Minority salience and the overestimation of individuals from minority groups in perception and memory

Rasha Kardosh[a,1,2], Asael Y. Sklar[b,2], Alon Goldstein[a], Yoni Pertzov[a], and Ran R. Hassin[c]

[a]Psychology Department, The Hebrew University of Jerusalem, Jerusalem 91905, Israel; [b]Edmond and Lily Safra Center for Brain Sciences, The Hebrew University of Jerusalem, Jerusalem 9190401, Israel; and [c]James Marshall Chair of Psychology, Psychology Department, The Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem, Jerusalem 9190401, Israel

Edited by Susan Fiske, Princeton University, Princeton, NJ; received September 14, 2021; accepted January 17, 2022

Our cognitive system is tuned toward spotting the uncommon and unexpected. We propose that individuals coming from minority groups are, by definition, just that—uncommon and often unexpected. Consequently, they are psychologically salient in perception, memory, and visual awareness. This minority salience creates a tendency to overestimate the prevalence of minorities, leading to an erroneous picture of our social environments—an illusion of diversity. In 12 experiments with 942 participants, we found evidence that the presence of minority group members is indeed overestimated in memory and perception and that masked images of minority group members are prioritized for visual awareness. These findings were consistent when participants were members of both the majority group and the minority group. Moreover, this overestimated prevalence of minorities led to decreased support for diversity-promoting policies. We discuss the theoretical implications of the illusion of diversity and how it may inform more equitable and inclusive decision-making.

minority salience | diversity | social perception | crowd perception | overestimation of minorities

In recent years, the realization that promoting diversity is a crucial stage on the road to an equitable society and that it is advantageous to our society as a whole has become more broadly accepted (1–5). While human societies are far from reaching this goal, we suggest that when we look at our social environments (e.g., workplaces, neighborhoods, or campuses), we overestimate the prevalence of minority group members, creating an illusion of diversity. Try to imagine walking down your campus lawn or the last conference you attended, and ask yourself what percentage of the students you encounter belong to a minority group? If you are anything like 82% of our participants, you vastly overestimate this percentage (yes, even if you yourself are a member of this minority group).

It is well established that, under many circumstances, our cognition is tuned toward spotting irregularities; our attention is drawn to unexpected events—prioritizing for us what is uncommon or unexpected (6–8). This tendency to pay attention to the unexpected is present from the lowest level of vision and from infancy (9, 10). Uncommon and unexpected objects pop out in perception (11, 12) and are also more likely to be remembered afterward (13–15). In addition, when sampling large groups of stimuli, those that differ from their surroundings attract attention and are therefore overweighted in our judgments (16). We suggest that, in many contexts*, individuals from a minority group are, by definition, uncommon and hence unexpected. Therefore, it stands to reason that, in these contexts, members of the minority group are salient; because they are unexpected, they stand out, automatically drawing attention (17–19). We hypothesize that this salience leads to an overestimation of their prevalence in perception and memory (20–22). Simply put, because we are more likely to notice and remember seeing them, we consistently overestimate the prevalence of minorities in our environments.

Cognitive salience often plays a role in human judgments (19, 23–26). Salience impacts risk evaluations: when judging the likelihood of the occurrence of rare risks (e.g., a car accident or an airplane crash), those that are salient are consistently overestimated in the participants' judgments (22, 27). Salience also affects social evaluations; it determines perceived attributes and amplifies evaluations of individuals (24, 28, 29). Finally, salience impacts the evaluation of groups. Research on illusory correlations established that the cooccurrence of two salient events, such as rare behaviors enacted by members of rare social groups, leads to a bias in judging the attributes of these social groups (19, 30–32). Here, we suggest that because individuals from minority groups are less frequent and, in some contexts, less expected, they are salient in our cognitive system, and this salience affects our judgments of the nature and structure of our social environment.

