When Sources Honestly Provide Their Biased Opinion: Bias as a Distinct Source Perception With Independent Effects on Credibility and Persuasion

Laura E. Wallace[1] 🆔, Duane T. Wegener[1], and Richard E. Petty[1]

Abstract

Anecdotally, attributions that others are biased pervade many domains. Yet, research examining the effects of perceptions of bias is sparse, possibly due to some prior researchers conflating bias with untrustworthiness. We sought to demonstrate that perceptions of bias and untrustworthiness are separable and have independent effects. The current work examines these differences in the persuasion domain, but this distinction has implications for other domains as well. Two experiments clarify the conceptual distinction between bias (skewed perception) and untrustworthiness (dishonesty) and three studies demonstrate that source bias can have a negative effect on persuasion and source credibility beyond any parallel effects of untrustworthiness, lack of expertise, and dislikability. The current work suggests that bias is an independent, but understudied source characteristic.

Perceptions that others are biased—that they have a skewed perspective—pervade many of society's most important dialogues. In the United States, every 4 years, presidential candidates accuse the media of being biased. Consider a Google Trends analysis of Google News searches in the United States including the word "biased." Such searches spiked around both the 2012 and 2016 Presidential elections (Figure 1). Thus, at times, perceptions of bias may have a nationwide impact in addition to the more disparate impact they have when people perceive bias in nonpolitical domains or in non-election years. Given the pervasiveness and potential impact of perceptions of bias, as well as the long history of persuasion research, it seems reasonable to assume that researchers would have clearly demonstrated consequences of perceiving message sources as biased.

Persuasion researchers have occasionally attempted to examine consequences of what might be considered source bias by manipulating source vested interest. Source vested interest refers to the communicator having something to gain by successfully persuading recipients (e.g., a store owner recommending shopping at his store). An interesting feature of this work has been that source vested interest effects have been discussed interchangeably as effects of source "bias" and "untrustworthiness." For example, Walster, Aronson, and Abrahams (1966) demonstrated that when sources

advocated against rather than for their own interests (e.g., a criminal advocating for more powerful prosecutors), they were perceived as more honest, which led to increased perceptions of source credibility and persuasion. However, Koeske and Crano (1968) described these findings as increasing perceived source objectivity (i.e., the opposite of bias). Relatedly, when describing a Hovland and Mandell (1952) study that examined consequences of source vested interest, Hovland, Janis, and Kelley (1953) referred to the speaker with a vested interest as "motivated" to take the position they did (which seems akin to bias), but they referred to the source without a vested interest as "honest" (which seems akin to trustworthiness), but also "fair" and "impartial" (which seem more akin to a lack of bias). Thus, in discussing source vested interest, early researchers often conflated the concepts of bias and untrustworthiness. Furthermore, it seems likely that these manipulations would have affected perceptions of both bias and untrustworthiness (see Wallace, 2019, for evidence of this). Thus, based on the previous research, it would be

[1]The Ohio State University, Columbus, USA

Corresponding Author:
Laura E. Wallace, Department of Psychology, The Ohio State University, 1835 Neil Avenue, Columbus, OH 43210, USA.
Email: wallace.748@osu.edu
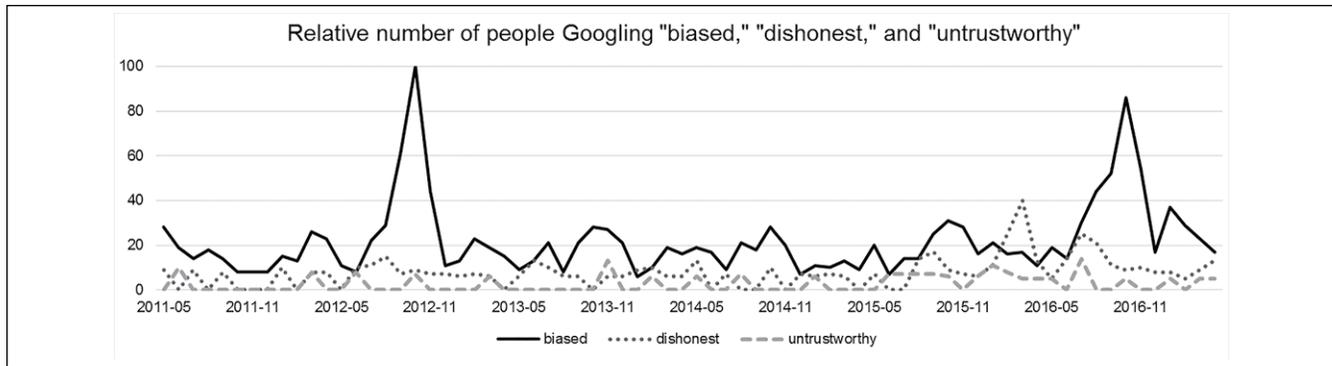
Figure 1. Google News trends analysis for "biased," "dishonest," and "untrustworthy" April 5, 2011 to April 5, 2017.
Note. Y axis represents search interest relative to the highest point on the chart for the time period.

difficult to draw conclusions about the effects of source bias on persuasion separate from effects of source untrustworthiness. Other than this source vested interest work, research in the domain of persuasion has largely overlooked perceptions of bias.[1]

Conceptual Differences Between Bias and Untrustworthiness

In its original conceptualization in the persuasion literature, trustworthiness was defined as "the communicator's intent to communicate the assertions that he considers most valid" (Hovland et al., 1953, p. 21). Consistent with this definition, we associate trustworthiness primarily with honesty, whereas we regard bias as having a perspective that is skewed (potentially by a motivation to hold a particular view or what Kruglanski, 1989, called a "need for specific closure"). Untrustworthy people intentionally present false information, whereas trustworthy, but biased people do not intend to deceive, but instead provide their honest, but skewed, perspective. Trustworthiness and bias should, therefore, represent two distinct impressions, although some antecedents, such as the person having a vested interest, might lead to perceptions of both dishonesty and bias. Of most importance, we postulate that there should be a variety of contexts in which bias and untrustworthiness are clearly separable and have independent effects. An initial suggestion that this may be the case appears in Figure 1 in which the lines representing Google News Searches for "untrustworthy" and "dishonest" do not show the same pattern as "biased," suggesting that "untrustworthy" and "dishonest" were not insults that were regularly thrown at the media during the 2012 and 2016 presidential campaigns the way that "biased" was. As an additional example, many grandparents are motivated to view their grandchildren positively. Thus, grandparents raving about their grandchildren might provide their honest, but skewed perspective.

Colloquially at least, one might say that if grandparents are biased, many would see them as "untrustworthy" sources

of information (i.e., their bias would make them untrustworthy). When doing so, such statements reflect use of the word "trustworthiness" to refer more broadly to the credibility of the information. In this colloquial example, "trustworthiness" is used to indicate that the source's information might be inaccurate due to the bias, not that the source might be intentionally dishonest, the most common use of trustworthiness in the persuasion literature. In accord with this view, the current research aims to show empirically that research participants' perceptions of trustworthiness are driven more by information about source honesty than by perceptions of source bias (Preliminary Studies 1a and 1b).

Bias as a Potential Contributor to Source Credibility

In the persuasion literature, trustworthiness has been discussed as a key component of credibility, which refers to the overall sense that a source can be relied upon. Trustworthiness, along with expertise, has been considered a pillar of credibility (Cooper, Blackman, & Keller, 2016; Hovland et al., 1953; McGuire, 1985; Petty & Cacioppo, 1981; Petty & Wegener, 1998). One reason that the distinction between bias and trustworthiness might matter is that bias could serve as a third basis for perceived source credibility separable from trustworthiness and expertise. Credibility of the source is perhaps the most studied variable in all of persuasion research. Over 60 years of research has indicated that credible sources are more persuasive than noncredible sources (Hovland et al., 1953), though there are a number of moderators of this effect and a number of potential processes involved (Chaiken, Liberman, & Eagly, 1989; Petty & Cacioppo, 1986). Furthermore, real-world situations, such as the testimony that Dr. Christine Blasey Ford and Brett Kavanaugh provided to the U.S. Senate as a part of Kavanaugh's nomination to the U.S. Supreme Court featured the credibility of the speakers as critical to the decision at hand. Given the influence that source credibility has on persuasion and its practical application to real-world situations, it is essential to

understand its underlying components. We propose that perceptions of the source as relatively (un)biased can contribute to perceptions of the source's credibility beyond any perception that the source is expert or trustworthy.[2] That is, recipients might not rely on a source because they perceive that the source has a skewed perspective (i.e., is biased), is uninformed about the topic (i.e., has low expertise), and/or is lying (i.e., is untrustworthy). As such, bias should also have its own independent consequences on persuasion beyond expertise or trustworthiness.

