Journal of Experimental Psychology: General
2017, Vol. 146, No. 4, 512–528

0096-3445/17/$12.00 http://dx.doi.org/10.1037/xge0000276

The "Common Good" Phenomenon: Why Similarities Are Positive and Differences Are Negative

Hans Alves, Alex Koch, and Christian Unkelbach
University of Cologne

Positive attributes are more prevalent than negative attributes in the social environment. From this basic assumption, 2 implications that have been overlooked thus far: Positive compared with negative attributes are more likely to be shared by individuals, and people's shared attributes (similarities) are more positive than their unshared attributes (differences). Consequently, similarity-based comparisons should lead to more positive evaluations than difference-based comparisons. We formalized our probabilistic reasoning in a model and tested its predictions in a simulation and 8 experiments ($N = 1{,}181$). When participants generated traits about 2 target persons, positive compared with negative traits were more likely to be shared by the targets (Experiment 1a) and by other participants' targets (Experiment 1b). Conversely, searching for targets' shared traits resulted in more positive traits than searching for unshared traits (Experiments 2, 4a, and 4b). In addition, positive traits were more accessible than negative traits among shared traits but not among unshared traits (Experiment 3). Finally, shared traits were only more positive when positive traits were indeed prevalent (Experiments 5 and 6). The current framework has a number of implications for comparison processes and provides a new interpretation of well-known evaluative asymmetries such as intergroup bias and self-superiority effects.

Keywords: valence asymmetries, person perception, comparison processes, information sampling

Supplemental materials: http://dx.doi.org/10.1037/xge0000276.supp

The social world is predominantly positive. People typically behave according to social norms prescribing positive behaviors, whereas negative, norm-incongruent behavior is the exception. Consequently, people represent others with predominantly positive attributes (Boucher & Osgood, 1969; Matlin & Stang, 1978; Parducci, 1968; Rothbart & Park, 1986; Sears, 1983). From this positivity prevalence assumption follows an important implication. If people's attributes are predominantly positive rather than negative, positive compared with negative attributes are more likely to be shared among people. Conversely, people's shared attributes (i.e., their similarities) are more likely to be positive than their distinct attributes (i.e., their differences). In the most general sense, people's similarities are more positive than their differences. We termed this probabilistic necessity the "common good" phenomenon. Thus, any evaluative process that relies either on similarities or differences between persons, groups, or any stimuli should be affected in the following way: Similarity-based evaluations are positively biased, whereas difference-based evaluations are negatively biased. These biases result from the assumed positivity prevalence in the environment, and not from faulty or biased cognitive processes. And because similarities and differences are the building blocks of comparison processes, our model provides alternative explanations for a number of well-known evaluative asymmetries (e.g., intergroup biases, self-serving biases) and allows new predictions in the domain of (social) comparisons.

In the remainder of the article, we first discuss the present work's scope and its significance for social–cognitive phenomena and their explanations. We then review existing evidence for the positivity prevalence assumption. Next, we present an attribute distribution model which delineates why (a) positive attributes are more likely to be shared by individuals, and (b) shared attributes (similarities) are more likely to be positive than unshared attributes (differences). We then present data from a simulation and eight trait-sampling experiments that test our model's predictions. Finally, we elaborate on our model's theoretical implications for evaluative comparison processes in general.

Scope of the Present Work

We present and explain a general characteristic of the social world that we termed the "common good" phenomenon: People's common attributes are more positive than their unique attributes. Those attributes that connect different people and that define their similarities are usually good attributes. Those attributes that distinguish different people and make them unique are often bad attributes. As provocative as this claim sounds, it follows directly from statistical reasoning.

The common good phenomenon is certainly intriguing in its own right, but it also has substantial implications for social–cognitive phenomena and their explanations. Social psychologists are ultimately interested in people's evaluations of different per-

This article was published Online First February 20, 2017.

Hans Alves, Alex Koch, and Christian Unkelbach, Social Cognition Center Cologne, University of Cologne.

The research was supported by a grant from the Deutsche Forschungsgemeinschaft (UN 273/4–1) awarded to Christian Unkelbach.

Correspondence concerning this article should be addressed to Hans Alves, Department Psychologie, Universität zu Köln, Richard-Strauss-Strasse. 2, 50931 Köln, Germany. E-mail: hans.alves@uni-koeln.de

sons, groups, or objects, and in the underlying mechanisms that alter these evaluations. One characteristic of evaluations is that they are always relative and thereby involve comparisons (Kahneman & Miller, 1986)—that is, targets are always compared with some standards. Comparisons, on the other hand, can be performed in two fundamentally different ways: They can rely on the similarities or on the differences between targets and standards (Hodges, 2005; Houston & Sherman, 1995; Kruschke, 1996, 2001, 2003; Mussweiler, 2003; Tversky, 1977). We suggest that similarity-based and difference-based evaluations generally differ in valence. If the positivity prevalence assumption is correct, similarity-based evaluations are positively biased, whereas difference-based evaluations are negatively biased. Again, as we will show, these biases are probabilistic necessities and not result from faulty or motivated reasoning. Hence, when social perceivers evaluate targets based on how they are similar to others, targets will appear more positive compared with when perceivers evaluate targets based on how they are different from others. This provides a new perspective on social perception. For example, consider the accepted notion in social, as well as folk, psychology that reminding people of their similarities increases interpersonal attraction and thereby helps to overcome social conflict. The current work suggests that this may simply be a probabilistic necessity.

Furthermore, if similarity-based and difference-based evaluations have different evaluative implications, this might renew our understanding of various well-known evaluative asymmetries in social psychology. For example, people evaluate in-group and majority-group members more favorably than out-group and minority-group members (Hewstone, Rubin, & Willis, 2002; Katz & Braly, 1935; Thalhammer, Zucha, Enzenhofer, Salfinger, & Ogris, 2001), they evaluate themselves more favorably than other people and the average person (Alicke & Govorun, 2005; Hoorens, 1993; Taylor & Brown, 1988), and they evaluate their current self more favorably than their past self (Wilson & Ross, 2000, 2001). Such effects are often explained by motivated or distorted information processing (Beauregard & Dunning, 1998; Guenther & Alicke, 2010; Tajfel & Turner, 1979). Once we have introduced our model, we will illustrate that these effects may arise simply because perceivers differentiate persons, groups, and objects based on their unique features in a predominantly positive social ecology.

More generally, the current work shows that such evaluative biases may be "innocent" in the sense that they do not result from faulty psychological processing, but probabilistically result from a given information ecology (Brunswik, 1956; Fiedler, 1996, 2000, 2014; Fiedler & Juslin, 2006; Gibson, 1979; Lewin, 1951). This also implies that the same process leads to opposite outcomes if the ecology changes, for example, when an ecological positivity prevalence is experimentally changed to negativity prevalence (see Experiments 5 and 6). However, to introduce our model in detail, we start with its central assumption that the social information environment is predominantly positive.

Positive Is Prevalent

In a world governed by prosocial norms, positive behavior and positive attributes occur more frequently (e.g., Clark & Clark, 1977). Most people we encounter every day behave according to social norms, and it is rare that we see somebody engaging in deviant behavior. Simply spoken, most people show behavior that that is nice, honest, and caring most of the time, rather than mean, deceptive, and cruel.

This positivity prevalence follows because positive behavior is usually reinforced, whereas negative behavior is sanctioned and people seek positive reinforcements (Thorndike, 1898). People also seek positive social encounters and avoid negative ones, and thereby create their own positive social environment (Denrell, 2005; Fazio, Eiser, & Shook, 2004; Walker, Skowronski, & Thompson, 2003).

People's mental representation of their social world mirrors this prevalence of positive behaviors. People generally show a strong tendency to evaluate others positively (Greenberg, Saxe, & Bar-Tal, 1978; Perlman & Oskamp, 1971; Rothbart & Park, 1986), and they expect others to behave positively in standard interactions (Anderson, 1981; Sears, 1983). People also use positive words more frequently than negative words in written and spoken language across different languages and cultures (Augustine, Mehl, & Larsen, 2011; Boucher & Osgood, 1969; Dodds et al., 2015; Zajonc, 1968). In particular, people use positive compared with negative person description words more frequently (Ric, Alexopoulos, Muller, & Aubé, 2013). And finally, people's subjective states also mirror the positivity prevalence, as most people seem to be happy most of the time in general (Diener & Diener, 1996), with some variation across cultures (Biswas-Diener, Vittersø, & Diener, 2005).

We therefore assume that people's social environment, as well as their mental representation of the same, is predominantly positive, meaning that positive attributes are more frequent. As we will show it follows statistically that positive compared with negative attributes are more likely to be shared by individuals, and that shared attributes (similarities) compared with unshared attributes (differences) are more likely to be positive, leading to a common good phenomenon.

Attribute Distribution Model

We make two assumptions. First, from all possible positive and negative attributes, some are present in a given person and some are absent. Second, we assume positivity prevalence, meaning that more positive than negative attributes are present on average. The same assumptions should apply to people's mental representation of others. Given these two assumptions, it follows mathematically that positive compared with negative attributes have a higher probability of being shared, and that shared compared with unshared attributes have a higher probability of being positive.

Figure 1 illustrates our reasoning with a feature logic (Tversky, 1977). Let us assume there are two target persons, X and Y. Their attributes are represented by the two columns on the left side of Figure 1. Both targets display 12 positive and six negative attributes of 20 possible positive and negative attributes, realizing positivity prevalence. In other words, in this ecology, targets display any positive attribute with a probability of $p(\text{pos}) = 0.6$ (i.e., 12 of 20) and any negative attribute with $p(\text{neg}) = 0.3$ (i.e., six of 20). Hence, both targets display twice as many positive as negative attributes.

It follows that positive compared with negative attributes that are present in one Target X are also twice as likely to be present in the other Target Y. This constitutes our first hypothesis that

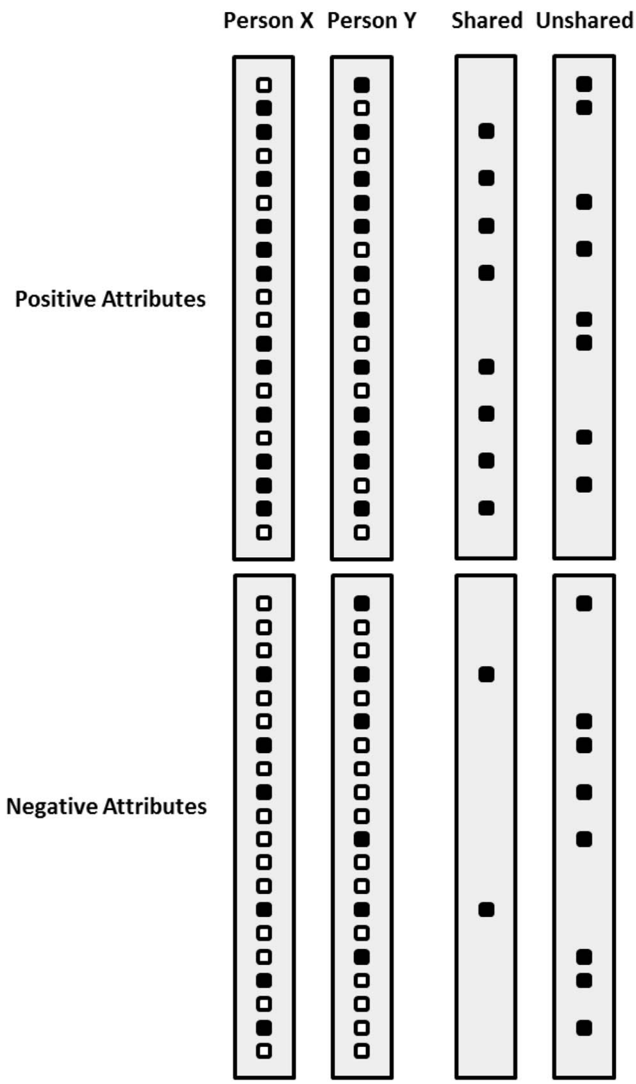

Figure 1. The two columns on the left depict the attribute profiles of two persons over 20 positive and 20 negative attributes. Filled rectangles symbolize presence of an attribute; nonfilled rectangles symbolize absence of an attribute. The ratio of positive to negative attributes in this example is 2:1 (i.e., 12 positive, 6 negative), indicating positivity prevalence. Shared attributes are those that are present in both persons, whereas unshared attributes are those that are present in one person but absent in the other person. The positive to negative ratio is amplified among the shared attributes (here, 4:1) and attenuated among the unshared attributes (here, 1:1).

positive compared with negative attributes should be more likely to be shared, that is, $p(\text{shared}|\text{positive}) > p(\text{shared}|\text{negative})$.

