Journal of Experimental Social Psychology 47 (2011) 1207–1213

Contents lists available at ScienceDirect

Journal of Experimental Social Psychology

journal homepage: www.elsevier.com/locate/jesp

Reports

No good deed goes unquestioned: Cynical reconstruals maintain belief in the power of self-interest☆

Clayton R. Critcher [a,*], David Dunning [b]

[a] Haas School of Business, University of California, Berkeley, USA
[b] Department of Psychology, Cornell University, USA

ARTICLE INFO

Article history:
Received 1 September 2009
Revised 28 April 2011
Available online 8 May 2011

Keywords:
Self-interest
Attribution
Belief perseverance
Person perception
Cynicism
Charitable giving

ABSTRACT

In four studies, we examined how people maintain beliefs that self-interest is a strong determinant of behavior, even in the face of disconfirming evidence. People reflecting on selfless behavior tend to reconstrue it in terms of self-interested motives, but do not similarly scrutinize selfish behaviors for selfless motives. Study 1 found that people react to new information that selfless behavior is common by interpreting it as more reflective of self-interest. Studies 2a and 2b, applying a Bayesian analysis, demonstrated that people see "too much" self-interest in seemingly selfless actions, given their prior beliefs, but see the predicted amount of self-interest in seemingly selfish actions. This demonstrates that people do not possess internally consistent belief systems, but rather undue cynicism. In Study 3, participants read about real philanthropists whose acts of generosity had been heralded by major news outlets. As participants spent more time considering why such philanthropy was performed, they formed more cynical impressions of the philanthropists' motives. Beyond offering insight into why belief in the norm of self-interest persists, these studies introduce a novel route by which beliefs resist disconfirmation.

Social perceivers expect others to be guided by self-interest (Miller, 1999), thinking, for example, that others are less troubled by lying (Kintz, 1977) or cheating on their tax returns (Wenzel, 2005) than are the social perceivers themselves. People incorrectly emphasize how much others' attitudes toward abortion policy or tobacco control will be self-interestedly guided by gender or smoking status, respectively (Miller & Ratner, 1998). Furthermore, they overestimate the proportion of their peers who will abuse their trust in an economic game (Fetchenhauer & Dunning, 2009, 2010).

Given these substantially inaccurate perceptions, why do people not observe others' attitudes and behaviors and learn that their own beliefs are overly cynical? One possibility is that people are simply not exposed to disconfirming evidence. For example, Holmes, Miller, and Lerner (2002) demonstrated that the norm of self-interest inhibits people from engaging in acts that appear purely selfless (Ratner & Miller, 2001), which limits the amount of selfless behavior people observe. As a second possibility, a misanthropic memory bias may lead people to disproportionately remember a person's negative or selfish actions (Ybarra, Stephan, & Schaberg, 2000).

Herein, we discuss another route by which people can maintain a belief in the power of self-interest even when they see evidence that should challenge their beliefs. We propose that people maintain such a belief due to an asymmetry in how they respond to seemingly selfless versus selfish acts, an asymmetry we term *attributional cynicism*. After further thought, people view behaviors that at first blush seem selfless to be more selfish. For example, once people observe a college woman buying flowers as part of a fraternity charity drive, an initial inclination to praise her ("how charitable!") may give way to a more cynical interpretation ("what a desperate maneuver to get an invitation to tonight's fraternity party"). This reflects attributional cynicism.[1] In contrast, selfish behavior is taken at face value as selfishly motivated. If a woman refuses to buy flowers, for example, one does not see the social perceiver search for altruistic reasons for the refusal (e.g., "she must have even more important charities to give to!").

Thus, in re-construing a seemingly selfless act as more selfish, people tend to shift their beliefs about the action that self-interest would likely promote in the specific situation they are considering. Because of this shift, people should no longer be so confident that self-interest would discourage an actor from performing the seemingly selfless action, thereby protecting their belief in the general prevalence of self-interest. We propose, in a sense, that people

☆ This research was partially supported by a National Science Foundation Graduate Research Fellowship and a Yale University Mellon Research Grant, awarded to Critcher. We thank Richard Eibach for his selfless contributions to this project. We thank Jeremy Cone, Linda Mohammadian, and Joshua Tabak for their assistance with data collection.

* Corresponding author.
E-mail address: ClaytonCritcher@haas.berkeley.edu (C.R. Critcher).

[1] The term *cynicism* refers merely to the content of the attributions, not to the process driving the attributions.

0022-1031/$ – see front matter © 2011 Elsevier Inc. All rights reserved.
doi:10.1016/j.jesp.2011.05.001

protect their belief in the norm of self-interest by seeing "too much" self-interest in seemingly selfless acts.