Overview of Studies

In the current research, we examined conceptual differences between bias and untrustworthiness and tested independent effects of bias on credibility and persuasion. In two preliminary experiments (Studies 1a and 1b), we tested conceptual differences between bias and untrustworthiness. Then, Study 2 examined whether bias would have independent effects on persuasion, beyond those explained by potentially related perceptions of trustworthiness, expertise, and liking. Studies 3a, 3b, 3c, and 4 examined effects of source bias on persuasion through perceptions of source credibility. Thus, the primary goals of our series of studies were to disassociate source bias and untrustworthiness and demonstrate that bias can have independent negative effects on credibility and, therefore, persuasion.

Preliminary Studies 1a and 1b

Across an initial two studies, we tested conceptual differences between source bias and untrustworthiness by presenting participants with descriptions of a source that were consistent with our hypothesized conceptualization of either bias or untrustworthiness. Then we measured the extent to which participants perceived the source as biased and untrustworthy. We hypothesized that when the source is motivated versus unmotivated to hold a particular position, this difference in description would drive differences in the perception of bias more so than the perception of trustworthiness. Conversely, when the source is described as honest versus dishonest, the difference between these descriptions would predict perceptions of trustworthiness more so than perceptions of bias. In essence, these preliminary studies assess our assumption about the different lay meanings of bias and trustworthiness.

Method

Participants. One hundred twenty-three participants in Study 1a and one hundred twenty participants in Study 1b were recruited through Amazon's Mechanical Turk.

Design and procedure. After consenting to participate, participants were asked to imagine that they were about to receive a persuasive message from a source. Then they read one of the four possible descriptions of the source and were asked the extent to which they would view them as biased and trustworthy. We manipulated whether participants were given a description of the source intended to convey bias or trustworthiness. We also manipulated whether the description conveyed the positive or the negative version of the trait (e.g., unbiased versus biased).[3] Manipulating the valence of the description in this way allowed us to examine the extent to which the dimension of interest captured the essence of the trait. For example, when a source is described as "motivated to take a particular position," participants should perceive the source as biased, but not particularly untrustworthy. Conversely, when a source is described as "open to supporting either position," participants should perceive the source as unbiased, but not necessarily trustworthy. The order in which trustworthiness and bias were measured was counterbalanced and order did not moderate the predicted Concept × Valence interactions on either bias or trustworthiness ($ps \geq .10$). Finally, participants were thanked for their time and debriefed. We provide descriptions of the operationalizations of the study variables below, but the exact study materials for every study are available in the stimulus file.

Independent variables. Because we manipulated both the concept (motivated to take a position or not, honest or not) and valence (positive, negative) of each description, participants read one of the four potential descriptions of the source. In the motivated reasoning/negative trait (biased) condition, participants read that the source was "motivated to take a particular position" (Study 1a) or "ideologically driven to take a particular position" (Study 1b). In the motivated reasoning/positive trait (unbiased) condition, participants read that the source was "open to supporting either position" (Study 1a) or "had no ideological drive to support one position or the other" (Study 1b). In the honesty/negative trait (untrustworthy) condition, participants read that the source was "willing to be dishonest" (Study 1a) or "willing to manipulate message recipients" (Study 1b). Finally, in the honesty/positive trait (trustworthy) condition, participants read that the source was "committed to being honest" (Study 1a) or "had the best interests of the message recipients in mind" (Study 1b).

These definitions composed the 2 × 2 design and allowed us to create two predictor variables to capture the concept and valence manipulations. The "concept" variable captured whether the source was described as motivated/unmotivated to take a position (coded as .5) or honest/dishonest (coded as −.5). The "valence" variable captured whether the source was described as positive (unmotivated or honest, coded as .5) or negative (motivated or dishonest, coded as −.5).

Dependent measures

Perceptions of bias. Perceptions of bias were measured with two items using a 9-point scale adapted from Kennedy

Table 1. Results of the Mixed GLM With Valence, Concept, and Trait Factors, Along With the Simple Two-Ways and Effects at Each Level of Concept in Studies 1a and 1b.

	Study 1a			Study 1b		
	F	p	η_p^2	F	p	η_p^2
Three-way model						
Valence	95.31	<.001	.45	27.51	<.001	.19
Concept	17.34	<.001	.13	6.47	.01	.05
Trait	.21	.65	.00	1.09	.30	.01
Valence × Concept	3.79	.05	.03	0.07	.79	.00
Valence × Trait	1.12	.30	.01	0.77	.38	.01
Concept × Trait	0.71	.40	.01	0.00	.99	.00
Valence × Concept × Trait	23.24	<.001	.16	20.36	<.001	.15
Within motivated cells						
Valence × Trait	6.86	.01	.10	13.37	.001	.18
Valence on bias simple effect	32.08	<.001	.35	22.73	<.001	.27
Valence on trust simple effect	10.99	.002	.16	1.04	.31	.02
Within honesty cells						
Valence × Trait	17.89	<.001	.23	7.42	.01	.12
Valence on bias simple effect	15.42	<.001	.21	1.42	.24	.03
Valence on trust simple effect	83.70	<.001	.59	21.88	<.001	.28

Note. GLM = generalized linear model.

and Pronin (2008). The items were "How much would you see the person as having a biased perspective?" (1 = *very biased*, 9 = *very unbiased*) and "How much would you see their opinion as a product of personal bias" (1 = *very biased*, 9 = *very unbiased*). The two items were correlated ($r = .76$, $p < .001$) and were averaged to create a composite index.

Perceptions of trustworthiness. We measured perceptions of trustworthiness on two 9-point scales adapted from Ziegler and Diehl (2011). The items were "How much would you see the person as trustworthy?" (1 = *very untrustworthy*, 9 = *very trustworthy*) and "How much would you perceive them as a trustworthy person" (1 = *very untrustworthy*, 9 = *very trustworthy*). The two items were highly correlated ($r = .92$, $p < .001$) and were averaged to create a composite index.

Results

If bias and trustworthiness indeed convey distinct constructs, a number of specific patterns should emerge. First, although the two concepts might relate, they should not be so highly correlated that they act as alternative measures of the same construct. Next, the differences in valence of the descriptions should correspond with differences in the measure that the descriptions were designed to convey more so than the measure that they were not.

Correlation. The measures of bias and trustworthiness correlated at $r = .50$, $p < .001$ in Study 1a and $r = .33$, $p < .001$ in Study 1b.[4] Within-cell correlations between bias and trustworthiness reveal an even smaller relation. In Study 1a, the

"committed to being honest" condition was the only condition in which bias and trustworthiness were significantly correlated ($r = .54$, $p = .002$). In the "dishonest" ($r = .19$, $p = .30$), "motivated" ($r = .35$, $p = .06$), and "not motivated" ($r = .07$, $p = .73$) conditions, there was no significant relation between bias and trustworthiness. In Study 1b, there were significant correlations between bias and trustworthiness for the "best interests" ($r = .45$, $p = .01$) and "not ideologically driven" ($r = .37$, $p = .04$) cells. However, there were no significant relations between bias and trustworthiness for the "willing to manipulate" ($r = .31$, $p = .09$) and "ideologically driven" cells ($r = .05$, $p = .80$), providing further evidence that bias and trustworthiness are not invariably related to each other. Although the exact r value at which bias and untrustworthiness might be considered the same construct is ambiguous, if they were the same, we would expect them to be more strongly and consistently related across the cells.