Although this prediction seems self-evident, it has an important consequence for the distribution of positive and negative attributes among shared and unshared attributes, which constitutes our second hypothesis, the common good phenomenon. The third column in Figure 1 depicts the attributes that are simultaneously present in both targets (shared attributes; similarities) and the attributes that are present in only one target (unshared attributes; differences). As the example shows, the initial ratio of positive to negative attributes (2:1) on the left is amplified among shared attributes (4:1). This amplification follows because the probability for the simultaneous presence of any positive attribute A_+ in the Targets X and Y is $p(t_{+_x} \cap t_{+_y}) = 0.6^*0.6 = 0.36$. The probability for the simultaneous presence of any negative attribute $p(t_{-_x} \cap t_{-_y}) = 0.3^*0.3 = 0.09$. Thus, the probability that two targets share a positive attribute is 4 times larger than the probability that the targets share a negative attribute. Consequently, shared attributes (similarities) are 4 times more likely to be positive than negative, and, generally, $p(\text{positive}|\text{shared}) > p(\text{negative}|\text{shared})$.

Among unshared attributes, on the other hand (the fourth column of Figure 1), the initial ratio of positive to negative attributes (2:1) is attenuated (1:1). This attenuation follows because the probability for the presence of any positive attribute A_+ in Target X and its simultaneous *absence* in Target Y is $p(A_{+_x} \setminus A_{+_y}) = 0.6^*0.4 = 0.24$. For any negative attribute A_-, this probability is $p(A_{-_x} \setminus A_{-_y}) = 0.3^*0.7 = 0.21$. Thus, the probability for any positive attribute to be present in only one person is only 1.1 times larger than the probability for a negative attribute to be present in only one person, reducing the initial positivity prevalence. Consequently, unshared attributes (differences) are about equally likely to be positive or negative in our example.

The model thus predicts that in a predominantly positive social world, people's similarities (shared attributes) are even more likely to be positive, whereas their differences (unshared attributes) are less likely to be positive; when social perceivers search for people's similarities, they should find more positive and less negative attributes compared with when they search for people's differences.

Figure 2 illustrates the statistical positivity amplification/attenuation that occurs when social perceivers sample similarities

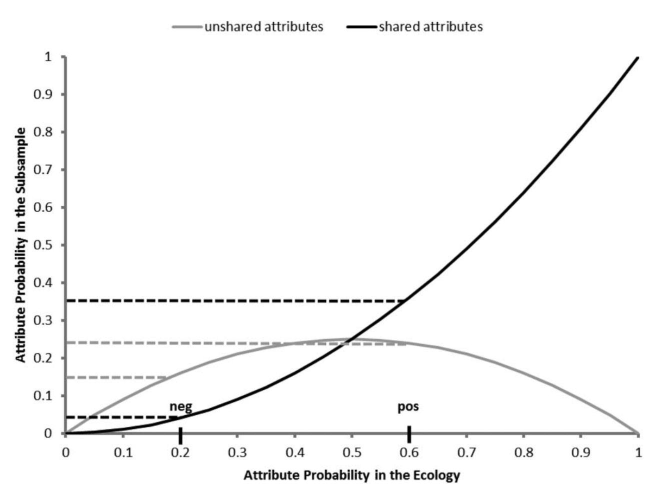

Figure 2. Attribute probabilities in samples of shared and unshared attributes (i.e., similarities vs. differences) as a function of ecological trait probabilities. The black curve depicts this relationship for traits that two targets share. Any given positive/negative ratio $ra \neq 1$ increases to ra^2. Here, a 2:1 ratio in the ecology ($p[\text{pos}] = 0.6$ vs. $p[\text{neg}] = 0.3$) is amplified to a 4:1 ratio among shared attributes ($p[\text{pos}] = 0.36$ vs. $p[\text{neg}] = 0.09$). The gray curve depicts the same relationship for unshared attributes. Any given positive/negative ratio that is unequal to 1 is attenuated. Here, a 2:1 ratio in the ecology is attenuated to a 1.1:1 ratio among unshared attributes ($p[\text{pos}] = 0.24$ vs. $p[\text{neg}] = 0.21$).

(shared attributes) and differences (unshared attributes). The *x*-axis depicts the ecological probability for positive and negative attributes, and the *y*-axis plots their probabilities of occurring in a sample of shared or unshared attributes. As Figure 2 shows, the prevalence amplification among shared attributes follows a simple rule: Any positive to negative ratio $ra \neq 1$ in the ecology is expected to increase to ra^2 among samples of shared attributes. The degree of attenuation among samples of unshared attributes is not as straightforward. As Figure 2 shows, the magnitude of attenuation depends on the specific probabilities of positive and negative attributes. However, any prevalence will be attenuated following an inverted "U"-shape function if the sampling process focuses on differences.

Further Model Specifications

Our model illustrates the common good phenomenon using a simple feature present–absent logic. Of course, the model abstracts from and simplifies reality. However, the probabilistic principle is highly robust regarding alternative attribute conceptualizations and additional assumptions: When the absence of a positive attribute is considered a negative attribute and vice versa; when attributes are conceptualized as continuous rather than binary (present vs. absent); when all attributes of the same valence do not have the same, but randomly varying base rates; when attributes are correlated; and even when both targets' attribute profiles are not independent of one another.

One restriction has to be made regarding attribute profile dependencies. With increasing dependency between the targets' attribute profiles, the changes of the initial ratios decrease. In real life, such dependencies might exist. For example, one might often encounter people with similar attributes (e.g., friends). However, the outlined statistical amplification and attenuation will only disappear when targets' profiles are either identical or completely different. In that case, shared as well as unshared attributes will not differ regarding their positive and negative attribute probabilities, and both will mirror the ecological base rates. As long as there is variation between targets, the basic principle will apply.

Overview of Empirical Demonstrations

We tested our model's predictions in a simulation and eight experiments (total $N = 1{,}181$). In all experiments, we asked participants to sample attributes (i.e., traits) of different target persons, such as personally known others or celebrities. Following our model, we made two main predictions regarding the trait sampling process. First, positive traits should be more likely to be shared among targets. We therefore asked participants to indicate which of a target person's positive and negative traits were shared by other targets (Experiments 1a and 1b). Second, shared compared with unshared traits should be more likely positive. We therefore asked participants to sample shared and unshared traits of two target persons and to indicate the valence of these traits (Experiments 2 and 4a). Trait valence was further validated by two independent participant samples (Experiment 4b). We also assessed the availability of positive and negative traits in participants' memory when sampling shared and unshared traits by measuring response latencies (Experiment 3). Experiments 5 and 6 tested a boundary condition for the proposed common good phenomenon as implied by our model, namely, that it should only occur if positive attributes are prevalent (cf. Figure 1). Thus, Experiments 5 and 6 featured predominantly positive as well as predominantly negative target persons. For all experiments, we report sample size, all data exclusions (if any), all manipulations, and all measures in the main text. Further methodological details can be found in the online supplemental materials. These include power analysis to determine sample size, participant gender, participant compensation, procedural specifics such as instructions, and experimental software.

Simulation

We implemented a simulation of the common good phenomenon using the Microsoft Visual Basic software. The simulation consisted of two steps that both introduced random variation.

Step 1: Simulation of the Ecology

The simulation followed our model's logic illustrated in Figure 1. Each simulation trial randomly generated attribute profiles for two targets. Each profile consisted of 20 positive and 20 negative attributes that could either be present or absent in a person.

Across nine different simulation conditions, we varied base rates for positive attributes (p[pos] = 0.6, 0.7, 0.8) and for negative attributes (p[neg] = 0.2, 0.3, 0.4) to simulate different degrees of positivity prevalence. For example, a base rate of $p(\text{pos}) = 0.6$ meant that of the 20 positive attributes, 12 randomly determined attributes were present and the other eight were absent. Once attribute profiles were generated, attributes present in both persons were defined as shared attributes, whereas attributes present in only one person were defined as unshared attributes (cf. Figure 1).

Step 2: Simulation of Attribute Sampling

Next, the simulation randomly sampled six attributes from the attribute profiles. In a similarity condition, the simulation sampled only from the shared attributes. In a difference condition, the simulation sampled from the unshared attributes. In a natural baseline condition, the simulation sampled from both targets' full attribute vectors. The dependent variable (DV) was the number of positive attributes in each condition. A common good phenomenon is indicated when the number of positive attributes in a sample exceeded three. The simulation consisted of 100 trials, drawing shared, unshared, and baseline attribute samples from the two targets. This simulation was repeated for all nine base rate combinations.

Results and Discussion

Figure 3 shows the simulation results. For all nine combinations of base rates, the shared sample always shows the largest average number of positive attributes, followed by the natural sample and the unshared sample. The ecological prevalence of positive attributes present in the natural sample is thereby amplified among shared attributes and attenuated among unshared attributes. Thus, implementing the mathematical assumption in a probabilistic simulation illustrates that the common good phenomenon is robust across a range of empirical parameters. We were thus confident