The present research is related to, but ultimately makes qualitatively different claims from past research arguing that people have a specialized ability for detecting self-interest or ulterior motivations in others. Past work has shown that people overcome the ubiquitous confirmation bias (e.g., Wason, 1983) when they are evaluating whether people are following social rules or cheating in a self-interested way (Cosmides, 1989; Cosmides, Tooby, Fiddick, & Bryant, 2005). In a separate program of research, Fein(1996) and Fein, McCloskey, and Tomlinson(1997) have found that people are quick to become suspicious of the ulterior motives of others, which can help them to avoid otherwise robust biases in social perception. In a sense, past work has shown sensitivity and sophistication in the perception of selfishness. Here, we test whether this system may work on overdrive not only to reason well about selfish behaviors, but instead to view the world more generally through a cynical lens.

Four studies tested whether people have a cynical bias when considering seemingly selfless behavior. In Study 1, we gave participants different information about the prevalence of honesty, which at first should seem selfless. We predicted that participants would react to news that most people are honest by seeing honesty as more a product of self-interest. That is, they would revise what behavior they believe self-interest motivates, rather than reducing how much self-interest in general they thought people possessed. Studies 2a and 2b tested whether attributional cynicism occurs without the presentation of new information about the occurrence or prevalence of behaviors. Instead, the studies rely on Bayesian principles to assess whether participants see "too much" (or "too little") self-interest in seemingly selfless or selfish behaviors. In Study 3, participants learned of real-life philanthropists who had been featured in major media outlets. We hypothesized that contemplating the altruistic acts of these individuals would lead participants to become more cynical when explaining these generous actions.

Study 1

Study 1 tested whether people maintain a belief in the norm of self-interest by reading self-interest into seemingly selfless actions—thus maintaining a belief in the overall prevalence of self-interest in others. Some participants were asked to respond to the fact that most college students are academically honest. We predicted participants would respond to such knowledge by: (a) increasing their belief that self-interest will lead students toward academic honesty (i.e., p[honesty | self-interest]), thereby (b) maintaining their belief in the self-interestedness of the population they judged (i.e., p[self-interest]). Some other participants were told, instead, that a majority of students cheat. We predicted that this information would lead to no revision whatsoever in people's beliefs about self-interest.

Method

Participants and design

Participants were 26 Cornell University undergraduates. Participants in this and future studies received extra course credit.

Procedure

Participants learned of a purported sociological study of academic dishonesty in America's colleges and universities. Six bullet points carefully explained what did and did not constitute academic dishonesty. Participants in the majority honest [dishonest] condition were told that the study had been completed, and that it was found that a clear majority (69%) of students had been academically honest [dishonest] in the past thirty days. Participants in a control condition did not receive any information about the results of the survey.

Table 1
Estimates that students are self-interested and would act honestly if self-interested (Study 1).

	p(honesty | self-interest)	p(self-interest)
Control	31.6%$_a$	76.9%$_a$
Majority dishonest	30.8%$_a$	81.8%$_a$
Majority honest	61.9%$_b$	75.9%$_a$

Note. Means in the same column not sharing a subscript differ at the $p<.001$ level. Those sharing a subscript do not differ at the $p<.05$ level.

At this point, all participants made two estimates by providing a percentage from 0 to 100:p(self-interest), "What percentage of students made their decision mostly due to self-interested (as opposed to selfless) motives?"; and p(academic honesty | self-interest), "Of those students whose decision was mostly driven by self-interested motives (as opposed to selfless motives), what percentage of them—in your best estimate—were academically honest?"[2]

Results and discussion

To determine whether learning behavioral base rates affected participants' theories about how self-interest guided behavior (p[honesty | self-interest]), we first performed a one-way ANOVA, which confirmed that participants' theories differed by condition, $F(2, 23) = 6.17$, $p = .01$. The means are presented in Table 1. Those told that a majority of students were academically honest became more confident that self-interest would lead students to be honest ($M = 61.9\%$) compared to those who were told the opposite ($M = 30.4\%$,), $t(23) = 3.01$, $p = .01$, $d = 1.26$. Compared to the control condition ($M = 31.6\%$), those in the majority honest condition displayed a shift in their belief that self-interested students would behave honestly, $t(23) = 3.04$, $p = .01$, $d = 1.27$, whereas those in the majority dishonest condition did not show evidence of a shift, $t<1$. Such shifts in participants' beliefs seemed to protect participants' beliefs in the norm of self-interest. Receiving information about the base rate of academic honesty failed to affect estimates of the percentage of respondents believed to have acted out of self-interest ($Ms = 75.9\%$ and 81.8%, for majority honest and dishonest conditions, respectively; $t(15) = 1.05$, $p>.31$).

In sum, even with grossly different information about how academically dishonest survey respondents were, participants did not revise their belief about how prevalent the motive of self-interest was. Rather, what participants did alter were their beliefs about which behavior that self-interest was likely to cause.

Studies 2a and 2b

Do the results of Study 1 merely reflect how people engage in Bayesian updating in light of new information, or instead do people twist their construals of seemingly selfless behaviors merely upon further contemplation? Belief updating in light of new information could be a perfectly reasonable response to new information. That is, people may be normative social perceivers with a strong prior belief in the prevalence of self-interest. If so, it would make sense for them to revise their theory to think self-interest is behind high rates of honesty.