Testing lay definitions of bias and untrustworthiness. We tested our hypotheses about the definitions of bias and untrustworthiness in a mixed-design generalized linear model (GLM) with the concept and valence manipulations as between-participant factors and the measured traits (bias vs. trustworthiness) as a within-participant factor. If the concept intended to capture bias (i.e., motivated perception) better predicted bias than trustworthiness and if the concept intended to capture trustworthiness (i.e., honesty) better predicted trustworthiness than bias, we would observe a three-way interaction between the concept, valence, and trait factors. Indeed, in this model, the Concept × Valence × Trait three-way interaction was significant (Table 1, Figures 2 and 3). To dissect the

Figure 2. Effects of the "motivated to take a position" and "dishonesty" dimensions on trustworthiness and a lack of bias.
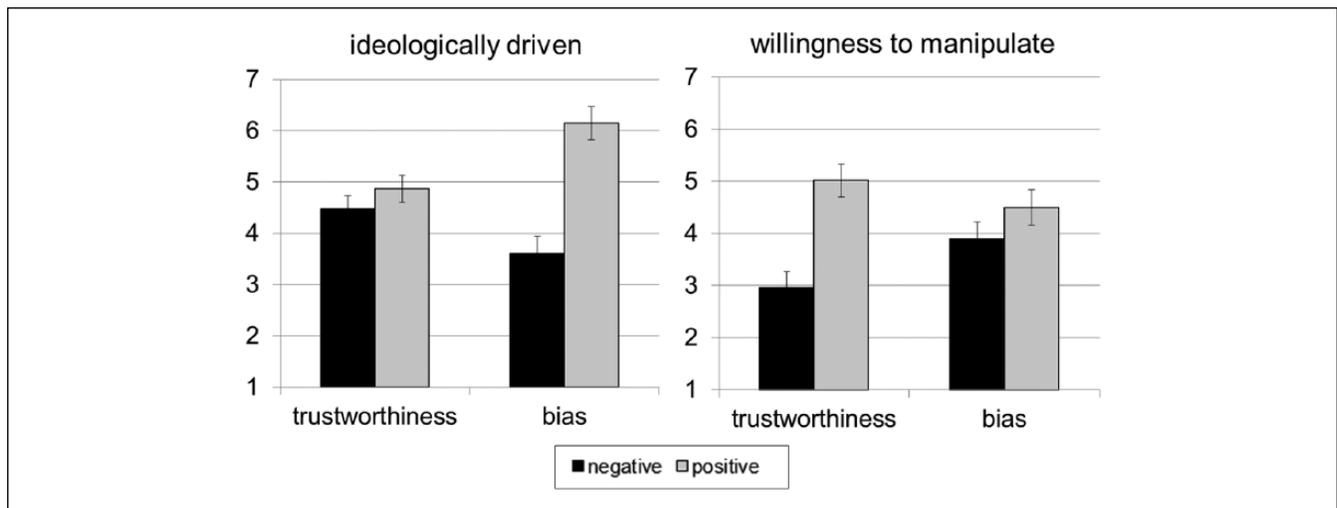
Figure 3. Effects of the "ideologically driven" and "willingness to manipulate" dimensions on trustworthiness and a lack of bias.

three-way interaction, we examined the simple two-way interactions between valence and trait type for each concept.

Motivation to take a position/ideologically driven. Among those participants who read information about the source's motivation to take a particular position, there was a significant interaction between valence and trait, such that knowing whether or not the source was motivated to take a particular position had a larger effect on bias than trustworthiness (Table 1, Figures 2 and 3).

Honesty/willingness to manipulate. Among those participants who read information about the source's honesty or willingness to manipulate, there was a significant interaction between valence and trait such that knowing whether or not the source was honest/willing to manipulate had a

larger effect on trustworthiness than bias (Table 1, Figures 2 and 3).

An alternative way to conduct these analyses is regressing bias and trustworthiness on each of the manipulations, as well as their interaction while controlling for the other perception. When conducted in this alternative manner, results support the same conclusions (see Online Appendix). Means, standard deviations, and covariances for key variables in each study are in the Online Appendix.

Discussion

Preliminary Studies 1a and 1b substantiated our conceptual distinction between bias and trustworthiness. The extent to which the source was motivated to hold a particular position primarily influenced perceptions of bias, whereas the extent

to which the source was honest primarily influenced perceptions of trustworthiness.

Study 2

In Study 2, we wanted to examine whether perceiving a source as biased would have a negative effect on persuasion beyond any effects of trustworthiness, or the other classic credibility variable, expertise. Research on the elaboration likelihood (Petty & Cacioppo, 1986) and heuristic-systematic (Chaiken et al., 1989) models of persuasion has demonstrated that source characteristics can affect persuasion through several different processes (see Petty & Wegener, 1998, for a review). Although we expect that source bias could influence persuasion through each of the processes outlined by these theories, in Study 2, we wanted to examine whether source bias would lead people to generate more negative thoughts in response to the source's claims (directional processing). Because we were interested in examining effects of bias beyond effects of trustworthiness and expertise, our stimuli were designed to hold expertise and trustworthiness at high levels, though we measured these concepts so they could be controlled.

Previous work has outlined two conditions that increase the likelihood of directional (biased) processing effects. The first is that people need sufficient motivation and ability to elaborate on the message (e.g., Petty, Schumann, Richman, & Strathman, 1993). Second, directional processing is most likely to occur when the quality of the arguments presented is ambiguous rather than clearly strong or weak (e.g., Chaiken & Maheswaran, 1994). When this occurs, characteristics of the source can disambiguate the message, leading to more positive thoughts with a positive source and more negative thoughts with a negative source. Thus, to increase the likelihood of directional processing in Study 2, we used moderately strong arguments and a relatively easy-to-process message.

In this study, we used a direct manipulation of source bias, mirroring many real-world situations, such as the self-descriptions of news organizations; for many years, Fox News labeled itself as "fair and balanced" and now as "Most Watched. Most Trusted." Politicians also often refer to the media or their political opponents as "biased" or "liars." This approach is also consistent with previous work using manipulations very similar to the constructs researchers wished to manipulate both in persuasion research (e.g., Priester & Petty, 1995) and in other domains, such as stereotyping and prejudice (e.g., Correll, Park, Judd, & Wittenbrink, 2002).

Method

Participants. Two hundred thirty-five Ohio State University undergraduates enrolled in an introduction to psychology course participated. Twenty-one participants were dropped who said that they did not read the passages. One hundred ninety-six participants completed all measures of interest.

Design and procedure. This study had a two-cell design in which we manipulated source bias. First, among other filler items, participants reported their premessage attitudes toward a university service program (Baker & Petty, 1994). Next, they received information about Dr. Brown, during which we randomly assigned participants to view him as biased or objective. Dr. Brown provided a message in favor of the university service program. Finally, participants reported the thoughts they had while reading the message, their postmessage attitudes toward the tuition program, and their perceptions of the source as biased, trustworthy, expert, and likable. When we designed this study, we were primarily concerned with creating a manipulation that affected perceptions of bias, but had relatively minimal effects on trustworthiness and expertise. As part of a battery of source perceptions adapted from another study, however, we also had a measure of source likeability. When reviews of an earlier version of this article raised the possibility that bias manipulations had their effect because of disliking the biased source, we included source liking as an additional control variable.

Independent and predictor variables

Premessage attitudes. Premessage attitudes were assessed with a single item, "How much would you support a university tuition plan through which students would have to work as part-time secretarial and maintenance staff in order to maintain current tuition levels?" (*1 = not at all, 9 = very much*).

Source bias. In both conditions, Dr. Brown was depicted as expert and trustworthy. Then participants read that "Dr. Brown has the reputation of being quite [biased and one-sided/objective and balanced] in his views on the university service program. He is a strong advocate for the program [and is not/but is always] open to considering other perspectives on this issue," in the biased and objective conditions, respectively.[5]

Dependent variables

Thoughts. After participants read the university service program message, they received thought listing instructions adapted from prior studies (see Wegener, Downing, Krosnick, & Petty, 1995). They were instructed to list up to six thoughts they had while reading the passage. Research assistants coded the thoughts for valence (1 = *positive*, 0 = *neutral*, and −1 = *negative*) and relevance or not to the university service program. An index of the thoughts was created by adding the valence of only the relevant thoughts and dividing by the total number of relevant thoughts.