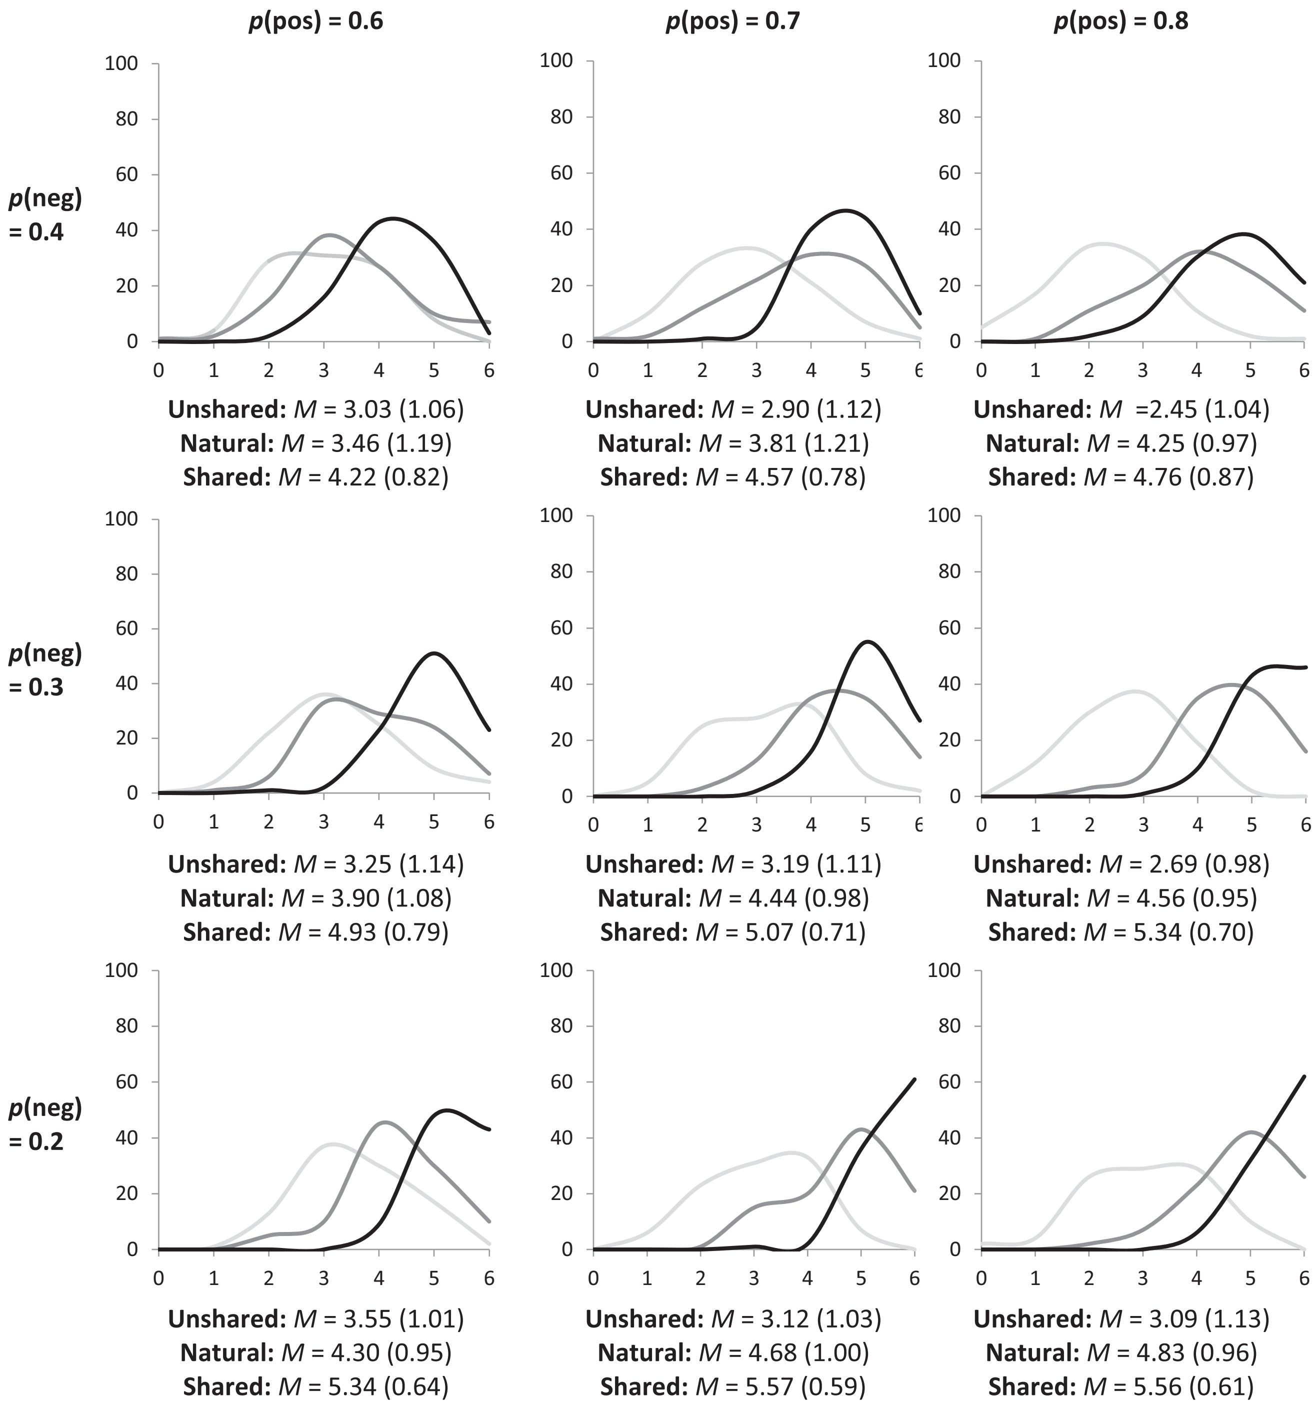

Figure 3. Simulated frequency (*y*-axis) distributions of positive traits (*x*-axis) in the unshared (light gray curves), natural (dark gray curves), and shared (black curves) samples for nine combinations of positive and negative trait base rates. Each graph plots how often a given number of positive traits was sampled in the three sampling conditions (unshared, natural, shared). Below each plot are the mean values and standard deviations within each simulation.

about the predicted effects and now turn to experimental tests of the model.

Experiment 1a

Experiment 1 tests whether positive traits are more likely to be shared by social targets than negative traits as predicted by our model. Participants provided the names of two persons they knew personally, and generated four positive and four negative character traits that described one of the persons. Then, participants indicated which of the traits also described the other person. We recruited 41 online participants from the United States. The only manipulation was the within-participants variation of trait valence. The DV was the number of shared traits.

Results

We computed the average frequency of positive and negative shared traits and compared these frequencies as a function of valence using a paired t test. As predicted, participants indicated that more positive than negative traits ($M_{pos} = 3.37$, $SD_{pos} = 0.86$ vs. $M_{neg} = 1.07$, $SD_{neg} = 1.17$) were shared by the targets, $t(40) = 9.08$, $p < .001$, $d = 2.24$. On average, for each positive trait that was present in one target person, the probability that it was shared by the other target person was $p(\text{shared}|\text{positive}) = .84$, whereas for each negative trait, the probability was only $p(\text{shared}|\text{negative}) = .27$.

Discussion

Participants indicated that more positive compared with negative traits that described one target also applied to the other target. Specifically, positive traits were about 3 times more likely to be shared than negative traits. Positive traits seem to define the similarities among these personally known targets, whereas negative traits constitute their differences.

Obviously, there are concerns regarding whether this first test really captures the attribute probabilities as specified in our model. First, participants could be motivated not to assign negative traits to people they know. We will return to the possibility of motivated responses in Experiment 2. Second, the results could be exclusive to comparisons of personally known others. For example, participants may have picked targets that are similar to themselves and thereby to each other. Those targets might then be similar regarding attributes that they share with participants and which participants therefore consider positive. To address the latter concern, Experiment 1b tested whether the positive compared with negative traits generated in Experiment 1a were also more likely to be shared by other participants' target persons.

Experiment 1b

Experiment 1b asked participants to provide the names of 10 persons they knew personally. Then, they were presented with one of the 41 different trait sets generated by participants in Experiment 1a, each consisting of four positive and four negative traits. Participants indicated, for each of their 10 target persons, whether the different traits applied to the target person or not. We recruited 82 online participants from the United States. The only manipulation was the within-participants variation of trait valence. The DV was again the number of shared traits. If our model is correct, positive compared with negative traits generated by participants in Experiment 1a should also be more frequently shared by target persons in the present experiment.

Results

We again tested the average number of shared traits as a function of valence using a paired t test. As predicted, participants indicated that more of the positive than negative traits applied to their target persons ($M_{pos} = 3.13$, $SD_{pos} = 0.63$ vs. $M_{neg} = 1.18$, $SD_{neg} = 0.65$), $t(81) = 16.59$, $p < .001$, $d = 3.06$. The mean probability for a positive trait to be shared by cross-experimental target pairs was $p(\text{shared}|\text{positive}) = .78$; the probability for negative traits was only $p(\text{shared}|\text{negative}) = .29$. As Figure 4 illustrates, the effect was evident for all 10 target persons that participants generated. This visual inference was confirmed statistically by a mixed model that specified participants and targets as random, with random error components for the intercept and valence effect (see Judd, Kenny, & Westfall, 2012). Contrast coding was used to code the fixed factor valence (.5 = positive; −.5 = negative). This mixed-model analysis affirmed the valence main effect: more positive compared with negative traits applied to the targets, $b = 1.96$, $t = 16.67$, $p < .001$. There was significant variation across participants regarding slopes and intercepts. Importantly, there was no significant variation across the ten target persons in the intercepts or in the slopes. Thus, across all 10 target persons that participants named, positive compared with negative traits were more likely to apply (see Figure 4).

Discussion

Participants indicated that more positive than negative traits applied to their targets, whereas traits were sampled from participants in Experiment 1a. Specifically, positive traits were about 2.6 times more likely to be shared than negative traits. Hence, the present results suggest that across different people's social mental representations, positive traits constitute the similarities and negative traits constitute the differences between persons. These results support our model's first prediction and its underlying probabilistic reasoning: Given that a trait is positive, it is more likely to be shared.

However, results of Experiment 1a and 1b might also reflect motivated responses; participants might be reluctant to assign negative traits to personally known others. Alas, when asking

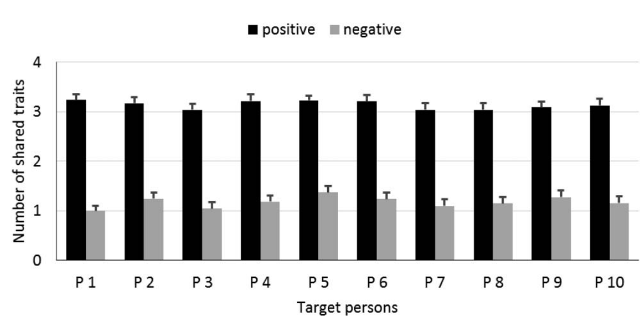

Figure 4. Number of positive and negative traits that apply to each of the 10 target persons that participants named. Error bars depict the standard error of the mean.

participants to assign positive traits and negative traits, our frequency-based explanation and the motivated reasoning explanation are inherently confounded. A stronger test is provided by the second prediction, which we are testing in the following experiments. Instead of testing whether positive compared with negative traits are more likely to be shared, we now test whether shared compared with unshared traits are more likely to be positive, which does not follow from motivated reasoning.

Experiment 2

Experiments 1a and 1b tested our first prediction: If it is good, it should be common. Experiment 2 tested the reversed prediction: If it is common, it should be good. Participants again named two personally known targets. We then manipulated whether participants sampled shared or unshared traits. Specifically, the shared condition asked for four traits that applied to both target persons simultaneously, and the unshared condition asked for two traits that applied to the first but not to the second target and two traits that applied to the second but not to the first target. The DV was the frequency of positive traits and negative traits, which should vary as a function of shared versus unshared traits. Participants also rated the valence of the traits they had generated on a scale ranging from 1 (*very negative*) to 7 (*very positive*). We recruited 73 online participants form the United States and randomly assigned them to one of the two conditions (shared or unshared traits).

Results

We coded traits with ratings between 1 and 3 as negative, and traits with ratings between 5 and 7 as positive. We excluded traits with ratings of 4 (neutral). Figure 5 shows the resulting frequencies. We analyzed these data with a Sampling Condition (shared vs. unshared) × Trait Valence (positive vs. negative) mixed ANOVA with repeated-measures on the latter factor. As Figure 5 indicates, trait valence yielded a main effect: Participants generated more positive than negative traits, which reflects the assumed positivity prevalence (M_{pos} = 2.86, SD_{pos} = 1.17 vs. M_{neg} = 0.78, SD_{neg} = 1.03), $F(1, 71) = 108.19$, $p < .001$, $\eta^2_p = .60$. Importantly, the analysis also found the predicted interaction effect. That is, the positivity prevalence was clearly stronger among shared traits than among unshared traits, $F(1, 71) = 30.41$, $p < .001$, $\eta^2_p = .30$.

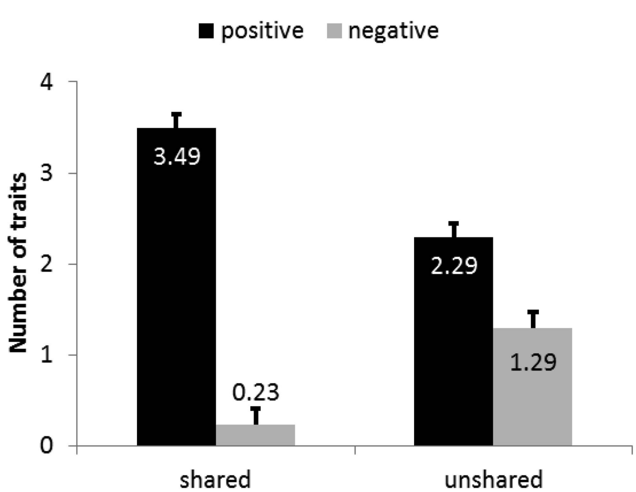

Figure 5. Mean frequencies of positive and negative traits among shared and unshared traits.