Studies 2a and 2b more precisely examine whether cynicism occurs even in the absence of new information. We accomplished this through a novel application of Bayes' Rule. More specifically, we proposed that an attribution can be expressed as a conditional probability—the likelihood that a behavior is driven by self-interested motives, p(self-interest | behavior X). According to Bayes' Rule, this estimate should be a function of three other beliefs: the likelihood that a person will be self-interested

[2] To avoid focusing participants on self-interest (versus selflessness), judgments in all studies were made about selflessness versus self-interest. For ease of presentation, we describe effects just referencing self-interest.

Table 2
The direct and extrapolated probability judgments by condition (Studies 2a and 2b).

	Study 2a		Study 2b	
	Frances (selfless)	Tyler (self–interested)	Frances (selfless)	Tyler (self–interested)
p(Frances)	56.77%	58.18%	56.67%	56.12%
p(self–interest)	53.45%	54.90%	57.98%	53.20%
p(Frances | self–interest)	30.31%	40.24%	30.53%	27.64%
p(self–interest | Frances)—direct	41.87%	---------	29.65%	---------
p(self–interest | Tyler)—direct	---------	78.38%	---------	82.13%

Note. Participants were not randomly assigned to condition until after having made the three judgments above the dotted line.

in deciding how to behave (*p*[self-interest]), the likelihood that the person will choose behavior X (*p*[behavior X]), and the likelihood that the person would choose behavior X if it were known that he or she would make such a decision out of self-interest (*p*[behavior X | self-interest]), such that:

$$p(\text{self}-\text{interest}|\textit{behavior X}) = p(\text{self}-\text{interest})^{*}p(\textit{behavior X}|\text{self}-\text{interest}) / p(\textit{behavior X}). \quad (1)$$

But would this normative relationship hold? Or would participants make a more cynical attribution for seemingly selfless behaviors than their other beliefs should allow for? By comparing participants' direct judgments of *p*(self-interest | selfless behavior) to one that can be extrapolated from the other elements of Bayes' Rule (the right side of Eq. (1)), we can assess whether participants were unjustifiably cynical given their prior beliefs. Thus, we could label each participant as overly cynical (direct estimate>extrapolated estimate), internally consistent or realistic (direct = extrapolated), or overly hopeful (direct<extrapolated). Because we provided no additional information about the target's actual behavior, there should be no systematic deviations between participants' direct attributions and those that can be extrapolated from Bayes' Rule.

Method

Participants and design

One hundred forty-six Yale and Cornell University students participated in Study 2a. Two hundred ninety-four Cornell University students participated in Study 2b.

Procedure

Participants read a short fictional story about a college student named Emma (see Appendix A). Emma had been invited by an acquaintance, Frances, to her family's vacation home on Martha's Vineyard for Labor Day weekend. Emma had accepted the offer. However, soon after, Emma's longtime crush Tyler invited Emma to be his date to a fraternity party the night Emma was scheduled to be vacationing with Frances. Thus, Emma faced a dilemma between a seemingly selfless action (honoring her commitment to Frances) and a seemingly selfish one (going with Tyler). Participants in Study 2a first provided three judgments: how likely it was that Emma would go with Frances versus Tyler; how likely it was that Emma's choice would be driven by selflessness or self-interest; and if Emma's choice were driven by self-interest, how likely it would be that she would go with Frances versus Tyler. Participants in 2b answered the same questions, but the more nebulous "self-interest" was replaced by "selfish motives, a concern primarily with her own benefits, interests, welfare, etc".

Finally, participants in both studies answered whether a decision by Emma to go with Frances or Tyler (manipulated between-participants) would have been driven by self-interest (or "selfish motives, a concern…") or selflessness (or "selfless motives"). We compared this direct judgment of self-interest to the judgment that could be extrapolated (using Bayes' Rule) from the three initial probability judgments. Crucially, we stressed that in responding to the conditional probability judgments (both *p*[self-interest | behavior X) and *p*[behavior X | self-interest]), participants did not actually know what Emma chose or what motivated her decision. In this way, participants' set of beliefs should not systematically depart from internal consistency. We of course expected few participants to display perfect internal consistency in their judgments, but bias can be detected by a systematic, asymmetric departure from Bayes' Rule.

Results and discussion

The average response to each individual judgment is listed in Table 2. Confirming our assumption that going on the date with Tyler would be considered the more selfish course of action, *p*(Frances | self-interest) was less than 50% for both Study 2a, $t(145) = -6.01$, $p < .001$; and Study 2b, $t(293) = -13.23$, $p < .001$. We used Bayes' Rule to compute each participant's extrapolated judgment of self-interest given the hypothetical outcome they judged. Because extrapolated judgments have no upper bound, whereas direct judgments have an upper bound of 100, we recoded extrapolated judgments greater than 100 as 100.