Note, however, that in addition to the mere statistical prevalence of minority groups that guides expectations, our knowledge of the structure of our society also shapes our expectations (33–35). Therefore, members of the minority group who do not

Significance

Our minds are tuned to the uncommon or unexpected in our environment. In most environments, members of minority groups are just that—uncommon. Therefore, the cognitive system is tuned to spotting their presence. Our results indicate that individuals from minority groups are salient in perception, memory, and visual awareness. As a result, we consistently overestimate their presence—leading to an illusion of diversity: the environment seems to be more diverse than it actually is, decreasing our support for diversity-promoting measures. As we try to make equitable decisions, it is important that private individuals and decision-makers alike become aware of this biased perception. While these sorts of biases can be counteracted, one must first be aware of the bias.

Author contributions: R.K., A.Y.S., A.G., Y.P., and R.R.H. designed research; R.K. performed research; R.K. analyzed data; and R.K., A.Y.S., and R.R.H. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence may be addressed. Email: rasha.kardosh@mail.huji.ac.il.

[2]R.K. and A.Y.S. contributed equally to this work.

This article contains supporting information online at http://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2116884119/-/DCSupplemental.

Published March 14, 2022.

*A social group can be a majority in one context and a minority in another. See Experiment 5 for elaboration on the role of social knowledge.

conform to knowledge-based expectation are also likely to be salient and hence overestimated. Generally, in many circumstances, both statistical prevalence and social knowledge lead us to overestimate the presence of minorities in our environments.

Support for the notion that we overestimate the prevalence of minority groups comes from national polls. When responders are asked to estimate the prevalence of national minorities, across many countries, polls consistently find that the prevalence of minority groups is overestimated. Black Americans make up 13% of the American population, but responders estimate they make up well over 30% (36–39). In France, Muslims constitute 7% of the population, but French responders estimated Muslims to be 30% of the French population; in fact, across Europe, the prevalence of Muslims is overestimated in each of the countries tested (40).

This overestimation of minority groups has adverse implications; previous research in social psychology has established that the perceived size of a national minority is associated with a host of crucial outcomes (41–48), among which are perceived threat and hostile attitudes toward minorities (41, 49), anti-immigrant attitudes, and perceived anti-White racial discrimination (50). In fact, the perceived size of a minority group is often more predictive of such attitudes than the actual size (48, 50). Taken together, this literature, with research originating from various countries, establishes that individuals who perceive the national minority as larger are more likely to hold adversarial attitudes toward this minority group. In addition, reminding participants in the United States about the growth in the racial minority share of the national population also leads to more adversarial attitudes (44, 50).

As we suggest in previous paragraphs, the overestimation observed in polls may be part of a broader phenomenon; it is but the end result of a robust, general, psychological mechanism that characterizes our experience of our immediate social environments as well: because of the way our cognitive system is built, rare events are psychologically salient. In many contexts, minorities tend to be rare and, hence, salient. Because of their salience, we perceive, remember, and consciously experience them more. These effects lead to an overestimation of the prevalence of minorities, creating an illusion of diversity. In other words, most of our social environments are likely to be less diverse than they seem to be in our minds.

To test our theoretical framework, we examine whether minority salience occurs in memory for one's real-life immediate social environments (Experiments 1A and 1B) and in new, controlled experimental social environments (Experiments 2A and 2B). We then test minority salience in perception (Experiments 3A, 3B, and 3C), and we examine whether minority salience operates at the level of prioritization for visual awareness, that is, whether individuals from minority groups are more readily available to visual awareness (Experiments 4A and 4B). Importantly, because minority salience results from the basic architecture of our cognitive system, we predicted that it also characterizes minority members themselves (Experiments 1B through 4B). We then argue that minority salience is not merely a statistical phenomenon, it is a social one. Hence, broader social knowledge would moderate the effect of minority salience. This hypothesis is examined in Experiment 5. Finally, in Experiment 6, we test the causal role of minority salience in decreasing support for diversity-promoting policies.