Discussion

When participants searched for shared traits between two people they knew, they almost exclusively sampled positive traits; this constitutes the proposed common good phenomenon. When participants searched for unshared traits, the frequency of sampled negative traits increased, whereas the frequency of positive traits decreased. Thus, searching for similarities versus differences differentially influenced the valence of the retrieved information. These effects directly follow from the assumed positivity prevalence as outlined in our probabilistic model. As the experiment did not include a baseline condition (i.e., the "natural" condition in the simulation), it is unclear, however, in which direction the sampling conditions deviated from a unrestricted baseline sampling. We return to this question in Experiment 4.

Experiment 2 also takes care of the motivational explanation for Experiments 1a and 1b. Although participants might be motivated to frequently assign positive traits, this motivation should be independent of the shared/unshared status of a given trait. That is, participants should be equally motivated to assign both shared and unshared positive traits to personally known others. The data show a clear increase in negative traits for unshared traits, which follows from our probabilistic model but not from a motivational explanation.

Again, we argue that the asymmetrical evaluative outcome results from the prevalence of positive traits in participants' memories. To test this claim, the next experiment moved from only investigating the evaluative outcome of the search process to investigating the search process itself. If positive traits are indeed prevalent in the mental representation of others, availability of positive and negative traits in memory should differ depending on whether participants search for shared or unshared traits.

Experiment 3

Our probabilistic model predicts that among shared traits, positive compared with negative traits have a higher base rate, whereas among unshared traits, the base rates are more even. Thus, when a perceiver searches for shared traits, positive compared with negative traits should be more available in memory, whereas this difference should be reduced when the perceiver searches for unshared traits (see simulations in Figure 3). To measure the availability of positive and negative traits among shared and unshared traits, we recorded participants' response latencies while they generated positive or negative traits that were either shared or unshared among two personally known target persons. Thus, Experiment 3 featured four experimental conditions that asked participants to generate positive/shared, negative/shared, positive/unshared or negative/unshared traits. In each condition, participants were instructed to sample four traits. In addition to the response latency measure, we also asked participants at the end of the experiment to rate the difficulty of trait sampling on a scale ranging from 1 (*very easy*) to 7 (*very difficult*). We recruited 114 students from the University of Cologne and randomly assigned them to one of the four conditions.

Results

Response latencies. As an index of trait availability, the computer program recorded the time that participants spent before entering each of the four traits into its respective textbox. The left half of Figure 6 shows the resulting mean latencies (in seconds) in the four conditions, which we analyzed using a Sampling Condition (shared vs. unshared) × Trait Valence (positive vs. negative) ANOVA. There was a significant main effect of valence: Participants were faster to generate positive compared with negative traits (M_{pos} = 50.92, SD_{pos} = 46.14 vs. M_{neg} = 71.13, SD_{pos} = 55.64), $F(1, 110) = 4.30$, $p = .040$, $\eta_p^2 = .04$. The sampling main effect did not reach conventional levels of significance, but, numerically, participants generated shared traits faster than unshared traits (M_{shared} = 51.70, SD_{shared} = 49.29 vs. $M_{unshared}$ = 70.29, $SD_{unshared}$ = 52.96), $F(1, 110) = 3.60$, $p = .061$, $\eta_p^2 = .03$. Crucially, as Figure 6 shows, participants generated positive compared with negative *shared* traits faster, but they took the same time to generate positive and negative traits that were *unshared*, resulting in a significant interaction, $F(1, 110) = 6.27$, $p = .014$, $\eta_p^2 = .06$.

Explicit ratings. The right half of Figure 6 shows participants mean rated difficulties across experimental conditions, which mirrored mean response latencies. Again, participants rated the task easier when they generated positive traits (M_{pos} = 4.32, SD_{pos} = 1.96 vs. M_{neg} = 5.87, SD_{neg} = 1.27), $F(1, 110) = 31.98$, $p < .001$, $\eta_p^2 = .23$. In addition, participants found it easier to generate shared compared with unshared traits (M_{shared} = 4.46, SD_{shared} = 2.09 vs. $M_{unshared}$ = 5.73, $SD_{unshared}$ = 1.21), $F(1, 110) = 20.86$, $p < .001$, $\eta_p^2 = .16$. Importantly, participants found it easier to generate positive shared than negative shared traits, but they found it equally difficult to generate positive and negative unshared traits, resulting in a significant interaction, $F(1, 110) = 25.00$, $p < .001$, $\eta_p^2 = .19$.

Discussion

Participants were faster and found it easier to generate positive compared with negative traits for personally known target persons, which follows from the assumed prevalence of positive traits. In addition, participants were numerically faster and found it significantly easier to generate shared compared with unshared traits, which is in line with research showing that finding matching features is easier than finding nonmatching features (Shiffrin & Schneider, 1977). Most importantly, participants were faster and found it easier to generate positive compared with negative traits when the traits were shared by the target persons, whereas there was no such difference for traits that were not shared by the target persons. This shows a differential availability of positive and negative traits among shared and unshared traits, which only follows from our model. As trait availability should reflect trait base rates, the results imply a stronger prevalence of positive traits among shared compared with unshared traits, in line with the common good phenomenon.

The previous experiments rendered first support for our model, but they omitted a critical test. Thus far, the results relied on the comparison of positive trait and negative trait frequency among shared and unshared traits. However, our model is more specific. The assumed positivity prevalence should be amplified by searching for shared traits *and* attenuated by searching for unshared traits. A baseline condition with an unconditional search should provide an estimate of the "natural" positivity prevalence, which, according to our model, is then amplified among shared and attenuated among unshared traits (cf. the "natural" condition in the simulations in Figure 3). The following experiment therefore replicated Experiment 2 and included an additional baseline condition (the "natural" condition).

Experiment 4a

Experiment 4a tested whether an initial prevalence of positive traits in participants' mental representation, as indicated by an unconditional search, is amplified among shared traits and attenuated among unshared traits. The experiment therefore included a "natural" baseline condition in which we simply asked participants to generate traits for two target persons without further specification.

Our model predicts that traits in the natural condition should be less likely to be positive than in the shared condition, but more likely to be positive than in the unshared condition. The natural condition affords another critical test: Although it was not specified whether traits had to be shared or unshared, participants might

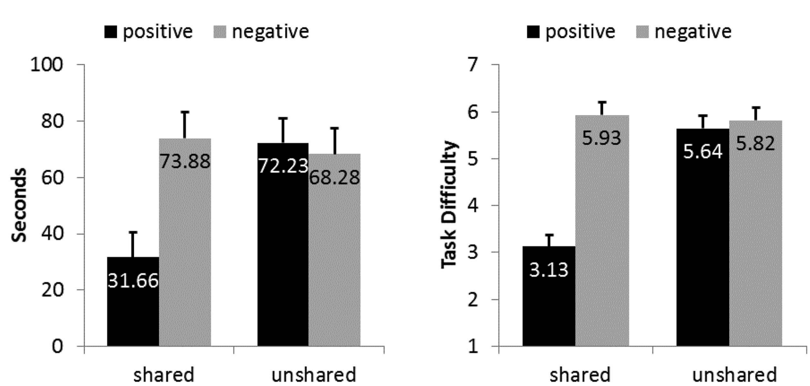

Figure 6. Mean response latencies (left half) and rated task difficulty for generating four traits as a function of trait valence (positive vs. negative) and sampling condition (shared vs. unshared). Error bars show the standard error of the mean.

still generate some traits repeatedly, thereby assigning the same trait to both persons. Traits assigned to both persons can be considered shared traits, and traits assigned to only one person can be considered unshared traits. Hence, we could compare the valence of shared and unshared traits within the natural condition. To increase the statistical power for this within-condition analysis, participants in the natural condition generated 12 traits, whereas participants in the shared and unshared conditions generated six traits. For comparisons across conditions, only the first six traits are included, making the three conditions factually equivalent. We recruited 176 students of the University of Cologne, who were randomly assigned to either the shared, unshared, or natural condition. The DV was again the frequency of positive and negative traits.

Results

Prior to analysis, we removed the data from two participants because they failed to provide personality traits and instead provided person descriptions such as “friend,” “student,” and “woman.”

For analyses between conditions, we calculated trait frequencies in the natural condition on the first six traits only. However, when we included all 12 trait ratings in the natural condition for between-participants analysis, all reported effects remained significant.

We again coded traits with a rating between 1 and 3 as negative traits, and traits with a rating between 5 and 7 as positive traits. Traits with a rating of 4 (neutral) were not considered. Figure 7 shows the resulting trait frequencies. We analyzed these data with a Sampling Condition (shared vs. natural vs. unshared) × Trait Valence (positive vs. negative) mixed ANOVA with repeated-measures on the latter factor. Again, there was a strong main effect of valence: Participants generated more positive than negative traits (M_{pos} = 4.78, SD_{pos} = 1.29 vs. M_{neg} = 0.67, SD_{neg} = 0.95), $F(1, 171) = 785.93$, $p < .001$, $\eta_p^2 = .82$. As Figure 7 suggests, there was also an interaction between sampling condition and trait valence, $F(2, 171) = 20.93$, $p < .001$, $\eta_p^2 = .20$. Planned contrasts confirmed that the difference between number of positive and negative traits was smaller in the natural than in the shared condition $F(1, 171) = 10.20$, $p = .002$, $\eta_p^2 = .06$, whereas it was larger in the natural compared with the unshared condition, $F(1, 171) = 11.00$, $p = .001$, $\eta_p^2 = .06$ (M_{shared} = 5.26, SD_{shared} = 1.38 vs. $M_{natural}$ = 4.12, $SD_{natural}$ = 2.31 vs. $M_{unshared}$ = 2.93, $SD_{unshared}$ = 1.97).

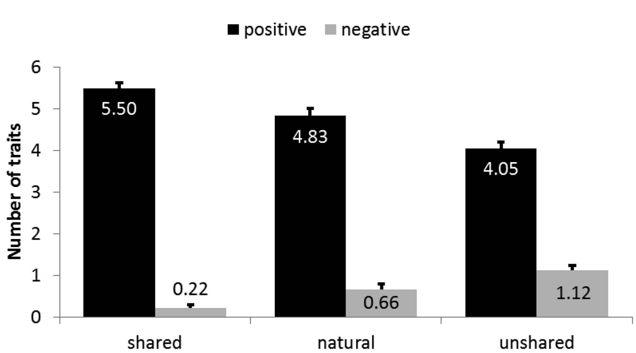

Figure 7. Mean trait frequencies of positive and negative traits in the shared, natural, and unshared condition. Error bars depict the standard error of the mean.

In the natural condition, participants were not instructed as to whether the traits they generated had to be shared or unshared traits. Nevertheless, 47 of the 59 participants in the natural condition generated at least one trait simultaneously for both targets. We considered these traits as shared traits, whereas we considered traits generated for only one target as unshared traits. For each participant, we divided the frequencies of positive and negative traits among shared and unshared traits by the total number of shared and unshared traits they had generated. The resulting index can be interpreted as the probability for positive/negative traits to occur among shared/unshared traits. Confirming previous results, positive traits were more likely to occur among shared traits (M = 0.91, SD = 0.26) than among unshared traits (M = 0.60, SD = 0.20), $t(46) = 7.27$, $p < .001$, $d = 1.34$. Conversely, negative traits were more likely to occur among unshared (M = 0.13, SD = 0.14) than among shared traits (M = 0.06, SD = 0.22), $t(46) = 2.09$, $p = .04$, $d = 0.32$.

Discussion

Experiment 4a fully replicated Experiment 2: Among shared traits, positive (negative) traits were more (less) frequent than among unshared traits. In addition, the “natural” condition provided an approximation of the natural prevalence of positive traits. This condition delivered a strong positivity prevalence, which was indeed *amplified* by searching for similarities (i.e., shared traits) and *attenuated* by searching for differences (i.e., unshared traits). The strong valence main effect may follow from multiple factors (e.g., both motivation and prevalence); the interaction follows only from our probabilistic model. Motivational explanations fail to explain the increase of negative traits and the decrease of positive traits from the shared to the natural to the unshared condition.