We then compared participants' direct attributions of self-interest, *p*(self-interest | behavior X), to the one extrapolated from their pre-existing beliefs. Extrapolated judgments were not normally distributed, so we used non-parametric analyses. If the direct attribution was greater than the extrapolated one, we labeled them as a "cynic." If the two beliefs were equal, they were labeled "realists," and if their direct attribution was less than their extrapolated estimate, they were labeled as "hopefuls." The percentage of cynics, realists, and hopefuls, by condition, is listed in Table 3. We of course do not expect all (or even many) participants to be perfectly internally consistent, but we can test for systematic bias by comparing the number of cynics and hopefuls. Thus, the tests below exclude realists.

Study 2a

We conducted a logistic regression to assess whether the percentage of cynics versus hopefuls differed by condition. And indeed, a significant interaction showed that those who considered the seemingly selfless act (going with Frances) were more likely to display cynicism (versus hopefulness) than were those who considered the seemingly self-interested action (going with Tyler), χ^2 (1, N = 122) = 6.55, $p = .01$. There were significantly more cynics (69%) than hopefuls (31%) in judging the seemingly selfless act, χ^2 (1, N = 67) = 9.33, $p = .002$, but just as many cynics (45%) as hopefuls (55%) in judging a seemingly selfish decision to go with Tyler, $\chi^2 < 1$.[3] In other words, people perceived more self-interest in the seemingly selfless decision than their prior beliefs permitted, whereas their attributions for a seemingly self-interested decision followed consistently from their prior beliefs.

[3] Note that because we excluded realists when performing these tests, the percentage of cynics and hopefuls differs from the overall percentages offered in Table 2.

Table 3
Percentage of participants who are cynics, realists, or hopefuls, by condition (Studies 2a and 2b).

	Study 2a		Study 2b	
	Frances (selfless)	Tyler (self-interested)	Frances (selfless)	Tyler (self-interested)
Cynics	60	37	51	33
Realists	13	18	14	26
Hopefuls	27	44	35	40
χ^2 (1)	9.33**	.46	4.65*	1.11
Extrapolated attribution	30.0%	80.0%	20.0%	85.0%
Amount of cynicism	+13.3%	−2.9%	+5.0%	−5.0%

Note. The chi-squared statistic comes from a test of whether the number of cynics and hopefuls differ. Cynics are participants for whom their direct judgment of self-interest is greater than their extrapolated judgment. Realists are those for whom their direct judgment equals their extrapolated judgment. Hopefuls are those for whom their direct judgment is less than their extrapolated judgment. The extrapolated attribution is the median extrapolated attribution of self-interest for those participants included in the chi-squared test. The amount of cynicism is the median inflation of the direct judgment compared to the extrapolated judgment (i.e., direct attribution–extrapolated attribution) for those participants included in the chi-squared test.
** $p<.01$.
* $p<.05$.

Study 2b

Using the more precise definition of self-interest, a logistic regression again found that participants were more likely to be cynical (versus hopeful) in judging a seemingly selfless act than a seemingly self-interested act, χ^2 (1, N = 233) = 5.01, $p = .03$. There were significantly more cynics (60%) than hopefuls (40%) in judging a seemingly selfless action, χ^2 (1, N = 124) = 4.65, $p = .03$, but just as many cynics (45%) as hopefuls (55%) in assessing the seemingly selfish act, χ^2 (1, N = 109) = 1.11, $p > .29$. Thus, merely upon considering why a selfless act might occur, people assume more self-interested and selfish motives than their prior beliefs predicted. Attributional cynicism did not emerge simply because of new information (as in Study 1), but reflected the cynical spin placed on seeming selfless behavior when considering why it might occur.

In addition, given the novelty of using Bayes' Rule as an idiographic standard for attributions, it is important that participants did not systematically deviate from Bayes' predictions when the bias was not expected, namely, when they were judging an act they already construed as selfish. Furthermore, the validity of our application of Bayes' Rule relies on participants interpreting the term "self-interest" in a consistent way across their judgments. Study 2b defined "self-interest" in a more precise way to assure such consistency, and evidence of attributional cynicism remained.

This latter finding in Study 2b is crucial, for the validity of Bayes' Rule relies on participants defining *self-interest* in a consistent way across their judgments. One might argue that participants in Studies 1 and 2a saw more self-interest after contemplating selfless behavior because they broadened their definition of the term to be more inclusive when considering a seemingly selfless act—to something akin to standard economic theory, which asserts that all freely chosen behavior is self-interested in order to maximize the person's utility (Becker, 1976; Swedborg, 1990). Thus, any observed cynicism would not reflect a shift in one's attribution for the act, but in one's definition of self-interest itself. However, in Study 2b, even once self-interest was more precisely defined as "selfish motives, a concern primarily with her own benefits, interests, welfare, etc.", evidence of attributional cynicism remained.

We wish to further stress the significance of Study 2's extension of Study 1. Psychologists are well aware of people's remarkable talent at accommodating new information into their belief systems, and it would be not be terribly novel if this research were merely an additional demonstration of that. But our claim is stronger. Studies 2a and 2b demonstrate that attributional cynicism does not merely emerge in an attempt to protect one's strong belief in the norm of self-interest, nor is it compatible with an argument that people may "preemptively" shape their beliefs to make certain that no data can challenge them.