Postmessage attitudes toward the university service program. Postmessage attitudes toward the university service
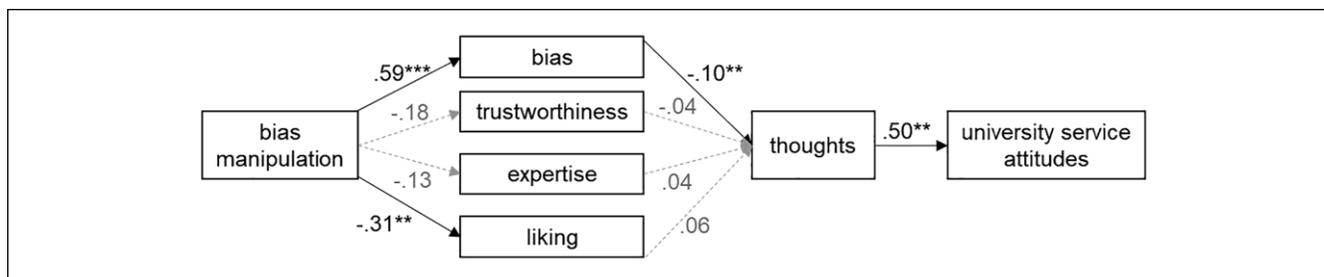
Figure 4. Parallel mediation analysis testing the effect of the bias manipulation on attitudes through perceptions of bias, trustworthiness, expertise, and liking through thoughts, controlling for premessage attitudes in Study 2.
*p < .05. **p < .01. ***p < .001.

program were measured with three 7-point items. As an example, one item was "How much is the university service program a good idea?" (1 = *not at all*, 7 = *very much*). The items were averaged to create an attitude composite, α = .95.

Perceptions of bias. Perceptions of bias were measured with three items anchored with, 1 = *not at all* to 9 = *very much*, and averaged to create a composite, α = .72. As an example, one item was "To what extent do you feel that Dr. Brown's position on the university service program is a product of personal bias?"

Perceptions of trustworthiness. Perceptions of trustworthiness were measured with three 9-point items anchored with "not at all" and "very much" and averaged to create a composite, α = .76. As an example, one item was, "To what extent does it seem like Dr. Brown is trustworthy?"

Perceptions of expertise. Perceptions of expertise were measured with three 9-point items adapted from Tormala, Briñol, and Petty (2006). As an example, one item was "How qualified did you think that Dr. Brown was to speak about the proposed tuition plan?" (1 = *not at all*, 9 = *very much*). The items were averaged to create a composite, α = .89.

Liking. Source liking was measured with three items on 9-point scales. As an example, one item was "How much do you like Dr. Brown?" (1 = *not at all*, 9 = *very much*). The items were averaged to create a composite, α = .90.

Results

Manipulation check. The biased source was seen as significantly more biased (*M* = 6.69, *SD* = 1.48) than the unbiased source (*M* = 5.50, *SD* = 1.60), *t*(194) = 5.40, *p* < .001, 95% CI = [0.76, 1.63], *d* = .78.

Mediation Analyses

We predicted that source bias should lead people to generate more negative thoughts in response to the message, which

should lead to less positive postmessage attitudes controlling for premessage attitudes (i.e., persuasion). We wanted to ensure that the effects of the bias manipulation on thoughts and attitudes were due to bias rather than other potentially related perceptions. As such, we conducted a parallel mediation analysis in which we first examined the effect of the bias manipulation on each of the perceptions. Next, we regressed thoughts on all of the perceptions and the bias manipulation. Finally, we regressed attitudes on thoughts, each of the perceptions, and the bias manipulation. In each of these analyses, we controlled for the effects of premessage attitudes because we were interested in attitude change. The results for each key pathway in this model are available in Figure 4. Importantly, the key paths from the bias manipulation to the bias measure to thoughts to attitudes were all significant. For each study, results for each predictor in each regression equation are available in the Online Appendix. We used 10,000 Monte Carlo simulations to calculate indirect effects through each of the perceptions that were entered as parallel mediators. This resulted in a significant indirect effect due specifically to source bias, 95% CI = [−0.062, −0.006]. There were not significant indirect effects through liking, 95% CI = [−0.027, 0.000]; trustworthiness, 95% CI = [−0.005, 0.018]; or expertise, 95% CI = [−0.012, 0.003], demonstrating that effects of perceived bias went beyond any effects of the bias manipulation on the other source perceptions.

Some readers may be concerned about using covariates in regression frameworks because measurement error can at times inflate the Type I error rates for incremental validity claims (Westfall & Yarkoni, 2016). As such, we also conducted SEM (structural equation modeling) analyses parallel to the primary analyses of interest. Those analyses result in the same conclusions reported in the text (see Online Appendix).

Discussion

Study 2 demonstrated that source bias versus objectivity can lead message recipients to process a message more negatively. Ultimately, this directional processing resulted in

reduced persuasion. Furthermore, these effects occurred beyond any effects of trustworthiness, expertise, and liking, suggesting that source bias is separable from those other traditional source characteristics.

Studies 3a, 3b, and 3c

As noted earlier, source credibility has long been conceptualized as the combination of trustworthiness and expertise. Studies 3a, 3b, and 3c were designed to test the hypothesis that a (lack of) bias, beyond effects of expertise and trustworthiness, should contribute to the perception of source credibility, which should then affect directional processing.

Method

Participants. Four hundred eighty-eight people (125 Mechanical Turk workers in Study 3a, 208 Mechanical Turk workers in Study 3b, and 155 Ohio State University undergraduates taking an introduction to psychology course in Study 3c) participated. Seventeen participants' data were excluded because they reported not reading about the source or the message. Four hundred seventy participants completed all measures of interest.

Design and procedure. After consenting to participate in the study, participants were asked to imagine that a political campaign was happening in their locality. They read that Cami was campaigning against Jim Smith for county commissioner. Cami was described as highly knowledgeable and trustworthy, as well as either biased or objective, depending on participants' random assignment to condition. Participants reported the extent to which they saw Cami as biased, expert, trustworthy, and credible. In Study 3c (collected after a previous round of reviews), they also reported how much they liked Cami. After source liking was raised as a potential reason for effects of source bias, we ran Study 3c as a replication of Study 3a to examine whether perceptions of source liking contribute to perceptions of source credibility in this paradigm, as Studies 3a and 3b did not contain measures of liking.

Next, Cami provided a message opposing Jim Smith. To increase the likelihood that directional processing effects would occur, in Study 3a, we used moderately strong arguments. In Study 3b, we also included a "strong argument" condition, which was intended to function as a condition in which directional processing would be less likely to occur (e.g., Chaiken & Maheswaran, 1994). However, the results demonstrated that directional processing occurred in both conditions, suggesting that the "strong" arguments were still ambiguous enough that source perceptions could influence message recipients' reactions.[6] As such, Study 3b functions as a replication of Study 3a in both argument quality conditions. For purposes of efficient presentation, we collapsed across all three studies and argument quality conditions.[7]

After reading the passage, participants reported the thoughts that had come to mind and their attitudes toward Jim Smith. Finally, participants were thanked for their participation and debriefed.

Independent variable: Perceptions of bias. Participants read that Cami was either objective or biased in her view of Jim Smith. We also attempted to hold perceptions of trustworthiness and expertise high across conditions.

Dependent measures

Source bias, trustworthiness, expertise, and liking. We measured perceptions of source bias, trustworthiness, expertise, and liking (in 3c) very similarly to Study 2. One exception was that in Study 2, the trustworthiness items were more abstract than the bias items. However, bias and trustworthiness could both be broad, applying to many topics, or specific, applying to only one or a few topics. Thus, we attempted to match the breadth of the bias and trustworthiness measures by using trustworthiness questions specific to the topic in the scenario. For example, one of the items was, "When Cami shares her opinion on Jim Smith with others, to what extent do you think she shares her honest opinion? ($1 = $ *very dishonest*, $9 = $ *very honest*)." The two bias items ($r = .82$), the two trustworthiness items, ($r = .67$), the two expertise items ($r = .73$), and the two liking items (in Study 3c, $r = .67$) were highly correlated. As such, each of these pairs of items were combined to create indices of bias, trustworthiness, expertise, and liking.