Furthermore, the common good phenomenon was visible not only across conditions but also within the natural condition. Traits assigned “naturally” to both persons were more positive than traits that were assigned to only one person. Hence, the effect was present even without any experimental manipulation.

We believe that Experiment 4a provides strong evidence for the common good phenomenon and for the proposed probabilistic explanation, as the data closely mirror the predictions from the sampling simulation presented in Figure 3.

A limitation of Experiments 2 and 4a arises from the fact that participants rated trait valence themselves. These ratings might be subject to a number of biases. For example, it is possible that the traits in the different conditions did not actually differ in valence but were simply interpreted more (less) positively when generated under a similarities (differences) instruction. Or as Experiment 3 showed the differential difficulty of trait generation in the two conditions, participants might evaluate traits that are difficult to generate more negatively. To validate participants’ trait ratings and the observed effects, we collected valence ratings from two independent samples of raters in Experiment 4b and repeated the same analyses from Experiment 4a with the new set of data.

Experiment 4b

Experiment 4b aimed at validating the trait ratings and observed effects from Experiment 4a. We conducted two different validations: An "individual" validation presented raters with all generated traits from Experiment 4a individually and asked them to rate their valence (on a 7-point scale), and an additional "holistic" validation presented raters with all traits that had been generated for each target and asked them to rate each target person's likability (on a 7-point scale). The latter task has the interesting implication that the common good phenomenon not only influences specific pieces of information, but indeed might change the overall evaluation of the target stimulus in a bottom-up fashion. Five research assistants participated in the individual validation, and 70 students of the University of Cologne participated in the holistic validation.

Results

Individual validation. We replaced participants' trait ratings from Experiment 4a with the trait valence ratings obtained from the independent raters in Experiment 4b and repeated analyses from Experiment 4a. Accordingly, degrees of freedom in the present analyses equal those from Experiment 4a. This analysis replicated Experiment 4a. Overall, most traits were positive (M_{pos} = 4.94, SD_{pos} = 1.16 vs. M_{neg} = 0.95, SD_{neg} = 1.12), $F(1, 170) = 672.41$, $p < .001$, $\eta_p^2 = .80$, and this effect was qualified by a Sampling × Trait Valence interaction, $F(2, 170) = 21.65$, $p < .001$, $\eta_p^2 = .20$. Planned contrasts showed that the positivity prevalence (difference between number of positive and negative traits) was smaller in the natural than in the shared condition, $F(1, 170) = 14.21$, $p = .003$, $\eta_p^2 = .08$, whereas it was larger in the natural compared with the unshared condition, $F(1, 170) = 7.85$, $p = .006$, $\eta_p^2 = .05$, (M_{shared} = 5.28, SD_{shared} = 1.31 vs. $M_{natural}$ = 3.86, $SD_{natural}$ = 2.25 vs. $M_{unshared}$ = 2.81, $SD_{unshared}$ = 2.34).

Holistic validation. We replaced participants' trait ratings from Experiment 4a with the holistic person ratings obtained from the independent raters and repeated the analyses from Experiment 4a. Figure 8 shows the mean likability ratings for the persons that our participants described in the shared, natural, and unshared conditions of Experiment 4a. A one-way ANOVA (sampling condition: shared vs. natural vs. unshared) on the mean likability ratings showed the corresponding effect of condition, $F(2, 170) = 35.45$, $p < .001$, $\eta_p^2 = .29$. Planned contrasts confirmed that persons described by shared traits were more likable than the persons described in the natural condition, $F(1, 170) = 11.79$, $p = .001$, $\eta_p^2 = .07$, and those were still more likable than persons described by unshared traits, $F(1, 170) = 24.58$, $p < .001$, $\eta_p^2 = .13$.

Discussion

Experiment 4b showed that previous results were not produced by possible confounds in the trait evaluations, for example, that easy-to-retrieve traits are evaluated more positively. The individual ratings from independent raters fully replicated Experiment 4a. Thus, searching for similarities or differences delivered traits that consensually differed in valence.

In addition, the holistic validation showed that target persons appear more or less likable depending on whether they are described by samples of shared or unshared traits. Participants rated the targets as overall likable, but they rated targets described by shared trait samples as more likable than targets described by "natural" trait samples, while they rated targets described by unshared trait samples as least likable. With regard to the model test, both Experiments 4a and 4b converge in showing that a given positivity prevalence is amplified among shared and attenuated among unshared traits, giving rise to the proposed common good phenomenon.

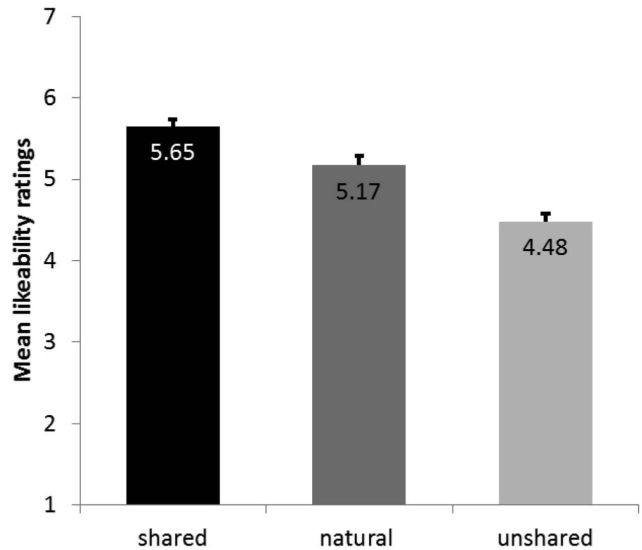

Figure 8. Mean likability ratings for target persons described in the shared, natural, and unshared conditions of Experiment 4a, based on participants' likability ratings in Experiment 4b. Error bars depict the standard error of the mean.

However, there are still at least two alternative explanations for the observed effects independent of our probabilistic explanation. First, participants self-sampled the target persons. The self-sampling procedure provides real target persons that are actually meaningful in participants' lives and guarantees positivity prevalence. Yet it comes with the cost of little control over possible confounds within this sampling process. For instance, participants might select people who are similar to themselves and therefore share attributes that participants evaluate positively. Likewise, target persons might be members of the same group (e.g., family, friend group). What these targets share are then group-typical attributes that participants probably evaluate positively. Experiment 1b addressed this by using target persons from a different set of participants. However, regarding the effects observed in Experiments 2 to 4, this concern remains.

A second concern arises because positive traits might be more likely to be shared because they are less diverse. Previous research suggests that positive information is less diverse than negative information. Here, this implies there might simply be a smaller number of distinct positive than negative traits, (Alves, Koch, & Unkelbach, 2016, 2017; Alves et al., 2015; Koch, Alves, Krüger, & Unkelbach, 2016; Koch, Imhoff, Dotsch, Unkelbach, & Alves, 2016; Unkelbach, 2012; Unkelbach, Fiedler, Bayer, Stegmüller, & Danner, 2008). If there are less distinct positive traits, this would also increase their likelihood of being shared relative to negative traits. This explanation would be different form the one suggested in our model: As Figure 1 clearly illustrates, the present model assumes the same number of possible positive and negative traits, and proposes differential frequency of positive and negative traits, not differential diversity, as the causal mechanism of the common good phenomenon.

To address both concerns, we moved from using self-sampled targets to experimenter-provided targets, avoiding possible confounds from specific target-participant relationships. To test our frequency-based explanation against a trait diversity explanation, we identified a critical test; our model predicts the common good phenomenon only when positive traits are prevalent. Thus, when negative traits are more frequent than positive traits, shared traits should be more negative than positive traits. The diversity explanation predicts that shared traits are still more likely to be positive, although they are less frequent in a predominantly negative stimulus ecology. In other words, our model predicts an interaction between sampling mode and stimulus ecology, whereas the diversity explanation predicts a main effect for sampling mode—that is, searching for similarities should lead to more positive traits in the sample independent of the frequency in the ecology. Experiments 5 and 6 will implement these tests.

Experiment 5

To compare shared and unshared traits among target persons that participants represent either with predominantly positive or predominantly negative traits, we required target persons that are evaluated differently by different people. Ideally, targets should be represented as predominantly positive by half of the participants and predominantly negative by the other half of participants.

We identified political figures in the United States as appropriate targets. The two-party system divides the U.S. population almost equally regarding their impressions about Republicans and Democrats. Thus, we asked U.S. participants to generate shared or unshared traits either for two well-known republicans (Mitt Romney and George W. Bush) or for two well-known democrats (Bill Clinton and Barack Obama). Participants indicated their liking of the target persons on a scale from -50 to $+50$. Subjective liking should indicate whether participants represented targets with predominantly positive or negative attributes. We predicted an interaction between participants' liking for the target persons and sampling mode. The diversity explanation predicts a main effect of sampling mode. We recruited 310 online participants from the United States who were randomly assigned to one of the four conditions that varied trait sampling mode (shared vs. unshared) and target persons (Republicans vs. Democrats); the DV was again frequency of positive and negative traits.

Results

Liking. We first separated participants regarding their liking for the target persons. We calculated each participant's mean liking of the two target persons. We considered liking values below zero as disliking target persons and values above zero as liking target persons. One hundred sixty of 310 participants indicated that they liked the target persons in their respective Republican and Democrats conditions, 143 indicated to dislike them, and seven participants produced an exact mean liking rating of 0. We excluded these seven participants from the following analyses. We quasi-experimentally split the remaining data file, consisting of 303 participants, into a "like" sample of 160 participants (52.8%) and a "dislike" sample of 143 participants (47.2%).

Traits. Figure 9 shows the mean frequencies of positive and negative traits in the shared and unshared conditions for participants who liked their targets and for participants who disliked their targets. We analyzed the difference between frequency of positive and negative traits with a Sampling Condition (shared vs. unshared) $\times$ Liking (liked targets vs. disliked targets) ANOVA. There was no main effect for sampling (shared vs. unshared; $M_{\text{shared}} = 1.31$, $SD_{\text{shared}} = 2.90$ vs. $M_{\text{unshared}} = 1.02$, $SD_{\text{unshared}} = 1.29$), $F(1, 299) = 0.00$, $p = .986$. There was a main effect for liking; participants generated more positive compared with negative traits for liked than for disliked targets ($M_{\text{like}} = 2.63$, $SD_{\text{like}} = 1.63$ vs. $M_{\text{dislike}} = -0.46$, $SD_{\text{dislike}} = 2.58$), $F(1, 299) = 166.95$, $p < .001$, $\eta_p^2 = .36$. Importantly, as Figure 9 suggests, there was an interaction between sampling condition and liking for this difference, $F(1, 299) = 32.21$, $p < .001$, $\eta_p^2 = .10$. Simple effects analysis showed that for participants who liked the target persons, the difference between number of positive and number of negative traits was larger for shared than for unshared traits, $F(1, 299) = 16.86$, $p < .001$, $\eta_p^2 = .05$, whereas the reverse was true for participants who disliked the target persons, $F(1, 299) = 15.43$, $p < .001$, $\eta_p^2 = .05$.

Discussion

Experiment 5 addresses the remaining alternative explanations. First, the common good phenomenon did not hinge on self-sampled target persons who might share positive traits with the participant and with each other (e.g., friends and family). Second, although the larger diversity of negative attributes as a possible alternative explanation also predicts shared traits to be more positive than unshared traits (e.g., Alves et al., 2016; Unkelbach et al., 2008), it does not predict the interaction between sampling mode and target valence. Experiment 5 clearly shows that shared compared with unshared traits are only more likely to be positive for liked targets, whereas the reverse is true for disliked targets. This interaction only follows from our present model.

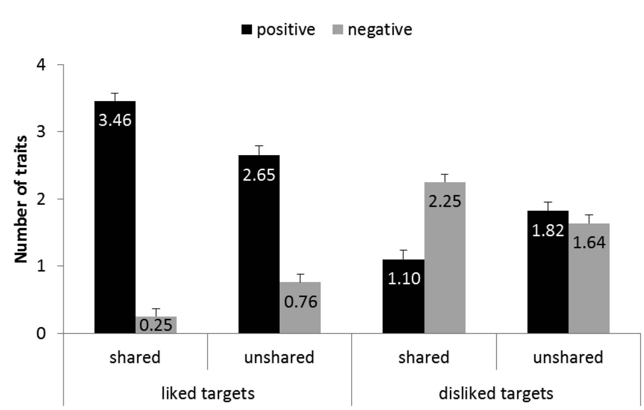

Figure 9. Frequency of positive and negative traits among shared and unshared traits separated for participants who liked versus disliked the target persons. Errors bars represent standard errors of the means.

