatable finding in social psychology” (1990, p. 138). We would tend to agree.

Appendix 1. Study 3 Stimuli Specifications

Cup Identifier	Stimuli Set	Container (g)	Contents (g)
!	A	113	168
@	A	113	148
#	A	113	296
\$	A	419	144
%	A	419	126
^	A	419	74
&	A	288	164
*	A	288	274
()	A	288	272
!	B	113	74
@	B	113	168
#	B	113	274
\$	B	419	296
%	B	419	144
^	B	419	272
&	B	288	148
*	B	288	164
()	B	288	126
!	C	113	272
@	C	113	144
#	C	113	126
\$	C	419	164
%	C	419	296
^	C	419	148
&	C	288	168
*	C	288	74
()	C	288	274

References

- Aristotle. (350 B.C.). *Physica*.
- Bless, H., Strack, F., & Schwarz, N. (1993). The informative functions of research procedures: Bias and the logic of conversation. *European Journal of Social Psychology*, 23(2), 149–165.
- Block, J., & Funder, D. C. (1986). Social roles and social perception: Individual differences in attribution and error. *Journal of Personality and Social Psychology*, 51(6), 1200–1207.
- Brosch, T., Schiller, D., Mojdehbakhsh, R., Uleman, J. S., & Phelps, E. A. (2013). Neural mechanisms underlying the integration of situational information into attribution outcomes. *Social Cognitive and Affective Neuroscience*, 8(6), 640–646.
- Cimpian, A., & Salomon, E. (2014). The inference heuristic: An intuitive means of making sense of the world, and a potential precursor to psychological essentialism. *Behavioral and Brain Sciences*, 37(05), 461–480. <https://doi.org/10.1017/S0140525X13002197>
- Darley, J. M., & Batson, C. D. (1973). “From Jerusalem to Jericho”: A study of situational and dispositional variables in helping behavior. *Journal of Personality and Social Psychology*, 27, 100–108.
- Epley, N., & Gilovich, T. (2001). Putting adjustment back in the the anchoring and adjustment heuristic: Differential processing of self-generated and experimenter-provided anchors. *Psychological Science*, 12(5), 391–396.
- Epley, N., Savitsky, K., & Gilovich, T. (2002). Empathy neglect: Reconciling the spotlight effect and the correspondence bias. *Journal of Personality and Social Psychology*, 83(2), 300–312.
- Fein, S. (1996). Effects of suspicion on attributional thinking and the correspondence bias. *Journal of Personality and Social Psychology*, 70(6), 1164–1184.

- Fein, S., Hilton, J. L., & Miller, D. T. (1990). Suspicion of ulterior motivation and the correspondence bias. *Journal of Personality and Social Psychology*, 58(5), 753–764.
- Galilei, G. (1632). *Dialogue concerning the two chief world systems*. Pisa, Italy.
- Gawronski, B. (2003). On difficult questions and evident answers: dispositional inference from role-constrained behavior. *Personality & Social Psychology Bulletin*, 29(11), 1459–1475.
<https://doi.org/10.1177/0146167203256375>
- Gawronski, B. (2004). Theory-based 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. T. (1994). Attribution and interpersonal perception. In A. Tesser (Ed.), *Advanced social psychology* (pp. 99–147). New York: McGraw-Hill.
- Gilbert, D. T. (1998). Ordinary personology. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (4th ed., Vol. 2, pp. 89–150). Boston, MA: McGraw-Hill.
- Gilbert, D. T., & Jones, E. E. (1986). Perceiver-induced constraint: Interpretations of self-generated reality. *Journal of Personality and Social Psychology*, 50(2), 269–280.
- Gilbert, D. T., Krull, D. S., & Pelham, B. W. (1988). Of thoughts unspoken: Social inference and the self-regulation of behavior. *Journal of Personality and Social Psychology*, 55(5), 685–694.
- Gilbert, D. T., & Malone, P. S. (1995). The correspondence bias. *Psychological Bulletin*, 117(1), 21–38.
- Jones, E. E. (1990). *Interpersonal perception*. New York: Freeman.
- Jones, E. E., & Harris, V. A. (1967). The attribution of attitudes. *Journal of Experimental Social*

Psychology, 3, 1–24.

Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness.

Cognitive Psychology, 3(3), 430–454.

Kestemont, J., Vandekerckhove, M., Ma, N., Van Hoeck, N., & Van Overwalle, F. (2012).

Situation and person attributions under spontaneous and intentional instructions: an fMRI study. *Social Cognitive and Affective Neuroscience*, 8(5), 481–493.

Krull, D. S., Loy, M. H.-M., Lin, J., Wang, C.-F., Chen, S., & Zhao, X. (1999). The fundamental attribution error: Correspondence bias in individualist and collectivist cultures.

Personality and Social Psychology Bulletin, 25(10), 1208–1219.

Lewin, K. (1931). The conflict between Aristotelian and Galileian modes of thought in contemporary psychology. *Journal of General Psychology*, 5, 141–177.

Lewin, K. (1951). *Field theory in social science*. New York: Harper & Row.

Miyamoto, Y., & Kitayama, S. (2002). Cultural variation in correspondence bias: The critical role of attitude diagnosticity of socially constrained behavior. *Journal of Personality and Social Psychology*, 83, 1239–1248.

Moore, D. A., & Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review*, 115(2), 502–517.

Moore, D. A., Swift, S. A., Sharek, Z., & Gino, F. (2010). Correspondence bias in performance evaluation: Why grade inflation works. *Personality and Social Psychology Bulletin*, 36(6).

<https://doi.org/10.1177/0146167210371316>

Morris, M. W., & Peng, K. (1994). Culture and cause: American and Chinese attributions for social and physical events. *Journal of Personality and Social Psychology*, 67(6), 949–971.

Nisbett, R., & Ross, L. (1991). *The person and the situation*. NY: McGraw Hill.

- Reeder, G. D. (1993). Trait-Behavior Relations and Dispositional Inference. *Personality and Social Psychology Bulletin*, 19(5), 586–593. <https://doi.org/10.1177/0146167293195010>
- Reeder, G. D., & Brewer, M. B. (1979). A Schematic Model of Dispositional Attribution in Interpersonal Perception. *Psychological Review*, 86(1), 61–79.
<https://doi.org/10.1037/0033-295X.86.1.61>
- Ross, L. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 10, pp. 173–220). New York: Academic.
- Ross, L., & Nisbett, R. E. (1991). *The person and the situation: Perspectives of social psychology*. New York: McGraw-Hill.
- Sabini, J., Siepmann, M., & Stein, J. (2001). The really fundamental attribution error in social psychological research. *Psychological Inquiry*, 12(1), 16–21.
<https://doi.org/10.1207/S15327965PLI1201>
- Schwarz, N. (2014). *Cognition and communication: Judgmental biases, research methods, and the logic of conversation*. Psychology Press.
- Selten, R. (1998). Axiomatic characterization of the quadratic scoring rule. *Experimental Economics*, 1(1), 43–61.
- Snyder, M., & Jones, E. E. (1974). Attitude attribution when behavior is constrained. *Journal of Experimental Social Psychology*, 10(6), 585–600.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124–1131.
- Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90(4), 293–315.

Van Boven, L., Kamada, A., & Gilovich, T. (1999). The perceiver as perceived: Everyday intuitions about the correspondence bias. *Journal of Personality and Social Psychology*, 77, 1188–1199.