Credibility. Credibility was measured with five items on a 9-point scale. As an example, participants were asked, "How much do you see Cami as a credible source about Jim Smith?" ($1 = $ *very noncredible*, $9 = $ *very credible*). The five credibility items were averaged to create an overall index of perceived credibility ($\alpha = .96$).

Thoughts. Thoughts were measured immediately after participants read Cami's message using the same procedure as in Study 2. In this study, participants also rated their own thoughts. We calculated a thought index the same as in Study 2, but using participant rated, rather than research assistant rated thoughts.

Attitudes toward Jim Smith. Attitudes toward Jim Smith were measured on three 9-point scales, and were averaged to create a composite, $\alpha = .97$. No premeasure was taken as Jim Smith was a novel attitude object.

Results

Manipulation checks. Participants viewed the source in the biased condition ($M = 6.35$, $SD = 2.33$) as significantly more biased than in the objective condition ($M = 4.91$, $SD = 1.98$), $t(468) = 7.26$, $p < .001$, 95% CI = [1.06, 1.84], $d = .67$.
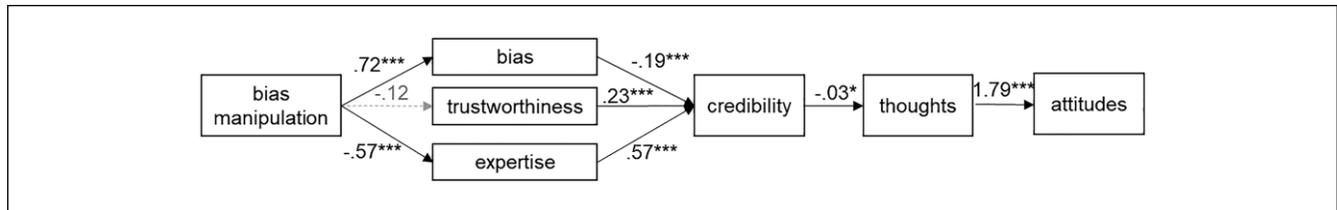
Figure 5. Parallel mediation analysis testing the effect of the bias manipulation on attitudes through perceptions of bias, trustworthiness, and expertise to credibility, then thoughts in Studies 3a, 3b, and 3c.
*p < .05. **p < .01. ***p < .001.

Effects of the bias manipulation on credibility. The biased source (M = 3.83, SD = 1.74) was viewed as significantly less credible than the objective source (M = 5.53, SD = 1.80), t(468) = −10.40, p < .001, 95% CI = [−2.02, −1.38], d = −.96.

Effect of source bias on attitudes mediated through credibility and thoughts. The primary purpose of these studies was to examine whether source bias would have a negative effect on credibility, which would then lead to reduced persuasion through its effects on participants' thoughts. Because Cami advocated against Jim Smith in the current study, perceiving her as biased should lead to more favorable thoughts and attitudes toward Jim Smith. Consistent with the analyses in Study 2, we entered our variables of interest into a parallel sequential mediation model (Hayes, 2013). The bias manipulation influenced bias, which decreased credibility, which then led to more positive thoughts about the candidate, which led to more positive attitudes toward Jim Smith, the candidate who Cami opposed (Figure 5). The mediation model tested parallel mediation paths through expertise and trustworthiness as well. Importantly, there was a significant indirect effect through perceptions of bias, 95% CI = [0.001, 0.017], providing further evidence that the effect on persuasion was due at least in part to perceptions of bias beyond the other source perceptions (this is a positive indirect effect because Cami opposed Jim Smith). In this study, there was also a significant indirect effect through expertise, 95% CI = [0.003, 0.039], but not trustworthiness, 95% CI = [–0.001, 0.005], suggesting that the manipulation also had an independent effect through perceptions of expertise.

Controlling for perceptions of liking. Because in the review process, liking was raised as a possible reason for the bias effects observed here, we also collected data including a measure of liking (Study 3c). Indeed, participants perceived the biased source (M = 5.34, SD = 1.05) as less likable than the objective source (M = 6.01, SD = 1.30), t(142) = −3.42, p = .001, 95% CI = [−1.06, −0.28], d = −.57. However, when we regressed perceptions of credibility on the bias manipulation controlling for liking, trustworthiness, and expertise, the bias manipulation continued to have a significant effect on perceptions of credibility, b = −.42, t(139) = −3.99,

p < .001, 95% CI = [−0.63, −0.021], r = .32, suggesting that the bias manipulation influenced credibility above and beyond any influences of the bias manipulation on liking of the source.

Discussion

Studies 3a-c demonstrated that viewing a source as biased led participants to perceive the source as less credible. This lack of credibility led participants to generate more positive thoughts about Jim Smith—the candidate the source opposed. These thoughts then translated to more positive attitudes toward Jim Smith. Importantly, all of these effects held controlling for perceptions of expertise and trustworthiness, the two classic pillars of source credibility, as well as controlling for likability (in Study 3c).

Study 4

Directly telling participants that a source is biased or objective mirrors many real-life situations in which people directly learn this information from the source or another person, but there are also times when people infer that a source is biased or objective without being told explicitly. In Study 4, we manipulated perceptions of bias by creating conditions in which participants would be more versus less likely to infer that a source was biased rather than directly telling them that a source was biased.

Analyses of the previous studies suggested that perceived bias contributes to source credibility and persuasion beyond source likeability. However, it is unclear whether people can perceive sources as biased without disliking them. Thus, we sought to conceptually replicate Studies 2 and 3a-c with an indirect manipulation that only influenced perceptions of bias, not liking, trustworthiness, or expertise. We did so by examining links between perceived source bias and credibility, and between credibility and persuasion. In this study, we described a scenario in which different people (one of which was our objective or biased source) advocated for different decisions in a two-option choice. Rather than assessing thoughts about each of the two options, we simply measured attitudes toward the advocated option and decisions for how to allocate resources across the two options.

Method

Participants. One hundred sixty-nine Ohio State University undergraduate students participated in this study for course credit. In the other studies in this package, we excluded participants who reported that they did not read about the source. Because there was no specific description of the source, but information about each of the people involved in discussing the options, this item did not apply as well to this paradigm, so we did not screen participants prior to analysis on the basis of this type of item.

Design and procedure. This study employed a two-cell design in which we manipulated the likelihood that a target source would be viewed as biased. After consenting to participate, participants read a conversation that aid workers had when trying to decide how to allocate resources at the beginning of an Ebola outbreak. The participants read that at the beginning of outbreaks like this, aid workers often have to make tough decisions about how to allocate limited resources. The aid workers only had enough resources to treat either Rutu, a rural area where the disease had started, or Poko, a city that could serve as a hub from which the disease would spread. The conversation included three aid workers: Steve, Roger, and Paul. Steve played a facilitating role, Roger argued that they should send their resources to Rutu, and Paul argued that they should send their resources to Poko. Roger was our target source and the transcript contained a manipulation of his potential bias (described shortly).

Roger and Paul introduced themselves as having advanced degrees in disaster management from prestigious universities so that participants would infer that they were highly expert across conditions. After participants read the transcript of the conversation, they were asked how biased, trustworthy, expert, and likable (counterbalanced) they thought Roger was. Then they reported how credible he seemed, rated their attitudes toward sending aid to Rutu, and made bipolar judgments about whether they wanted to allocate aid to Rutu or Poko first and what percentage of resources they wanted to send to each location. Finally, participants were thanked for their time and debriefed about the purpose of the study.