One might still argue that participants' liking for the target persons was determined post hoc on the basis of ratings they provided at the end of the experiment. This rating was possibly influenced by the preceding trait-generation task. Depending on whether participants sampled positive or negative shared traits, participants might have concluded that they liked or disliked the targets. In addition, targets were similar to one another regarding their political party affiliation. Hence, the political figures in the current experiment are similar regarding traits associated with their party affiliation, which may be evaluated positively or negatively depending on participants' own party preferences. Hence, the interaction effect could still hinge on these specific target pairs. Experiment 6 addresses both concerns.

Experiment 6

Experiment 6 aimed at replicating the previous Experiment 5 with a larger, representative stimulus set of consensually liked and disliked target persons. Following the idea of representative stimulus sampling (Wells & Windschitl, 1999; Westfall, Kenny, & Judd, 2014), we collected the 10 most frequently mentioned liked and disliked celebrities from an independent participant sample (see the Appendix for the list of all 20 targets). Then, we randomly sampled the target pairs from these liked and disliked celebrities. Although this approach results in a more representative target set, and thereby increases experimental validity, the fully data-driven approach also introduces substantial noise into the trait sampling task resulting from unusual target pairs. For example, some participants were asked to provide traits shared by Justin Bieber and Adolf Hitler (two negative targets). As we felt committed to ruling out the possibility that the interaction between sampling mode and target valence hinges on specific target pairs, we accepted this increase in unsystematic variance. Experiment 6 again included four conditions in which participants were asked to generate four traits that were either shared or unshared between two consensually liked or disliked targets. We recruited 310 online participants from the United States who were randomly assigned to one of the four experimental conditions that varied trait sampling mode (shared vs. unshared) and target valence (positive vs. negative).

Results

Figure 10 shows the mean frequencies for positive and negative traits in the four conditions. We analyzed the difference between number of positive and negative traits in each condition with a Sampling Mode (shared vs. unshared) × Target Valence (positive vs. negative) ANOVA. There was a main effect for sampling mode ($M_{\text{shared}} = 1.13$, $SD_{\text{shared}} = 3.03$ vs. $M_{\text{unshared}} = 0.67$, $SD_{\text{unshared}} = 2.56$), $F(1, 306) = 4.01$, $p = .046$, $\eta_p^2 = .01$, and a main effect for target valence ($M_{\text{like}} = 2.72$, $SD_{\text{like}} = 1.61$ vs. $M_{\text{dislike}} = -0.85$, $SD_{\text{dislike}} = 2.60$), $F(1, 306) = 215.83$, $p < .001$, $\eta_p^2 = .41$. As predicted, there was again an interaction between sampling mode and target valence, $F(1, 306) = 5.28$, $p = .022$, $\eta_p^2 = .02$. Simple effects analysis showed that for liked targets, the difference between number of positive and number of negative traits was larger for shared than for unshared traits, $F(1, 306) = 9.17$, $p < .01$, $\eta_p^2 = .03$, whereas there was no such difference for disliked target persons, $F(1, 306) < 1$, $p = .884$.

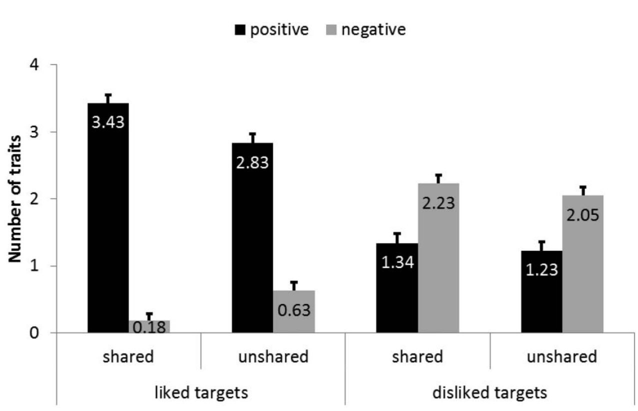

Figure 10. Frequency of positive and negative traits among shared and unshared traits for liked and disliked target persons. Errors bars represent standard errors of the means.

Discussion

Experiment 6 replicated Experiment 5, as trait valence was a function of sampling mode (shared vs. unshared) and target valence (positive vs. negative). Importantly, target persons were not self-sampled by participants, and their valence was not determined by each participant individually in a post hoc manner. Targets were instead generated in a bottom-up fashion by an independent sample of participants, and their valence was consensually determined. However, this procedure also creates substantial noise in the target pairs because of very unusual comparisons.

Although the observed interaction effect between sampling mode and target valence clearly supports our frequency-based model, one finding deserves further attention, as it deviates from Experiment 5: The simple effect of sampling mode was only significant for liked targets, whereas it was not significantly reversed for disliked targets. Searching for disliked targets' similarities did not produce more negative traits than searching for their dissimilarities. This might hint to an actual influence of negative traits larger diversity (e.g., Alves et al., 2016; Unkelbach et al., 2008). Given a larger universe of negative traits, participants might still be able to find unique negative traits among disliked targets. However, the interaction effects we found in the previous and the current experiment can only be explained by our frequency-based model. Results thereby suggest that the main part of the common good phenomenon is a function of the proposed prevalence of positive traits.

General Discussion

We started with the assumption that people live in a predominantly positive social world. Most people behave positively and display positive attributes most of the time, whereas norm-violating behavior and attributes are rare. The present work shows that this positivity prevalence has intriguing implications for person perception, in particular, and for comparison processes, in general. We formalized these implications in an attribute distribution model.

The model makes two central predictions: Positive attributes are more likely to be shared by individuals, and people's shared attributes (similarities) are more likely to be positive than their unshared attributes (differences). More specifically, the positivity prevalence is amplified among shared attributes and attenuated among unshared attributes. Consequently, when perceivers search for targets' similarities, they are more likely to find positive compared with negative traits. When perceivers search for differences between targets, the likelihood that they find negative attributes increases. As a result, people's social perception is characterized by a common good phenomenon—common attributes are more positive than unique attributes.

We showed this theoretically derived phenomenon in a simulation and confirmed it empirically in eight experiments. Table 1 provides an overview of the effects that we predicted and found, along with sample and effect sizes for the respective experiments. These experiments ruled out several alternative explanations by varying different aspects of the experimental designs. We varied the types of target persons by using personally known others (Experiments 1 to 4), politicians (Experiment 5), and celebrities (Experiment 6). We employed different independent variables, namely, trait valence (Experiments 1 and 3) and sampling mode (Experiments 2 and 6). We used different dependent measures, including trait assignment (Experiment 1), valence ratings from participants (Experiments 2, 4a, 5, and 6) and independent raters (Experiment 4b), and response latencies and task difficulty ratings (Experiment 3).

In order to estimate the average effect sizes for our main predictions, we conducted a mini meta-analysis using fixed effects, in which the mean effect sizes were weighted by sample size (Goh, Hall, & Rosenthal, 2016). All correlations were then Fisher's z transformed for analyses and converted back to Pearson correlations for presentation. Analysis of Experiments 1a and 1b revealed that our first predicted effect—namely, that positive compared with negative traits are more likely to be shared by targets—was large and highly significant ($r = .86$, $z = 13.77$, $p < .001$). Analysis of Experiments 2 to 4 showed that our second predicted effect—namely, that shared compared with unshared traits are more likely to be positive—was also large and highly significant ($r = .63$, $z = 9.27$, $p < .001$). Finally, analysis of Experiments 5 and 6 showed a smaller, but still significant, interaction effect, which confirms the boundary condition for the common good phenomenon as implied by our model: Shared traits are only more likely to be positive when they are overall more frequent ($r = .24$, $z = 5.90$, $p < .001$). Based on the distribution of the p values across all eight experiments, the estimated achieved statistical power was 99% (p-curve.com; Simonsohn, Nelson, & Simmons, 2014). The analysis included eight tests of our main predictions.

Implications

The common good phenomenon has a number of implications for social perception and might renew the understanding of widely observed evaluative asymmetries in the social domain. We begin with a general implication for social distinctiveness. The common good phenomenon implies that those attributes that make people distinguishable are likely to be negative ones. Conversely, those attributes that connect people are very likely to be positive. According to optimal distinctiveness theory, people's social identity is driven by two contradicting motivations, namely, need for affiliation and need for distinctiveness (Brewer, 1991). To fulfill both needs, people want to be similar to their social peers but also different from them. In light of the present work, the latter motivation comes with a fundamental problem: Distinct attributes are likely to be negative. Though not impossible, displaying distinct attributes that are also positive constitutes a challenge in a world in which positive attributes are prevalent. Although it is easy to fulfill the need for distinctiveness by engaging in norm-violating acts such as cursing, lying, and stealing, it is difficult to achieve the same through norm-congruent acts like being kind, honest, and helpful. Hence, people with a chronically high need for distinctiveness should be especially prone to engage in negative behavior. Likewise, in situations that trigger a need for distinctiveness, people should also show norm violations. The need for affiliation, on the other hand, should drive people toward norm-congruent, and thus desirable, behavior.

The reverse implication is also noteworthy: People's negative attributes should be more likely to be considered typical or characteristic of them because they make them distinct. Hence, nega-

Table 1
Descriptions of Observed Effects, Sample Sizes, and Effect Sizes Across the Eight Experiments

Experiment	Effect	n	Effect size
1a	Positive compared with negative traits were more likely to be shared by two personally known target persons	41	$d = 2.24$
1b	Positive compared with negative traits from Experiment 1a's targets were more likely to be shared by 10 target persons from a different group of participants	82	$d = 3.06$
2	Searching for targets' shared traits resulted in more positive traits than searching for targets' unshared traits	73	$\eta_p^2 = .30$
3	Among shared traits, positive traits were more available than negative traits, whereas among unshared traits, positive and negative traits were about equally available	114	RL: $\eta_p^2 = .06$ ER: $\eta_p^2 = .19$
4a	Searching for targets' shared traits resulted in more positive traits compared with a baseline condition, whereas searching for targets' unshared traits resulted in less positive traits compared with baseline	176	$\eta_p^2 = .20$
4b	Confirmed results from Experiment 4a with trait and target ratings from an independent sample of participants	75	$\eta_p^2 = .20$
5	Amplification/attenuation of positivity among political targets depended on participants' liking for the targets	310	$\eta_p^2 = .10$
6	Amplification/attenuation of positivity among a set of consensually liked and disliked targets depended on targets' likeability	310	$\eta_p^2 = .02$

Note. RL = response latencies measure; ER = explicit ratings in Experiment 3.

tive attributes should receive greater weighting in person perception, which is often observed in person perception research. People are being judged primarily on their negative attributes because these attributes distinguish them best (see Fiske, 1980, and Skowronski & Carlston, 1989, for a related idea). The common good phenomenon thereby implies a general distinctiveness/desirability conflict.

Social evaluation. The current framework has a number of implications for social evaluations. When evaluations depend on the similarities or differences of targets and standards (Hodges, 2005; Houston & Sherman, 1995; Kahneman & Miller, 1986; Mussweiler, 2003; Tversky, 1977), the present work suggests a hitherto undiscovered asymmetry: Evaluations based on differences will be negatively skewed, whereas evaluations based on similarities will be positively skewed.