Results

Experiments 1A and 1B. In Experiments 1A and 1B, we tested whether students overestimate the percentage of Palestinian Israeli students (Arab students) at the Hebrew University. All participants were students at the Hebrew University. A total of 32 Jewish Israeli students participated in Experiment 1A, and 30 Palestinian Israeli students participated in Experiment 1B[†]. We asked students to recall instances of walking in the main hallway of the Hebrew University campus (Mount Scopus) and to reply to a few questions about the typical experience (e.g., "Is there a long line for the cafeteria?"). Afterward, the participants were asked to estimate "out of all the people you encounter along the walk, how many are Arab (0 to 100%)?". Finally, the participants were asked to estimate the percentage of Arab students in the university. The actual percentages of the latter were obtained from the Hebrew University student enrollment offices.

While the actual percentage of Arab students studying in this campus is 9.28%, both Jewish Israeli participants' average estimate (M = 31.56%, SD = 9.28, Experiment 1A) and Palestinian Israeli participants' average estimate (M = 35.81%, SD = 14.28, Experiment 1B) were significantly larger (t (31) = 13.58, $P < 0.001$, d = 2.40 and t (25) = 9.47, $P < 0.001$, d = 1.86, respectively). In fact, each of the 62 Jewish and Palestinian participants overestimated the prevalence of Arab students in the Hebrew University. Interestingly, this university-wide estimate was highly correlated with the percentage of Arab students participants remembered seeing in the corridor ($r = 0.66$, $P < 0.001$). While this is only a correlation, it suggests that prevalence estimates do rely on representation in memory.

It is important to note that Palestinian Israelis (Arabs) comprise 21% of the population in Israel. So, while this group in underrepresented by about 50% in the student population, in the students' judgments, it is overrepresented by roughly 30%.

We conducted a joint analysis of both experiments to test whether the effect differed significantly between Palestinian versus Jewish participants. The results indicate that both Palestinian Israeli participants and Jewish Israeli participants overestimated the percentage of Arab students in a similar manner (t (56) = 1.37, $P = 0.18$). This result may suggest that the overestimation is not influenced by in-group/out-group affiliation, a suggestion we further test in the next experiments.

Experiments 2A and 2B. Experiments 2A and 2B examine this effect in controlled experimental settings in which the statistical prevalence of the minority group displayed is systematically controlled. In these experiments, American participants viewed matrices of faces of Black American and White American individuals (Fig. 1). Experiments 2A and 2B were identical, except for the participant pool; White Americans (n = 100) participated in Experiment 2A, and Black Americans (n = 100) participated in Experiment 2B[‡]. The participants were informed that they will be shown images of American college students. Each image consisted of a matrix of 100 faces of Black Americans and White Americans. The overall percentage of Black American faces viewed throughout the study was set to be 25% in one condition and 45%[§] in the second condition. In the 25% condition, the percentages of Black American faces in each matrix were uniformly distributed from 10 to 40%. In the 45% condition, the percentages in each matrix varied between 30 and 60% in uniform distribution. The locations of the faces were randomly distributed. The participants viewed 20 matrices, each displayed for 2 s, and had an incidental task—to indicate, for each matrix, whether it included mostly female or male faces. This unrelated task intended to emulate natural

[†]Preregistration, including stopping rules for data collection, is available at https://aspredicted.org/blind.php?x=/5CN_22K and https://aspredicted.org/TNQ_Y41.

[‡]Preregistration available at https://aspredicted.org/blind.php?x=/LAT_ZPX and https://aspredicted.org/blind.php?x=/BT2_KH4.

[§]The 45% condition was included in order to test whether the overestimation may stem from a tendency to respond toward the middle of the scale. As the results show, the estimates are significantly over 50% (t (47) = 4.64, $P < 0.001$), indicating that a response bias toward the center of the scale cannot account for the observed overestimation.

Fig. 1. An example of a matrix used in Experiments 2, 3, and 5. Each matrix was displayed for 2 s. The percentages of Black American faces in each matrix were manipulated according to the experiment and condition. The locations of the faces were randomly distributed.

settings (after all, we do not normally go around the world counting/estimating minorities). After viewing all matrices, the participants were asked to estimate the percentage of African American faces they had seen throughout the study and the percentage of White American faces.