Perceptions of bias. We manipulated perceptions of bias though two exchanges. First, in the biased, but not objective, condition, Roger's introduction included that he had done his Peace Corps service in Rutu. In addition, in the biased condition, after Roger argued in favor of sending resources to Rutu, Paul insinuated that Roger's experience with Rutu might be coloring his perception. In both conditions, Paul asked Roger whether he really thought that they should send resources to Rutu first. Roger replied that he really thought so and provided some reasons. Full details of the exchange are available in the Online Appendix.

Dependent measures

Perceptions of bias. To reduce measurement error, we included four bias items in this study. All items were measured on a 9-point scale anchored with 1 = *very biased* and 9 = *very unbiased*, and were similar to the items used in our other studies. The items were highly related, $\alpha = .96$, and were averaged together to create an index. We recoded this variable before analyses so that the direction of this measure would be consistent across studies.

Perceptions of trustworthiness. Perceptions of trustworthiness were measured with four items anchored with 1 = *not at all*, and 9 = *very much*. Similar to Studies 3a-c, we attempted to equate the scope of the bias and trustworthiness items by making the trustworthiness items situation-specific. The items were highly related, $\alpha = .95$, and were averaged to create an index.

Perceptions of expertise. Perceptions of expertise were measured with four items measured on 9-point scales anchored with 1 = *not at all* to 9 = *very much*. These items were similar to those in the previous studies, highly related to each other, $\alpha = .94$, and averaged to create an index.

Perceptions of likeability. We asked how much participants liked Roger with four items that were similar to the previous studies. These items were measured on a 9-point scale anchored with 1 = *not at all* to 9 = *very much*. These items were highly related, $\alpha = .93$, and were averaged to create an index.

Perceptions of credibility. Perceptions of source credibility were measured with three items, each on 9-point scales anchored with, 1 = *very noncredible* to 9 = *very credible*. The three items were averaged to create an index, $\alpha = .95$.

Attitudes toward sending aid to Rutu. Attitudes toward sending aid to Rutu were assessed using three items on 9-point scales anchored with, 1 = *not at all* to 9 = *very much*. These items were highly related, $\alpha = .93$, and were averaged to create an index of attitudes. As with the prior study, no initial attitudes were taken because this was an unfamiliar topic.

Allocating resources to Rutu or Poko first. Participants also responded to two bipolar items intended to assess the extent to which they wanted to allocate resources to Rutu or Poko first. As an example, they were asked "How much do you think the aid team should send resources to Rutu or Poko first?" (1 = *send resources to Poko first*, 9 = *send resources to Rutu first*). These items were highly correlated, $\alpha = .93$, and were averaged together.

Percentage of resources allocated between Rutu and Poko. Finally, using a sliding scale, we asked participants what percent of resources they would allocate to Rutu ver-
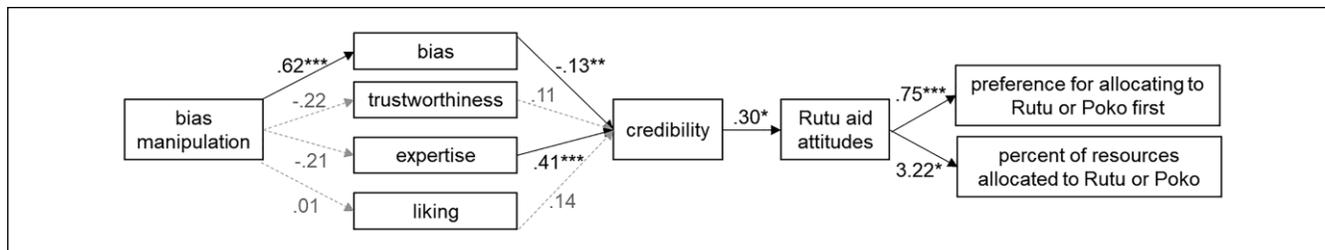
Figure 6. Parallel mediation analysis testing the effect of the bias manipulation on the resource allocation decisions through perceptions of bias, trustworthiness, expertise, and likeability to credibility, then attitudes, then allocation beliefs in Study 4.
*p < .05. **p < .01. ***p < .001.

sus Poko. Because this item was not highly related to the items asking where they wanted to allocate aid first (*r* = .27, *p* < .001), we treated this as a separate outcome.

Results

Manipulation checks. The bias manipulation had a significant effect on perceptions of bias with the objective source being seen as less biased (*M* = 4.18, *SD* = 1.83) than the biased source (*M* = 5.42, *SD* = 1.86), *t*(167) = 4.37, *p* < .001, 95% CI = [0.68, 1.80], *d* = .68.

Effect of bias on credibility. The bias manipulation had a significant effect on perceived source credibility, with participants viewing the source as more credible in the unbiased condition (*M* = 6.31, *SD* = 1.49) compared with the biased condition (*M* = 5.81, *SD* = 1.43), *t*(167) = −2.20, *p* = .03, 95% CI = [−0.93, −0.05], *d* = −.34.

Effect of bias on attitudes mediated through credibility. The primary purpose of this study was to again examine whether bias would have a negative effect on credibility, which would then lead to reduced persuasion. As such, we tested a mediation model in which bias reduced perceptions of credibility, which then reduced the extent to which participants thought it was a good idea to send resources to Rutu. Because it is possible that our manipulation influenced more than one source perception, in this mediation model as in the previous studies, we entered bias, trustworthiness, expertise, and liking as parallel mediators to examine whether our manipulation was primarily having its influence through perceptions of bias (Figure 6). First, we examined the influence of the manipulation on perceptions of bias, trustworthiness, expertise, and liking. The manipulation only significantly influenced perceptions of bias. Next, we regressed credibility on each of the perceptions and the bias manipulation. Only perceptions of bias and expertise significantly predicted perceptions of credibility.[8] Next, we regressed attitudes toward sending aid to Rutu on perceptions of credibility, bias, trustworthiness, expertise, and liking, and the bias manipulation and found that only perceptions of credibility were related to participants'

attitudes. We used 10,000 Monte Carlo simulations to test for parallel mediation through each of the source perceptions on attitudes. There was only a significant indirect effect through perceptions of source bias, 95% CI = [−0.060, −0.003], not through perceptions of source trustworthiness, 95% CI = [−0.027, 0.003], expertise, 95% CI = [−0.073, 0.002], or liking, 95% CI = [−0.012, 0.013].

We also extended beyond attitudes in this study to examine the use of those attitudes in related allocation decisions. Thus, we extended our mediation model to include each allocation outcome. We regressed participants' preferences for allocating resources to Rutu or Poko first on attitudes and all of the source perceptions, as well as the bias manipulation. Attitudes significantly predicted wanting to allocate resources to Rutu first. We conducted the same regression model with the percentage decision as the outcome. The more participants liked sending aid to Rutu, the larger the percentage of resources they allocated to Rutu. We used 10,000 Monte Carlo simulations to test for parallel mediation through each of the perceptions on both allocation outcomes. When we examined indirect effects on which place participants wanted to allocate resources first, there was a significant indirect effect through perceptions of source bias, 95% CI = [−0.046, −0.002], but not through perceptions of source trustworthiness, 95% CI = [−0.021, 0.001], expertise, 95% CI = [−0.054, 0.001], or liking, 95% CI = [−0.009, 0.009]. When we examined an indirect effect on the percentage of resources that participants wanted to allocate between the two places, there was a significant indirect effect through perceived source bias, 95% CI = [−0.232, −0.005], but not trustworthiness, 95% CI = [−0.096, 0.008], expertise, 95% CI = [−0.265, 0.005], or liking, 95% CI = [−0.037, 0.045]. Once again, these results suggest that the bias manipulation was operating through perceptions of source bias and not any of these other possible perceptions.

Discussion

This study demonstrated that people can infer source bias without also perceiving the source as dislikeable, untrustworthy, or inexpert. Moreover, in the absence of these other perceptions, perceiving a source as biased negatively affected

perceptions of source credibility, which had downstream consequences for how persuasive the source was.