Another aspect of comparison processes is that they entail a *direction* of comparison, as a given target is compared with some standard. Extensive research has shown that the evaluation of a target is typically driven by its unique attributes, whereas the shared attributes of target and standard are cancelled out (Houston & Sherman, 1995; Tversky & Gati, 1978; see Hodges, 2005, for an overview). At the same time, a standard is typically evaluated based on its complete set of attributes regardless of whether they are shared or not shared by the target. The greater impact of the target's unique attributes has been demonstrated for judgments of similarity (Srull & Gaelick, 1983; Tversky & Gati, 1978), typicality (Hodges & Hollenstein, 2001), feature change (Agostinelli, Sherman, Fazio, & Hearst, 1986), and evaluations (Hodges, 1997; Hodges, Bruininks, & Ivy, 2002; Houston, Sherman, & Baker, 1989; Wänke, Schwarz, & Noelle-Neumann, 1995).

If unique attributes are more likely to be negative, as suggested here, comparison targets have an evaluative disadvantage, whereas standards enjoy an advantage, as long as target and standard are predominantly positive. This rather abstract notion becomes more meaningful if one considers what usually determines whether an object or a person is treated as target or standard in comparison processes. Comparison standards are prototypical stimuli (Hodges, 2005), which are more familiar to the perceiver (Karylowski, 1990), which more frequently occur (Polk, Behensky, Gonzalez, & Smith, 2002), and which are encountered first (Houston et al., 1989). We can therefore expect evaluative disadvantages for nonprototypical, unfamiliar, infrequent, and new persons, groups, and objects merely based on probabilistic reasoning. This, in turn, provides a new explanation for various prominent evaluative asymmetries in the social domain. We present three examples that demonstrate how the current framework can inform theories about intergroup bias, self–other comparisons, and temporal comparisons.

Intergroup bias. Intergroup bias describes the tendency to evaluate members of one's own group (in-group) more favorably than members of other groups (out-groups). This tendency reveals itself in behavior (discrimination), attitude (prejudice), and cognition (stereotyping; Hewstone et al., 2002; Mackie & Smith, 1998). It is assumed that this bias is driven by the motivation to maintain or achieve a positive social identity (e.g., Rubin & Hewstone, 1998), optimal distinctiveness (Leonardelli & Brewer, 2001), or social dominance (Sidanius & Pratto, 1999), or to reduce uncertainty through in-group identification (Hogg & Abrams, 1993).

However, the common good phenomenon may also contribute to the formation of intergroup bias. The in-group typically serves as a standard of comparison, whereas out-groups and their members constitute comparison *targets*, which is in line with the notion that targets are nonprototypical, unfamiliar, rare, and new stimuli. We can then expect out-groups and their members to be associated with and evaluated based on their unique attributes, that is, attributes they do not share with in-groups (Kruschke, 2003; Sherman et al., 2009). Whereas in-groups enjoy positive evaluations reflecting the general positivity prevalence in the social world, out-groups are associated with and judged based on their unique features, which have a smaller probability of being positive. Our model thereby provides a simple framework for intergroup biases that arise as a natural consequence of the information environment and the perceiver's attempt to differentiate social groups. The same principle is also applicable to the evaluative advantage of majority groups over minority groups (Katz, & Braly, 1935; Thalhammer et al., 2001).

The self as the standard. People are most familiar with themselves, and it is fair to assume that the self serves as a comparison standard in most social comparisons (Dunning & Hayes, 1996). If we assume that the representation of the self accurately mirrors the prevalence of positive attributes in the social world, and that social targets other than the self are evaluated based on how they are different from the self, people should generally perceive themselves as superior compared with their social environment. This widely observed phenomenon is often explained by the human motivation to maintain a positive self-view (Alicke & Govorun, 2005; Hoorens, 1993; Taylor & Brown, 1988). The current framework provides a statistical explanation. Further, with increasing social distance or dissimilarity between a social target and the self, the perception of self-superiority should increase. This follows from the notion that perceived dissimilarity triggers difference-based comparison (Mussweiler, 2003). Likewise, the observation that people perceive their current self as superior to their past self (Wilson & Ross, 2000, 2001) might follow from the same argument beyond people's motivation to enhance the current self. That is, the self in the here and now serves as a standard that past selves are compared with based on their unique attributes.

Overcoming evaluative biases. The current model also suggests a way to overcome these discussed evaluative biases by focusing on the similarities of standards (i.e., the in-group, the majority, the self) and targets (i.e., the out-group, the minority, the past). Related strategies have been proposed before, but their effects were usually interpreted as motivational (e.g., Byrne & Clore, 1967; Montoya & Horton, 2013). For example, intergroup bias is reduced when members of different groups conceive of themselves to be part of the same, higher order group (common group identity; e.g., Gaertner, Dovidio, Anastasio, Bachman, & Rust, 1993). Further, focusing on how someone is similar to the self generally increases liking for that person (Berscheid, 1985; Klohnen & Luo, 2003; Montoya, Horton, & Kirchner, 2008). Highlighting interpersonal similarity is commonly assumed to reduce evaluative biases because it reduces people's self-enhancement motives. Such explanations for example refer to Heider's (1958) balance theory or Festinger's (1957) concept of cognitive dissonance. Alternatively, the current work suggests that highlighting similar attributes means highlighting positive attributes, which should naturally lead to a more positive representation of social targets.

Conclusion

If positive attributes are prevalent, people's common (shared) attributes are more positive than their unique (unshared) attributes. This common good phenomenon implies that similarity-based social comparisons are positively biased, whereas difference-based comparisons are negatively biased. Evaluative disadvantages for social targets might therefore result from the perceiver's need to distinguish different persons, groups, and objects in a predominantly positive environment.

References

Agostinelli, G., Sherman, S. J., Fazio, R. H., & Hearst, E. S. (1986). Detecting and identifying change: Additions versus deletions. *Journal of Experimental Psychology: Human Perception and Performance, 12,* 445–454. http://dx.doi.org/10.1037/0096-1523.12.4.445

Alicke, M. D., & Govorun, O. (2005). The better-than-average effect. In M. D. Alicke, D. A. Dunning, & J. I. Krueger (Eds.), *The Self in Social Judgment* (pp. 85–106). New York, NY: Psychology Press.

Alves, H., Koch, A., & Unkelbach, C. (2016). My friends are all alike—The relation between liking and perceived similarity in person perception. *Journal of Experimental Social Psychology, 62,* 103–117. http://dx.doi.org/10.1016/j.jesp.2015.10.011

Alves, H., Koch, A., & Unkelbach, C. (2017). Why good is more alike than bad: Processing Implications. *Trends in Cognitive Sciences, 21,* 69–79. http://dx.doi.org/10.1016/j.tics.2016.12.006

Alves, H., Unkelbach, C., Burghardt, J., Koch, A. S., Krüger, T., & Becker, V. D. (2015). A density explanation of valence asymmetries in recognition memory. *Memory & Cognition, 43,* 896–909. http://dx.doi.org/10.3758/s13421-015-0515-5

Anderson, N. H. (1981). *Foundations of information integration theory*. New York, NY: Academic Press.

Augustine, A. A., Mehl, M. R., & Larsen, R. J. (2011). A positivity bias in written and spoken English and its moderation by personality and gender. *Social Psychological and Personality Science, 2,* 508–515. http://dx.doi.org/10.1177/1948550611399154

Beauregard, K. S., & Dunning, D. (1998). Turning up the contrast: Self-enhancement motives prompt egocentric contrast effects in social judgments. *Journal of Personality and Social Psychology, 74,* 606–621. http://dx.doi.org/10.1037/0022-3514.74.3.606

Berscheid, E. (1985). Interpersonal attraction. In G. Lindzey & E. Aronson (Eds.), *Handbook of social psychology* (3rd ed., pp. 413–484). New York, NY: Random House.

Biswas-Diener, R., Vittersø, J., & Diener, E. (2005). Most people are pretty happy, but there is cultural variation: The Inughuit, the Amish, and the Maasai. *Journal of Happiness Studies, 6,* 205–226. http://dx.doi.org/10.1007/s10902-005-5683-8

Boucher, J., & Osgood, C. E. (1969). The Pollyanna hypothesis. *Journal of Verbal Learning & Verbal Behavior, 8,* 1–8. http://dx.doi.org/10.1016/S0022-5371(69)80002-2

Brewer, M. B. (1991). The social self: On being the same and different at the same time. *Personality and Social Psychology Bulletin, 17,* 475–482. http://dx.doi.org/10.1177/0146167291175001

Brunswik, E. (1956). *Perception and the representative design of psychological experiments* (2nd ed.). Berkeley, CA: University of California Press.

Byrne, D., & Clore, G. L., Jr. (1967). Effectance arousal and attraction. *Journal of Personality and Social Psychology, 6,* 1–18. http://dx.doi.org/10.1037/h0024829

Clark, H. H., & Clark, E. V. (1977). *Psychology and language: An introduction to psycholinguistics*. New York, NY: Harcourt Brace Jovanovich.

Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Hillsdale, NJ: Erlbaum.

Denrell, J. (2005). Why most people disapprove of me: Experience sampling in impression formation. *Psychological Review, 112,* 951–978. http://dx.doi.org/10.1037/0033-295X.112.4.951

Diener, E., & Diener, C. (1996). Most people are happy. *Psychological Science, 7,* 181–185. http://dx.doi.org/10.1111/j.1467-9280.1996.tb00354.x

Dodds, P. S., Clark, E. M., Desu, S., Frank, M. R., Reagan, A. J., Williams, J. R., . . . Danforth, C. M. (2015). Human language reveals a universal positivity bias. *Proceedings of the National Academy of Sciences of the United States of America, 112,* 2389–2394. http://dx.doi.org/10.1073/pnas.1411678112

Dunning, D., & Hayes, A. F. (1996). Evidence for egocentric comparison in social judgment. *Journal of Personality and Social Psychology, 71,* 213–229. http://dx.doi.org/10.1037/0022-3514.71.2.213

Fazio, R. H., Eiser, J. R., & Shook, N. J. (2004). Attitude formation through exploration: Valence asymmetries. *Journal of Personality and Social Psychology, 87,* 293–311. http://dx.doi.org/10.1037/0022-3514.87.3.293

Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford, CA: University Press.

Fiedler, K. (1996). Explaining and simulating judgment biases as an aggregation phenomenon in probabilistic, multiple-cue environments. *Psychological Review, 103,* 193–214. http://dx.doi.org/10.1037/0033-295X.103.1.193

Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to judgment biases. *Psychological Review, 107,* 659–676. http://dx.doi.org/10.1037/0033-295X.107.4.659

Fiedler, K. (2014). From intrapsychic to ecological theories in social psychology: Outlines of a functional theory approach. *European Journal of Social Psychology, 44,* 657–670. http://dx.doi.org/10.1002/ejsp.2069

Fiedler, K., & Juslin, P. (2006). Taking the interface between mind and environment seriously. In K. Fiedler & P. Juslin (Eds.), *Information sampling and adaptive cognition* (pp. 3–29). New York, NY: Cambridge University Press.

Fiske, S. T. (1980). Attention and weight in person perception: The impact of negative and extreme behavior. *Journal of Personality and Social Psychology, 38,* 889–906. http://dx.doi.org/10.1037/0022-3514.38.6.889

Gaertner, S. L., Dovidio, J. F., Anastasio, P. A., Bachman, B. A., & Rust, M. C. (1993). The common ingroup identity model: Recategorization and the reduction of intergroup bias. *European Review of Social Psychology, 4,* 1–26. http://dx.doi.org/10.1080/14792779343000004

Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston, MA: Houghton, Mifflin.

Goh, J. X., Hall, J. A., & Rosenthal, R. (2016). Mini meta-analysis of your own studies: Some arguments on why and a primer on how. *Social and Personality Psychology Compass, 10,* 535–549.

Greenberg, M. S., Saxe, L., & Bar-Tal, D. (1978). Perceived stability of trait labels. *Personality and Social Psychology Bulletin, 4,* 59–62. http://dx.doi.org/10.1177/014616727800400112

Guenther, C. L., & Alicke, M. D. (2010). Deconstructing the better-than-average effect. *Journal of Personality and Social Psychology, 99,* 755–770. http://dx.doi.org/10.1037/a0020959

Heider, F. (1958). *Psychology of interpersonal relations*. New York, NY: Wiley.