In other words, up front, someone with a strong belief in the power of self-interest should be able to construe any possible future behavior as self-interested. Someone using this strategy would have judged p(self-interest) to be high, and p(behavior X) as fairly close to p(behavior X | self-interest). This is because, according to such a belief system, knowing that a person will behave in a self-interested manner does not provide information. And this social perceiver would not be identified as a cynic by our methods; they would have set up their belief system in a way to accommodate most any behavior. We imagine that many academic economists fit this description. Instead, we suggest that as one moves from considering a choice between a seemingly selfish and selfless act to considering what motivates a particular selfless action, the very construal of this act shifts. Thus, combining across these steps, people start with the sense that the norm of self-interest helpfully distinguishes likely from unlikely behavior. But these boundaries are blurred after further thought, and such blurring does not discredit the belief in self-interest; it maintains it.

Study 3

Study 3 built on the prior studies in two key ways. First, participants judged real-world behaviors that had been featured in the news media because of the extraordinary selflessness these behaviors reflected. Observing cynicism in reaction to these stories would provide evidence for the robustness of attributional cynicism. Second, we aimed to more conclusively test whether it was extra thought, not simply belief updating in light of new information, that produces cynicism. Thus, all participants learned of the occurrence of ten philanthropic acts, but participants spent extra time contemplating only five of these behaviors before going on to judge all ten behaviors. We expected that this added thought would prompt cynicism about the five people participants considered compared to those participants who did not engage in additional consideration about the same targets.

Method

Participants and design

Two hundred twenty-two Cornell University undergraduates participated for extra credit toward their grades in psychology and human development classes.

Materials

We constructed profiles of ten major philanthropists. Of the ten, five made their mark by donating large sums of money to various organizations; the other five donated their time or services. We selected the money philanthropists by randomly choosing five profiles from *Slate Magazine*'s 2009 "Slate 60," an annual list extolling the year's biggest money philanthropists (Stonesifer, 2010). We selected the time philanthropists by searching Google News for instances of the phrase "donated time" and selecting appropriate

philanthropists that had been featured in the media. For each philanthropist, we created a brief profile that included the philanthropist's picture, his or her source of wealth, and a three-sentence description of the person's philanthropy (see Fig. 1). So results would not be contaminated by a general bias for or against wealthy individuals, we included a fictitious source of wealth for each time philanthropist.

Procedure

Participants were first exposed to all ten profiles, one at a time, in a random order. Then, participants were told about an attributional thought task: They would be exposed to a random five of the ten profiles again. For these five philanthropists, participants would list up to six reasons why the philanthropist might have performed each act of philanthropy. Just as in real-life judgment contexts, the motivations for giving were not explicitly mentioned in the profiles. After all, person perception is largely an exercise in (reasoned) speculation. To be certain that participants were aware of this, we added, "Because you only have sparse information on which to base these theories, feel free to write down reasons that might be true, even if the information doesn't allow you to explicitly confirm it." *Money* participants completed the attributional thought task for the five money philanthropists, and *time* participants completed the attributional thought task for the five time philanthropists. By this point, no mention had been made of selfishness or selflessness.

Next, to assess perceptions of selflessness, participants saw all ten philanthropists with a brief phrase reminding participants of each target's philanthropic contribution. Participants ordered the 10 philanthropists from who they thought was most motivated by selflessness to who they felt was most motivated by selfishness. They then saw all ten profiles again in a new random order. Participants were to indicate to what extent pure selflessness was the likely motivator behind each particular target's act(s) of philanthropy by providing a response from 1 (pure selfishness) to 9 (pure selflessness).

Finally, participants returned to the reasons they had generated earlier. Participants coded each one for whether it reflected a self-interested or selfish motivation, a selfless motivation, or one that was not easily classifiable on this dimension.

Results and discussion

Although there is no standard like Bayes' Rule by which to identify how much cynicism should be present in participants' attributions, it is notable just how cynical participants' explanations for philanthropy were. The average participant generated more selfish reasons for the philanthropy ($M = 11.25$, $SD = 5.80$) than selfless reasons ($M = 9.70$, $SD = 4.86$), paired $t(199) = 2.55$, $p = .01$. A minority of reasons ($M = 3.10$, $SD = 3.24$) were not easily classifiable.

To determine whether attributional thought prompted participants to form more cynical impressions, we tested how the extra thought condition impacted rankings and ratings of self-interest. For the rankings, for each participant we calculated the average selfishness ranking given the money philanthropists and the time philanthropists. We then submitted these average rankings to a 2 (extra thought: time or money) × 2 (philanthropists: time or money) mixed-model ANOVA, with only the second variable measured within-subjects. The expected interaction emerged, $F(1, 220) = 16.36$, $p < .001$. This reflected that after giving more thought to why either money or time philanthropists donated their resources, participants' rankings shifted so that philanthropists in the category they focused on became, on average, 0.52 "rankings" more selfish.