General Discussion

We have argued that bias represents a distinct, but understudied source perception. Studies 1a and 1b empirically tested the conceptual distinction between bias and trustworthiness. Descriptions relevant to a source's honesty primarily affected perceptions of trustworthiness, but descriptions relevant to a person's motivation to hold a particular position primarily affected perceptions of bias. This distinction was shown to be important within the persuasion domain because bias had independent negative effects on source credibility (Studies 3a-c and 4) and persuasion (Studies 2, 3a-c, and 4) beyond the effects of trustworthiness, expertise, or liking. This suggests that source bias represents a third, understudied contributor to source credibility, a classic persuasion variable.

In our goal to distinguish bias from other potentially related variables, we began with these effects because we thought it was a fairly large oversight that social psychologists were unable to speak to the role of perceptions of bias in perceptions of credibility and persuasion. This is especially true in light of current conversations about the distinction between "biased" and "fake" news. Indeed, many textbooks and classic texts in psychology and related fields, such as communications (Cooper et al., 2016; Fiske & Dupree, 2014; Gass & Seiter, 2018; Hovland et al., 1953; McGuire, 1985; Petty & Cacioppo, 1981; Petty & Wegener, 1998) describe source credibility as the combination of trustworthiness and expertise. If practitioners were to use this criteria to select "credible" sources, the current work suggests that they might still employ sources that suffer from a lack of credibility if those sources are perceived as biased.

Implications

This article lays the groundwork for the distinction between bias and untrustworthiness by demonstrating that they have independent effects on perceived credibility and persuasion. However, we believe that there may be times when bias and untrustworthiness may have *differing* influences on some variables. For example, people might hold different expectations for how consistent sources will be in their position taking (Wallace, Wegener, & Petty, under review). Knowing that a source is biased might lead people to expect that the source will maintain the same position. However, simply knowing that a source is untrustworthy might not lead to these same expectations. These different expectations might have important implications for persuasion; previous research has suggested that when a source takes an unexpected position, it can increase persuasion (Eagly, Wood, & Chaiken, 1978). When combined with differential expectancies for biased versus untrustworthy sources, this would result in more positive expectancy-based influences on persuasion when a biased rather than an untrustworthy source changes positions. In a related vein, there may be circumstances in which receivers construe source bias positively, which might result in a positive effect on persuasion.

Bias and untrustworthiness may also have differing consequences for selective exposure following advocacy from a biased versus untrustworthy source. With each of these sources, recipients should feel that the information they received is lacking. However, the types of "problems" with the information is different for biased (but honest) sources versus dishonest (but objective) sources. In the case of biased sources, the information they present might only support one side of the issue, but at least people can treat the information as useful for understanding that side of the issue. Thus, accuracy-motivated recipients could seek out information that supports the opposite side from that which the biased source advocates to gain a more complete view of the issue. In contrast, simply knowing that a source is dishonest suggests that one cannot count on the accuracy of the information presented about either side of the issue, muting the directional information seeking that might occur when the source is viewed as biased rather than dishonest.

The distinction between bias and trustworthiness could also be relevant to person perception more generally. For example, traditional models of person perception have focused on perceptions of competence and warmth (with warmth sometimes divided into morality and sociability; e.g., Cuddy, Fiske, & Glick, 2008; Leach, Ellemers, & Barreto, 2007). The items used to capture morality tap directly into perceptions of trustworthiness (e.g., honest, sincere) but they do not generally include perceptions of bias per se. However, recent research has identified that people also categorize social groups along the dimension of conservative to progressive beliefs (Koch, Imhoff, Dotsch, Unkelbach, & Alves, 2016), which seems similar to perceptions of one type of bias. Perceptions of bias might be particularly important in an intergroup domain in which individuals would be concerned that others would have biases against their group separate from concerns that they might be untrustworthy. To the extent that perceptions of bias can be separated from liking, trustworthiness, and expertise (competence) of message sources (see also Wallace, 2015), it may be important for person perception research more generally to consider bias separately from morality (trustworthiness), sociability, or competence.

Relationships represent an additional domain in which the distinction between bias and trustworthiness might be important. It is relatively easy to imagine that people would desire honesty in a relationship partner. However, it is less clear that they would prefer objectivity in a relationship partner; in fact, it seems that relationships often thrive when partners are favorabley biased in their views of one another (Murray, Holmes, & Griffin, 1996).

Relatedly, the interpersonal consequences of admitting (at least some types of) bias versus dishonesty could be quite different. As just one example, consider the recent psychological research methods discussion. Srivastava (2018) argued that it is better to discuss concerns about researchers' data analyses as the result of unintentional biases rather than "planned malfeasance." This seems related to the distinction that we are drawing here. Questionable research practices could occur because of skewed perceptions (bias) or dishonesty (untrustworthiness) on the part of the researcher. Srivastava suggests that the assumptions others make about why QRPs occurred and whether they frame these concerns as about bias or untrustworthiness should have implications for how the authors receive the concerns and the authors' reputations. The distinction between bias and untrustworthiness may inform how to have more productive conversations about research practices, in additional to applying more broadly to giving feedback about others' behavior.

Statistical Power

Because we were dealing with novel effects, we ensured that samples were somewhat larger than in most similar persuasion studies. Furthermore, when possible, we used measures validated in similar research, and we sought effects that replicated across a number of studies (thereby lessening concerns about individual study power, as the consistent pattern of results across studies makes it less likely that we are dealing with false positive results; see Fabrigar & Wegener, 2016). Therefore, we do not believe that there should be concerns about potential lack of power for tests of the primary effects in the presented studies, especially independent bias effects on credibility.

Notes

1. One key exception to the general tendency overlook source bias is the work by Eagly, Wood, and Chaiken (1978). This research examined perceptions of source knowledge bias (similar to the bias concept we discuss) and reporting bias (linked directly to trustworthiness). However, this work did not result in routinely distinguishing between bias and trustworthiness, perhaps because the point of the Eagly et al. (1978) research was that violating expectations can enhance persuasion regardless of whether expectancy confirmation implied bias or untrustworthiness.

2. It may seem intuitive that bias and expertise refer to distinct constructs. Expertise refers to the amount of knowledge a person has on the topic. People might assume that individuals with high levels of knowledge might be less likely to have biased knowledge. However, short of possessing all knowledge about a topic, the amount of knowledge seems clearly separable from potential slants in that knowledge. In this article, we discuss biases that might stem from motivations, but it is also possible that biases could stem from slants in knowledge (Eagly et al., 1978; Petty & Cacioppo, 1986).

3. We could have crossed motivation and honesty in a 2 (motivation: high vs. low) $\times$ 2 (honesty: high vs. low) design. However, we wanted to initially focus on the particular elements that underlie bias or trustworthiness without information about the other components of credibility, as we thought this would provide the purest test of these conceptualizations. This focus parallels some previous studies of source credibility that have generally manipulated only one component without information about other traits (e.g., Briñol, Petty, & Tormala, 2004; Hovland & Mandell, 1952; Pratkanis, Greenwald, Leippe, & Baumgardner, 1988; Priester & Petty, 1995; Tormala, Briñol, & Petty, 2006).

4. To assess what the correlation would be in the absence of measurement error, we used the correction for attenuation formula, in which the correlation between two variables is divided by the square root of the product of their reliabilities. Based on this formula, the correlation between the bias and trustworthiness latent variables was $r = .55$ in Study 1a and $r = .37$ in Study 1b.

5. In Studies 2-3c, we describe the biased source as "one-sided" to indicate that they view the attitude object as more extreme than it actually is. Biases might commonly lead people to one-sided views. For example, an ideological bias might lead someone to view a political candidate for their party as exclusively positive, when in fact that candidate possesses both positive and negative qualities. However, it is also possible that an ideological bias would lead someone to view a candidate for their political party as relatively neutral when in fact that candidate possesses very few positive and many negative qualities. In this case, the recipient would be relatively two-sided in their view, yet still biased. Whether one-sidedness corresponds to bias depends on whether the attitude-object possesses relatively mixed qualities.

6. In a regression with the argument quality manipulation, the bias manipulation, their interaction, and the trustworthiness and expertise measures predicting thoughts, argument quality had a main effect on thoughts (i.e., more favorable thoughts with strong than

ambiguous arguments), $p = .001$, but did not interact with the bias manipulation, $p = .42$. In the same regression model predicting attitudes, argument quality had a main effect on attitudes, $p < .001$, but did not interact with the bias manipulation, $p = .53$.