We hypothesized that the percentage of Black American faces would be overestimated, while the percentage of White American faces would be underestimated. The participants indeed vastly overestimated the percentage of Black American faces. In Experiment 2A, in the 25% Black Americans condition, the mean estimate was 43.22% (SD = 15.57; t (49) = 8.28, $P < 0.001$, $d = 1.17$), and in the 45% condition, the mean estimate was 58.85% (SD = 13.23; t (47) = 7.25, $P < 0.001$, $d = 1.05$). In Experiment 2B, the mean estimates were 43.36% (SD = 17.23; t (46) = 7.31, $P < 0.001$, $d = 1.07$) and 56.18% (SD = 16.27; t (48) = 4.81, $P < 0.001$, $d = 0.69$) in the 25 and 45% conditions, respectively (Fig. 2).

As in Experiments 1A and 1B, the lack of differences between White and Black American participants ($F(1, 190) = 0.32$, $P > 0.5$) suggests that a participants' ingroup plays no role in the effect. Overall, 83.7% of participants showed overestimation in Experiment 2A and 78.1% in Experiment 2B.

Complementarily to the overestimation of Black American faces, when asked about the percentage of White American faces, the participants significantly underestimated the percentage in Experiment 2A by 18.55% in the 25% Black Americans condition and by 12.13% in the 45% condition (SD = 13.87; t (49) = −9.37, $P < 0.001$, $d = 1.34$; SD = 13.68; t (47) = −6.14, $P < 0.001$, $d = 0.89$, respectively) and in Experiment 2B by 13.37% in the 25% Black Americans condition and by 7.02% in the 45% condition (SD = 16.13; t (45) = −5.62, $P < 0.001$, $d = 0.83$; SD = 16.51; t (48) = −2.98, $P = 0.005$, $d = 0.43$, respectively).

The results of Experiments 2A and 2B thus conceptually replicate the findings observed in Experiments 1A and 1B regarding the recalled prevalence of minority group members in a carefully controlled experimental context. These results additionally show the complementary effect, an underestimation of the prevalence of the majority group.

Experiments 3A, 3B, and 3C. Four experiments so far tested the overestimation of minority group members' prevalence based on memory: participants were asked to provide estimates of minority groups' prevalence, either based on their experience on campus or at the end of the experiment after viewing all matrices of faces. Experiments 3A and 3B tested whether this overestimation would be evident even when estimates are based on perception—participants were asked to provide estimates immediately after viewing each matrix. Participants (all White Americans in Experiment 3A [$n = 50$] and Black Americans in Experiment 3B [$n = 50$])[¶] viewed 100 images of face matrices similar to those used in Experiment 2A and 2B that included 10 to 50% Black American faces (overall, 30% across the experiment). Unlike previous experiments, participants were asked to provide an estimate immediately after each matrix was displayed. To prevent participants from trying to count the number of relevant faces in each matrix (or use other strategies), we included filler trials; after viewing each matrix for 2 s, participants were asked to estimate one of the four categories—Black Americans, White Americans, women, and men. The trials on which participants estimated the percentage of men/women served as filler trials and were excluded from analysis. To test overestimation in perception, we calculated the difference between a participant's estimate and the actual percentage displayed in the matrix on each trial.

We expected an overestimation when participants estimated the percentage of Black Americans and underestimation when participants rated the percentage of White Americans. Finally, to replicate the previous results and examine the association between estimates in perception and memory, participants were asked to estimate the overall prevalence of White and African Americans from memory at the end of the study.