One limitation of the rankings measure is that participants may have simply forgotten the information presented in the profiles of the five philanthropists for which they did not generate attributions. To address this possibility, it would be important to observe a similar interaction on participants' ratings of the each target's selflessness, which they made only after rereading each full profile. Consistent with the attributional cynicism hypothesis, the same 2 (extra thought) × 2 (philanthropists) interaction emerged on participants' average selfishness ratings of each target, $F(1, 216) = 19.69$, $p < .001$ (see Fig. 2). Those who had generated reasons why money philanthropists engaged in their acts of charity rated the money philanthropists as more selfish than those who did not,

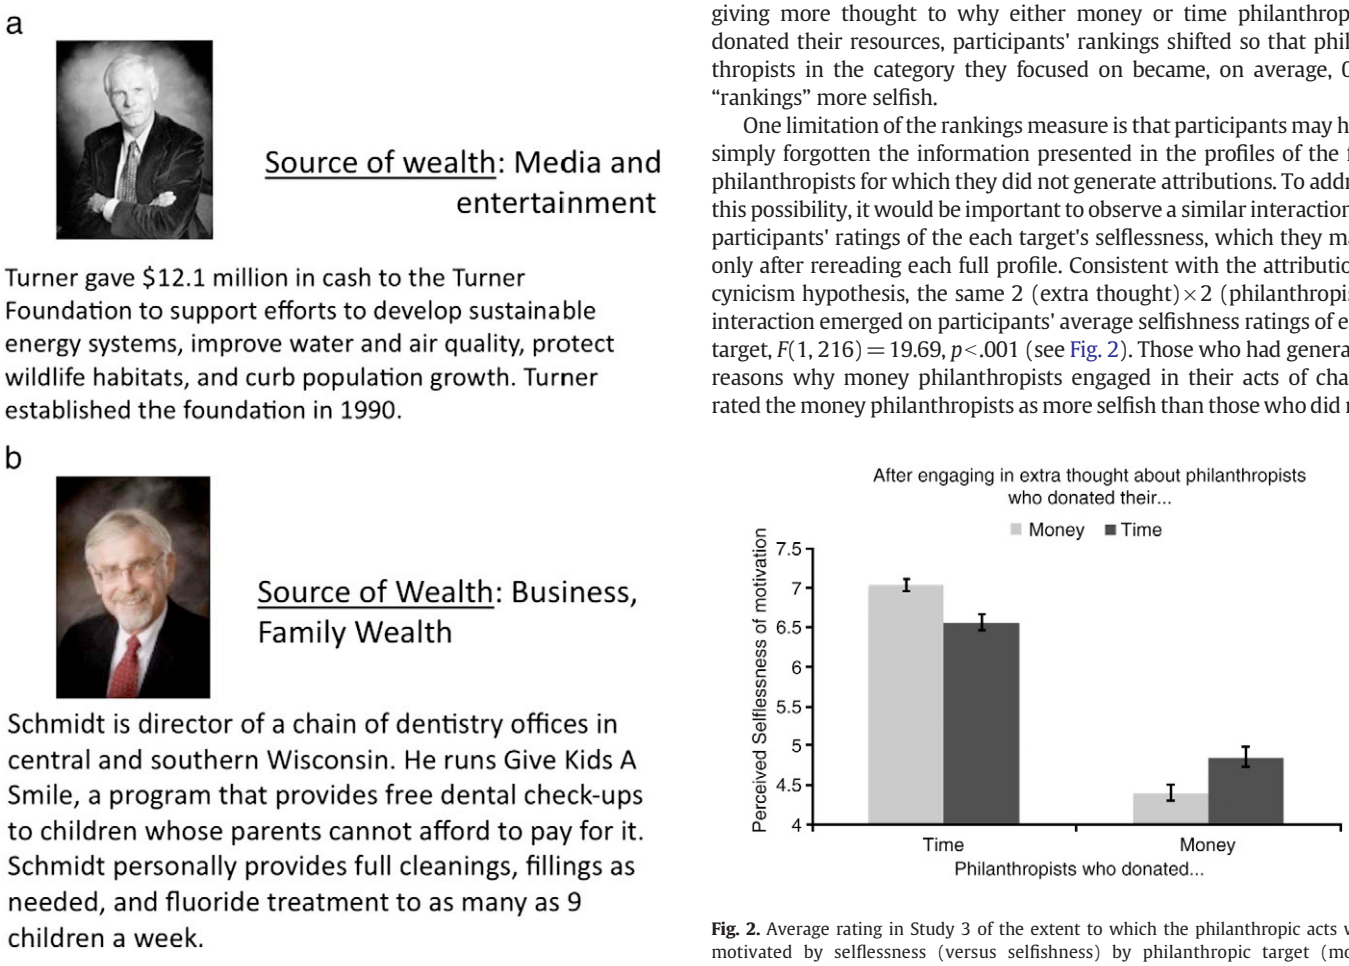

Fig. 1. Two example profiles (Study 3): (a) a money philanthropist, and (b) a time philanthropist.