7. Including "study" as a factor did not moderate any of the paths in the mediation model from bias --> credibility --> thoughts --> attitudes, ps $> .13$.

8. Despite the lack of significant effect of trustworthiness on credibility, the effect was in the expected direction and reasonably close to significant ($p = .11$). In addition, if source liking is removed from the model to just examine effects of trustworthiness, bias, and expertise on credibility, the relation between trustworthiness and credibility reaches significance, $b = .14$, $t(168) = 2.09$, $p = .04$.

Supplemental Material

Supplemental material for this article is available online.

ORCID iD

Laura E. Wallace [iD] https://orcid.org/0000-0002-2298-7569

References

Baker, S. M., & Petty, R. E. (1994). Majority and minority influence: Source-position imbalance as a determinant of message scrutiny. *Journal of Personality and Social Psychology*, *67*, 5-19.

Briñol, P., Petty, R. E., & Tormala, Z. L. (2004). Self-validation of cognitive responses to advertisements. *Journal of Consumer Research*, *30*, 559-573.

Chaiken, S., Liberman, A., & Eagly, A. H. (1989). Heuristic and systematic information processing within and beyond the persuasion context. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 212-252). New York, NY: Guilford Press.

Chaiken, S., & Maheswaran, D. (1994). Heuristic processing can bias systematic processing: Effects of source credibility, argument ambiguity, and task importance on attitude judgment. *Journal of Personality and Social Psychology*, *66*, 460-473.

Cooper, J., Blackman, S. J., & Keller, K. (2016). *The science of attitudes*. New York, NY: Routledge.

Correll, J., Park, B., Judd, C. M., & Wittenbrink, B. (2002). The police officer's dilemma: Using ethnicity to disambiguate potentially threatening individuals. *Journal of Personality and Social Psychology*, *83*, 1314-1329.

Cuddy, A. J. C., Fiske, S. T., & Glick, P. (2008). Warmth and competence as universal dimensions of social perception: The stereotype content model and the BIAS Map. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 40, pp. 61-149). San Diego, CA: Elsevier Academic Press.

Eagly, A. H., Wood, W., & Chaiken, S. (1978). Causal inferences about communicators and their effect on opinion change. *Journal of Personality and Social Psychology*, *36*, 424-435.

Fabrigar, L. R., & Wegener, D. T. (2016). Conceptualizing and evaluating the replication of research results. *Journal of Experimental Social Psychology*, *66*, 68-80.

Fiske, S. T., & Dupree, C. (2014). Gaining trust as well as respect in communicating to motivated audiences about science topics.

Proceedings of the National Academy of Sciences, 111(Suppl. 4), 13593-13597.

Gass, R. H., & Seiter, J. S. (2018). *Persuasion, social influence, and compliance gaining* (6th ed.). New York, NY: Routledge.

Hayes, A. F. (2013). *Introduction to mediation, moderation, and conditional process analysis: A regression-based approach*. New York, NY: Guilford.

Hovland, C. I., Janis, I. L., & Kelley, H. H. (1953). *Communication and persuasion: Psychological studies of opinion change*. New Haven, CT: Yale University Press.

Hovland, C. I., & Mandell, W. (1952). An experimental comparison of conclusion-drawing by the communicator and by the audience. *The Journal of Abnormal and Social Psychology*, *47*, 581-588.

Kennedy, K. A., & Pronin, E. (2008). When disagreement gets ugly: Perceptions of bias and the escalation of conflict. *Personality and Social Psychology Bulletin*, *34*, 833-848.

Koch, A., Imhoff, R., Dotsch, R., Unkelbach, C., & Alves, H. (2016). The ABC of stereotypes about groups: Agency/socio-economic success, conservative-progressive beliefs, and communion. *Journal of Personality and Social Psychology*, *110*, 675-709.

Koeske, G. F., & Crano, W. D. (1968). The effect of congruous and incongruous source-statement combinations upon the judged credibility of a communication. *Journal of Experimental Social Psychology*, *4*, 384-399.

Kruglanski, A. W. (1989). *Lay epistemics and human knowledge: Cognitive and motivational bases*. New York, NY: Springer.

Leach, C. W., Ellemers, N., & Barreto, M. (2007). Group virtue: The importance of morality (vs. competence and sociability) in the positive evaluation of in-groups. *Journal of Personality and Social Psychology*, *93*, 234-249.

McGuire, W. (1985). Attitudes and attitude change. In G. Lindzey & E. Aronson (Eds.), *Handbook of social psychology* (3rd ed., pp. 233-346). New York, NY: Random House.

Murray, S. L., Holmes, J. G., & Griffin, D. W. (1996). The benefits of positive illusions: Idealization and the construction of satisfaction in close relationships. *Journal of Personality and Social Psychology*, *70*(1), 79-98.

Petty, R. E., & Cacioppo, J. T. (1981). *Attitudes and persuasion : Classic and contemporary approaches*. Dubuque, IA: Brown.

Petty, R. E., & Cacioppo, J. T. (1986). *Communication and persuasion: Central and peripheral routes to attitude change*. New York, NY: Springer.

Petty, R. E., Schumann, D. W., Richman, S. A., & Strathman, A. J. (1993). Positive mood and persuasion: Different roles for affect under high- and low-elaboration conditions. *Journal of Personality and Social Psychology*, *64*, 5-20.

Petty, R. E., & Wegener, D. T. (1998). Attitude change: Multiple roles for persuasion variables. In D. Gilbert, S. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (4th ed., pp. 323-390). New York, NY: McGraw-Hill.

Pratkanis, A. R., Greenwald, A. G., Leippe, M. R., & Baumgardner, M. H. (1988). In search of reliable persuasion effects: III. The sleeper effect is dead. Long live the sleeper effect. *Journal of Personality and Social Psychology*, *54*, 203-218.

Priester, J. R., & Petty, R. E. (1995). Source attributions and persuasion: Perceived honesty as a determinant of message scrutiny. *Personality and Social Psychology Bulletin*, *21*, 637-654.

Wallace, L.E. (2015). *Antecedents and consequences of perceiving a source as biased* (Unpublished master's thesis). Ohio State University, Columbus, Ohio.

Wallace, L.E. (2019). *Distinguishing perceptions of bias from perceptions of untrustworthiness: Independent perceptions with shared as well as unshared consequences and antecedents* (Unpublished doctoral dissertation). Ohio State University, Columbus, Ohio.

Wallace, L.E., Wegener, D.T., & Petty, R.E. (under review). Influences of Source Bias that Differ from Source Untrustworthiness: When Flip-Flopping is More and Less Surprising.

Srivastava, S. (2018). What if we talked about p-hacking the way we talk about experimenter effects? *The Hardest Science.* Retrieved from https://thehardestscience.com/2018/10/19/what-if-we-talked-about-p-hacking-the-way-we-talk-about-experimenter-effects/

Tormala, Z. L., Briñol, P., & Petty, R. E. (2006). When credibility attacks: The reverse impact of source credibility on persuasion. *Journal of Experimental Social Psychology*, *42*, 684-691.

Walster, E., Aronson, E., & Abrahams, D. (1966). On increasing the persuasiveness of a low prestige communicator. *Journal of Experimental Social Psychology*, *2*, 325-342.

Wegener, D. T., Downing, J., Krosnick, J. A., & Petty, R. E. (1995). Measures and manipulations of strength-related properties of attitudes: Current practice and future directions. In R. E. Petty & J. A. Krosnick (Eds.), *Attitude strength: Antecedents and consequences* (pp. 455-488). Mahwah, NJ: Lawrence Erlbaum.

Westfall, J., & Yarkoni, T. (2016). Statistically controlling for confounding constructs is harder than you think. *PLoS ONE*, *11*(3), e0152719.

Ziegler, R., & Diehl, M. (2011). Mood and multiple source characteristics: Mood congruency of source consensus status and source trustworthiness as determinants of message scrutiny. *Personality and Social Psychology Bulletin*, *37*, 1016-1030.