To test for overestimation, we examined whether the difference between participants' estimation and the actual prevalence was significantly different from zero. Overall, participants significantly overestimated the percentage of Black Americans immediately after viewing the images (Experiment 3A: mean = 7.14%, SD = 7.76; t (49) = 6.51, $P < 001$, d = 0.92. Experiment 3B = mean = 9.58%, SD = 9.02; t (49) = 7.51, $P < 0.001$, d = 1.06). Complementarily, participants significantly underestimated the percentage of White Americans in images immediately after perceiving them (Experiment 3A: mean = −9.98%, SD = 5.60; t (49) = −12.60, $P < 0.001$, d = 1.78; Experiment 3B: mean = −12.87%, SD = 6.35; t (48) = −14.17, $P < 0.001$, d = 2.03). Notably, the estimates provided by Black American participants in Experiment 3B did not differ from the estimates provided by White American participants in Experiment 3A (all Fs < 1).

On memory estimates, replicating the effect found in Experiments 2A and 2B, participants significantly overestimated the prevalence of Black Americans when recalling it at the end of the experiment (overestimation M = 15.02%, SD = 12.20, t (47) = 8.53, $P < 0.001$, d = 1.23 and M = 13.47%, SD = 12.12, t (46) = 7.62, $P < 0.001$, d = 1.59 in Experiments 3A and 3B, respectively) and significantly underestimated the prevalence of White Americans (M = −5.19%, SD = 9.61; t (47) = −3.74, $P = 0.001$, d = 0.54 and M = −9.42%, SD = 12.40, t (47) = −4.53, $P < 0.001$, d = 0.76 in Experiments 3A and 3B, respectively; Fig. 2). Notably, the participants' memory-based overestimations at the end of the experiment were significantly larger than their overestimation immediately after perceiving the images (Experiment 3A: mean difference = 7.79%, SD = 10.23, t (47) = 5.27, $P < 0.001$, d = 0.76; Experiment 3B: mean difference = 3.96%, SD = 10.90, t (46) = 2.49, $P = 0.016$, d = 0.36).

[¶]Preregistration available at https://aspredicted.org/blind.php?x=/QHS_6LH, https://aspredicted.org/blind.php?x=/164_M3D, and https://aspredicted.org/blind.php?x=/QXN_QGR.