Fig. 2. Average rating in Study 3 of the extent to which the philanthropic acts were motivated by selflessness (versus selfishness) by philanthropic target (money philanthropists vs. time philanthropists) and extra thought condition (money or time). Evidence of attributional cynicism is observed when the selflessness ratings for a target group decline after engaging in extra thought about that target group.

$t(208.10) = 2.85$, $p = .005$. Those who focused their attributional thought on the time philanthropists rated those philanthropists as more selfish than those who did not, $t(216) = 3.67$, $p < .001$.

Further buttressing our account that it was the attributions that participants generated that were responsible for the shift in participants' ratings and rankings, the greater the proportion of self-interested reasons that participants in the money [time] condition generated, the more they ranked those targets as selfish, $r(98) = .25$, $p = .01$ [$r(98) = .30$, $p = .002$], and rated the targets' motivations as selfish, $r(98) = .22$, $p = .03$ [$r(97) = .32$, $p = .001$].

These results provide the most straightforward support for our hypothesis that it is mere attributional thought (and not necessarily new information) that leads people to twist seemingly selfless behaviors into more selfish ones. Unlike in Studies 2a and 2b, we could not identify at the level of the individual whether attributional cynicism was displayed. Nonetheless, the differences in rankings and ratings at the group-level confirm that further thought about seemingly selfless acts is characterized by cynicism. In addition to providing support for our hypotheses, these results offer a cautionary tale to those who wish to heap well-intentioned praise upon altruists. The more people are prompted to consider the selflessness of others, the more cynical such perceivers may become.

General discussion

In this manuscript, we examined how people may retain their beliefs about the power of self-interest even in the face of what should be disconfirming evidence. We suggested that after being engaged in thought about why people act in an apparently selfish manner, people reconstrue that behavior more as a product of selfishness than they had before. No such shift would occur when people considered selfish behaviors.

Our studies provide evidence for this attributional cynicism. Upon learning that the prevalence of seemingly selfless behaviors is especially high, people do not revise their beliefs about the overall prevalence of selfishness in the world. Instead, they decide that seemingly selfless behaviors must be selfish after all (Study 1). Studies 2 and 3 moved beyond this initial finding to demonstrate that such revision is not simply a necessary consequence of receiving new empirical information. Instead, mere attributional thought prompted people to see more self-interest in a seemingly selfless act. Their attributions, in contrast, for a seemingly self-interested act do not systematically deviate from their prior beliefs in a more selfless direction (Studies 2a and 2b).Finally, Study 3 found that such cynicism is particularly robust. As participants thought more about leading philanthropists—people selected by media outlets for their unusually kind, generous acts—their thoughts became increasingly cynical. Notably, participants more frequently appealed to selfish than selfless motives when explaining the philanthropists' actions.

We have discussed bias by highlighting the way in which one's ultimate attributions did not follow from one's prior beliefs. But one could describe this effect in the opposite way, placing "bias" in one's prior beliefs—an unjustifiably definitive view of what self-interest predicts that is "corrected" with further thought. Regardless of where bias is placed (i.e., in the prior or posterior beliefs), what is clear is that the difference between prior and posterior beliefs indicates a counternormative dynamic process. For those cases in which participants received no new information, no Bayesian would endorse the systematic revisions in theory that participants made. This dynamic process leads people to maintain a belief in the predictive value of the self-interest construct, but then to fudge their interpretation of data in evaluating the theory's accuracy. Before considering a behavior, our participants believed that self-interest was a useful concept predicting which behaviors were more likely than others (e.g., dishonesty over honesty). But, after consideration, they shifted toward believing that self-interest could predict anything (both honesty and dishonesty). This left participants with "confirmed" views of the power of self-interest, regardless of the data (real or hypothetical). However, with faith in the power of self-interest affirmed, people can approach a new situation all-too-eager to once again rely on their confidence that seemingly selfish behavior (rather than seemingly selfless ones) will emerge (e.g., Miller & Ratner, 1998).[4]

These observations lead us to comment on psychological processes that lead to belief perseverance. The existing literature on belief perseverance has focused on ways in which people maintain false beliefs because they are not exposed to belief-disconfirming information (Holmes et al., 2002; Prentice & Miller, 1993), have poorer recall for belief-consistent information (Chapman & Chapman, 1969; Mischel, Ebbesen, & Zeiss, 1976), see more belief-consistent than belief-inconsistent points in a data display (Anderson, 1995), and hold belief-inconsistent data to a more stringent standard of validity (Lord, Ross, & Lepper, 1979; Mischel et al., 1976). Unlike in these past demonstrations, we have found that the belief-inconsistent nature of data is not a given, for the data themselves can be reconstrued in belief-consistent terms.

To illustrate the difference between these strategies, imagine a researcher who performs a number of studies to test a theory. Half of the studies produce hypothesis-consistent results; half directly counter her hypotheses. Past research has suggested that she may maintain belief in her theory because she misremembers that most of her studies were hypothesis-consistent; because her research assistants never show her the hypothesis-inconsistent results; or because she holds her hypothesis-inconsistent studies to a higher standard of rigor, finding they contain methodological flaws. In contrast, we suggest that she may "realize" that in actuality, her theory would have predicted those hypothesis-inconsistent results all along. Although we have focused on this process in the context of perceptions of the self-interestedness of others, we believe that this basic process can be used to account for many instances of belief perseverance.