Hewstone, M., Rubin, M., & Willis, H. (2002). Intergroup bias. *Annual Review of Psychology, 53,* 575–604. http://dx.doi.org/10.1146/annurev.psych.53.100901.135109

Hodges, S. D. (1997). When matching up features messes up decisions: The role of feature matching in successive choices. *Journal of Personality and Social Psychology, 72,* 1310–1321. http://dx.doi.org/10.1037/0022-3514.72.6.1310

Hodges, S. D. (2005). A feature-based model of self-other comparisons. In M. D. Alicke, D. A. Dunning, & J. I. Krueger (Eds.), *The self in social judgment* (pp. 131–153). New York, NY: Psychology Press.

Hodges, S. D., Bruininks, P., & Ivy, L. (2002). It's different when I do it: Feature matching in self-other comparisons. *Personality and Social Psychology Bulletin, 28,* 40–53. http://dx.doi.org/10.1177/0146167202281004

Hodges, S. D., & Hollenstein, T. (2001). Direction of comparison in typicality judgments. *Social Cognition, 19,* 601–624. http://dx.doi.org/10.1521/soco.19.6.601.20887

Hogg, M. A., & Abrams, D. (1993). Towards a single-process uncertainty-reduction model of social motivation in groups. In M. A. Hogg & D. Abrams (Eds.), *Group motivation: Social psychological perspectives* (pp. 173–190). Hertfordshire, UK: Harvester Wheatsheaf.

Hoorens, V. (1993). Self-enhancement and superiority biases in social comparison. *European Review of Social Psychology, 4,* 113–139. http://dx.doi.org/10.1080/14792779343000040

Houston, D. A., & Sherman, S. J. (1995). Cancellation and focus: The role of shared and unique features in the choice process. *Journal of Experimental Social Psychology, 31,* 357–378. http://dx.doi.org/10.1006/jesp.1995.1016

Houston, D. A., Sherman, S. J., & Baker, S. M. (1989). The influence of unique features and direction of comparison on preferences. *Journal of Experimental Social Psychology, 25,* 121–141. http://dx.doi.org/10.1016/0022-1031(89)90008-5

Judd, C. M., Westfall, J., & Kenny, D. A. (2012). Treating stimuli as a random factor in social psychology: a new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology, 103,* 54.

Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review, 93,* 136–153. http://dx.doi.org/10.1037/0033-295X.93.2.136

Karylowski, J. J. (1990). Social reference points and accessibility of trait-related information in self-other similarity judgments. *Journal of Personality and Social Psychology, 58,* 975–983. http://dx.doi.org/10.1037/0022-3514.58.6.975

Katz, D., & Braly, K. W. (1935). Racial prejudice and racial stereotypes. *The Journal of Abnormal and Social Psychology, 30,* 175–193. http://dx.doi.org/10.1037/h0059800

Klohnen, E. C., & Luo, S. (2003). Interpersonal attraction and personality: What is attractive—Self similarity, ideal similarity, complementarity or attachment security? *Journal of Personality and Social Psychology, 85,* 709–722. http://dx.doi.org/10.1037/0022-3514.85.4.709

Koch, A., Alves, H., Krüger, T., & Unkelbach, C. (2016). A general valence asymmetry in similarity: Good is more alike than bad. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 42,* 1171–1192. http://dx.doi.org/10.1037/xlm0000243

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

Kruschke, J. K. (1996). Base rates in category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22,* 3–26. http://dx.doi.org/10.1037/0278-7393.22.1.3

Kruschke, J. K. (2001). The inverse base-rate effect is not explained by eliminative inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27,* 1385–1400. http://dx.doi.org/10.1037/0278-7393.27.6.1385

Kruschke, J. K. (2003). Attention in learning. *Current Directions in Psychological Science, 12,* 171–175. http://dx.doi.org/10.1111/1467-8721.01254

Leonardelli, G. J., & Brewer, M. B. (2001). Minority and majority discrimination: When and why. *Journal of Experimental Social Psychology, 37,* 468–485. http://dx.doi.org/10.1006/jesp.2001.1475

Lewin, K. (1951). *Field theory in social science: Selected theoretical papers* (D. Cartwright, Ed.). Oxford, UK: Harpers.

Mackie, D. M., & Smith, E. R. (1998). Intergroup relations: Insights from a theoretically integrative approach. *Psychological Review, 105,* 499–529. http://dx.doi.org/10.1037/0033-295X.105.3.499

Matlin, M. W., & Stang, D. J. (1978). *The Pollyanna principle: Selectivity in language, memory, and thought*. Cambridge, MA: Schenkman.

Montoya, R. M., & Horton, R. S. (2013). A meta-analytic investigation of the processes underlying the similarity-attraction effect. *Journal of Social and Personal Relationships, 30,* 64–94. http://dx.doi.org/10.1177/0265407512452989

Montoya, R., Horton, R. S., & Kirchner, J. (2008). Is actual similarity necessary for attraction? A meta-analysis of actual and perceived similarity. *Journal of Social and Personal Relationships, 25,* 889–922. http://dx.doi.org/10.1177/0265407508096700

Mussweiler, T. (2003). Comparison processes in social judgment: Mechanisms and consequences. *Psychological Review, 110,* 472–489. http://dx.doi.org/10.1037/0033-295X.110.3.472

Parducci, A. (1968). The relativism of absolute judgements. *Scientific American, 219,* 84–90. http://dx.doi.org/10.1038/scientificamerican1268-84

Perlman, D., & Oskamp, S. (1971). The effects of picture content and exposure frequency on evaluations of Negroes and whites. *Journal of Experimental Social Psychology, 7,* 503–514. http://dx.doi.org/10.1016/0022-1031(71)90012-6

Polk, T. A., Behensky, C., Gonzalez, R., & Smith, E. E. (2002). Rating the similarity of simple perceptual stimuli: Asymmetries induced by manipulating exposure frequency. *Cognition, 82,* B75–B88. http://dx.doi.org/10.1016/S0010-0277(01)00151-2

Ric, F., Alexopoulos, T., Muller, D., & Aubé, B. (2013). Emotional norms for 524 French personality trait words. *Behavior Research Methods, 45,* 414–421. http://dx.doi.org/10.3758/s13428-012-0276-z

Rothbart, M., & Park, B. (1986). On the confirmability and disconfirmability of trait concepts. *Journal of Personality and Social Psychology, 50,* 131–142. http://dx.doi.org/10.1037/0022-3514.50.1.131

Rubin, M., & Hewstone, M. (1998). Social identity theory's self-esteem hypothesis: A review and some suggestions for clarification. *Personality and Social Psychology Review, 2,* 40–62. http://dx.doi.org/10.1207/s15327957pspr0201_3

Sears, D. O. (1983). The person-positivity bias. *Journal of Personality and Social Psychology, 44,* 233–250. http://dx.doi.org/10.1037/0022-3514.44.2.233

Sherman, J. W., Kruschke, J. K., Sherman, S. J., Percy, E. J., Petrocelli, J. V., & Conrey, F. R. (2009). Attentional processes in stereotype formation: A common model for category accentuation and illusory correlation. *Journal of Personality and Social Psychology, 96,* 305–323. http://dx.doi.org/10.1037/a0013778

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review, 84,* 127–190. http://dx.doi.org/10.1037/0033-295X.84.2.127

Sidanius, J., & Pratto, F. (1999). *Social dominance: An intergroup theory of social hierarchy and oppression*. New York, NY: Cambridge University Press. http://dx.doi.org/10.1017/CBO9781139175043

Simonsohn, U., Nelson, L. D., & Simmons, J. P. (2014). P-curve: A key to the file-drawer. *Journal of Experimental Psychology: General, 143,* 534–547. http://dx.doi.org/10.1037/a0033242

Skowronski, J. J., & Carlston, D. E. (1989). Negativity and extremity biases in impression formation: A review of explanations. *Psychological Bulletin, 105,* 131–142. http://dx.doi.org/10.1037/0033-2909.105.1.131

Srull, T. K., & Gaelick, L. (1983). General principles and individual differences in the self as a habitual reference point: An examination of self–other judgments of similarity. *Social Cognition, 2,* 108–121. http://dx.doi.org/10.1521/soco.1983.2.2.108

Tajfel, H., & Turner, J. C. (1979). An integrative theory of intergroup conflict. In W. G. Austin & S. Worchel (Eds.), *The social psychology of intergroup relations* (pp. 33–47). Monterey, CA: Brooks/Cole.

Taylor, S. E., & Brown, J. D. (1988). Illusion and well-being: A social psychological perspective on mental health. *Psychological Bulletin, 103,* 193–210. http://dx.doi.org/10.1037/0033-2909.103.2.193

Thalhammer, E., Zucha, V., Enzenhofer, E., Salfinger, B., & Ogris, G. (2001). *Attitudes toward minority groups in the European Union: A special analysis of the Eurobarometer 2000 survey*. Vienna, Austria: European Monitoring Centre on Racism and Xenophobia.

Thorndike, E. L. (1898). Animal intelligence: An experimental study of the associative processes in animals. *The Psychological Review: Monograph Supplements, 2,* i–109.

Tversky, A. (1977). Features of similarity. *Psychological Review, 84,* 327–352. http://dx.doi.org/10.1037/0033-295X.84.4.327

Tversky, A., & Gati, I. (1978). Studies of similarity. In E. Rosch & B. Lloyd (Eds.), *Cognition and categorization* (pp. 79–98). Hillsdale, NJ: Erlbaum.

Unkelbach, C. (2012). Positivity advantages in social information processing. *Social and Personality Psychology Compass, 6,* 83–94. http://dx.doi.org/10.1111/j.1751-9004.2011.00407.x

Unkelbach, C., Fiedler, K., Bayer, M., Stegmüller, M., & Danner, D. (2008). Why positive information is processed faster: The density hypothesis. *Journal of Personality and Social Psychology, 95,* 36–49. http://dx.doi.org/10.1037/0022-3514.95.1.36

Walker, W. R., Skowronski, J. J., & Thompson, C. P. (2003). Life is pleasant—And memory helps to keep it that way! *Review of General Psychology, 7,* 203–210. http://dx.doi.org/10.1037/1089-2680.7.2.203

Wänke, M., Schwarz, N., & Noelle-Neumann, E. (1995). Asking comparative questions: The impact of the direction of comparison. *Public Opinion Quarterly, 59,* 347–372. http://dx.doi.org/10.1086/269481

Wells, G. L., & Windschitl, P. D. (1999). Stimulus sampling and social psychological experimentation. *Personality and Social Psychology Bulletin, 25,* 1115–1125. http://dx.doi.org/10.1177/01461672992512005

Westfall, J., Kenny, D. A., & Judd, C. M. (2014). Statistical power and optimal design in experiments in which samples of participants respond to samples of stimuli. *Journal of Experimental Psychology: General, 143,* 2020–2045. http://dx.doi.org/10.1037/xge0000014

Wilson, A. E., & Ross, M. (2000). The frequency of temporal-self and social comparisons in people's personal appraisals. *Journal of Personality and Social Psychology, 78,* 928–942. http://dx.doi.org/10.1037/0022-3514.78.5.928

Wilson, A. E., & Ross, M. (2001). From chump to champ: People's appraisals of their earlier and present selves. *Journal of Personality and Social Psychology, 80,* 572–584. http://dx.doi.org/10.1037/0022-3514.80.4.572

Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology, 9,* 1–27. http://dx.doi.org/10.1037/h0025848

Appendix

List of the 10 Most Frequently Mentioned Negative and Positive Target Persons That Were Used in Experiment 6

Negative targets	Positive targets
Adolf Hitler	Abraham Lincoln
Donald Trump	John F. Kennedy
George W. Bush	Elvis Presley
Osama Bin Laden	Martin Luther King
Saddam Hussein	Oprah Winfrey
Joseph Stalin	Taylor Swift
Kim Jong Un	George Washington
Justin Bieber	Michael Jordan
Fidel Castro	Beyoncé Knowles
Kanye West	Jesus Christ

Received August 12, 2016
Revision received December 20, 2016
Accepted December 23, 2016 ■