Appendix A

Emma and Frances were acquaintances their freshman year of college. Upon returning to college sophomore year, they found to their surprise that they were living across the hall from each other. Frances, who tends to be shy and socially inhibited, hoped to find a friend in Emma. Frances decided to invite Emma to her family's lavish Martha's Vineyard shorefront vacation home for Labor Day weekend. Emma, who has never vacationed outside of her home state of New Jersey, enthusiastically accepted, and Frances immediately began making plans to ensure that the trip would be maximally conducive to bonding.

The day before Emma and Frances are set to leave for Frances' beach house, Emma is invited by Tyler—a guy she has had a crush on since the beginning of freshman year—to be his date to a party at his fraternity, the very night she is set to be at Martha's Vineyard with Frances! Knowing Tyler is highly sought after, she fears that by turning him down, the chances of his asking her out again are next to nothing. Also, she worries about the consequences of offending Frances, who she knows has been working hard to make their trip a success.

References

Anderson, C. A. (1995). Implicit personality theories and empirical data: Biased assimilation, belief perseverance and change, and covariation detection sensitivity. *Social Cognition*, *13*, 25–48.

[4] It is not the unfalsifiability alone of the belief system that makes the process counternormative. For example, an unfalsifiable belief system that all behavior is self-interested would produce internally consistent beliefs.

Becker, G. (1976). The economic approach to human behavior. In G. Becker (Ed.), *The economic approach to human behavior* (pp. 3–14). Chicago: The University of Chicago Press.

Chapman, L. J., & Chapman, J. P. (1969). Illusory correlation as an obstacle to the use of valid psychodiagnostic signs. *Journal of Abnormal Psychology*, *74*, 271–280.

Cosmides, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition*, *31*, 187–276.

Cosmides, L., Tooby, J., Fiddick, L., & Bryant, G. A. (2005). Detecting cheaters. *TRENDS in Cognitive Science*, *9*, 505–506.

Fein, S. (1996). Effects of suspicion on attributional thinking and the correspondence bias. *Journal of Personality and Social Psychology*, *70*, 1164–1184.

Fein, S., McCloskey, A. L., & Tomlinson, T. M. (1997). Can the jury disregard that information? The use of suspicion to reduce the prejudicial effects of pretrial publicity and inadmissible testimony. *Personality and Social Psychology Bulletin*, *23*, 1215–1226.

Fetchenhauer, D., & Dunning, D. (2009). Do people trust too much or too little? *Journal of Economic Psychology*, *30*, 263–276.

Fetchenhauer, D., & Dunning, D. (2010). Why so cynical? Asymmetric feedback underlies misguided skepticism in the trustworthiness of others. *Psychological Science*, *21*, 189–193.

Holmes, J. G., Miller, D. T., & Lerner, M. J. (2002). Committing altruism under the cloak of self-interest: The exchange fiction. *Journal of Experimental Social Psychology*, *38*, 144–151.

Kintz, B. L. (1977). College student attitudes about telling lies. *Bulletin of the Psychonomic Society*, *10*, 490–492.

Lord, C., Ross, L., & Lepper, M. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, *37*, 2098–2109.

Miller, D. (1999). The norm of self-interest. *The American Psychologist*, *54*, 1053–1060.

Miller, D. T., & Ratner, R. K. (1998). The disparity between the actual and assumed power of self-interest. *Journal of Personality and Social Psychology*, *74*, 53–62.

Mischel, W., Ebbesen, E. B., & Zeiss, A. M. (1976). Determinants of selective memory about the self. *Journal of Consulting and Clinical Psychology*, *44*, 92–103.

Prentice, D. A., & Miller, D. T. (1993). Pluralistic ignorance and alcohol use on campus: Some consequences of misperceiving the social norm. *Journal of Personality and Social Psychology*, *64*, 243–256.

Ratner, R., & Miller, D. (2001). The norm of self-interest and its effects on social action. *Journal of Personality and Social Psychology*, *81*, 5–16.

Stonesifer, P. (2010). The 2009 Slate 60: The largest American charitable contributions of the year [Electronic version]. *Slate* Retrieved April 21, 2011, from. http://www.slate.com/id/2243495 February 8.

Swedborg, R. (1990). *Economics and sociology*. Princeton, NJ: Princeton University Press.

Wason, P. C. (1983). Realism and rationality in the selection task. In J. St. B. T. Evans (Ed.), *Thinking and reasoning: Psychological approaches* (pp. 44–75). London: Routledge and Kegan Paul.

Wenzel, M. (2005). Misperceptions of social norms about tax compliance: From theory to intervention. *Journal of Economic Psychology*, *26*, 862–883.

Ybarra, O., Stephan, W. G., & Schaberg, L. (2000). Misanthropic memory for the behavior of group members. *Personality and Social Psychology Bulletin*, *26*, 1515–1525.