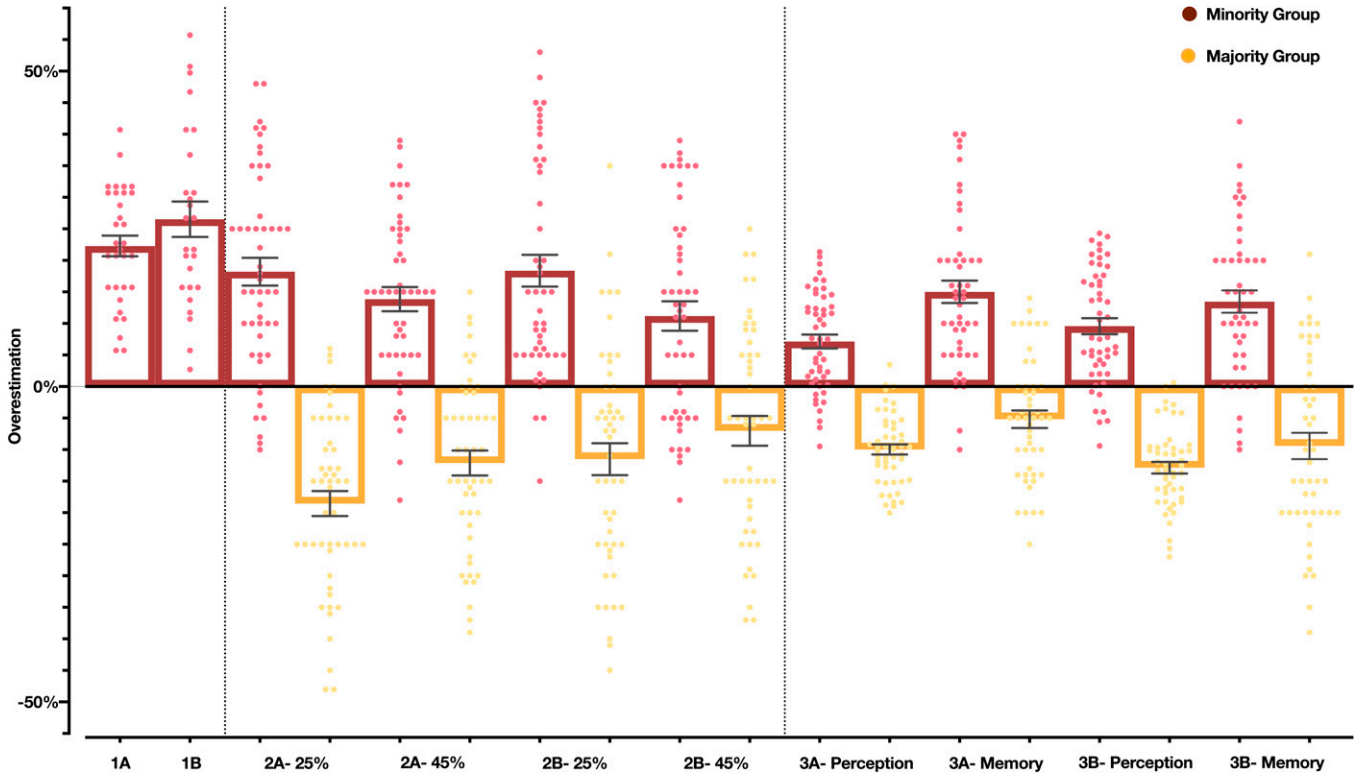

Fig. 2. The results of Experiments 1A to 3B. Zero denotes accurate estimates, positive values denote overestimation, and negative values underestimation. In Experiments 1A and 1B, the participants estimated the percentage of Arab students at the Hebrew University. In Experiment 2, the participants estimated the percentage of Black American faces (red) and White American faces (yellow) from memory in two conditions; overall prevalence is 25 or 45%. In Experiment 3, the participants estimated the percentage of Black American faces and White American faces in perception and memory. The participants in Experiments 1A, 2A, and 3A belong to the majority group (Jewish Israelis or White Americans). The participants in Experiments 1B, 2B, and 3B belong to the minority groups (Palestinian Israelis or Black Americans). Overall, 82.1% of participants overestimated the percentage of the minority group.

Taken together, these findings indicate minority salience leads to both an immediate overestimation of the minority just following perception and an additional amplification of this overestimation when recalling the overall sequence of images from memory.

A repeated measures ANOVA testing the overestimation of Black Americans across the five actual prevalence conditions (10 to 50%) indicated that the overestimation was largest when the actual prevalence of the minority group was lowest ($F(4, 176) = 28.62$, $P < 0.001$, $\eta_p^2 = 0.39$ and $F(4, 144) = 41.90$, $P < 0.001$, $\eta_p^2 = 0.54$ in Experiment 3A and 3B, respectively, *SI Appendix*, Fig. 1). In many real-life situations, minorities are less prevalent than in our experiments. For example, Black Americans comprise 13% of college students in the United States (51) and about half of that in some elite colleges. In such cases, our results indicate the overestimation would likely be even larger than in most of the conditions we tested so far (see Experiment 6 for estimates on a 5% minority prevalence).

Finally, to examine the overestimation of minority group members' prevalence in the perception of more realistic scenes, Experiment 3C employed photos from naturalistic settings. Participants ($n = 50$, all White Americans) viewed nine pictures which included a large number of individuals (e.g., a picture of a busy subway station). All pictures included both Black American and White American individuals, and participants were asked to estimate what percentage of the individuals in each image were Black Americans immediately after viewing it for 2 s. Replicating the results of Experiments 3A and 3B, the prevalence of Black Americans was significantly overestimated in both perception (M = 3.67%, SD = 8.74, t (43) = 2.79, $P = 0.008$, d = 0.42) and in memory (M = 11.26%, SD = 16.13, t (43) = 4.63, $P < 0.001$, d = 0.70). Thus, the overestimation of the prevalence of minority groups occurs even in images of naturalistic scenes one might experience in everyday life.

Experiments 4A and 4B. So far, our experiments demonstrated that we overestimate the prevalence of minority group members both in memory- and in perception-based judgments. In Experiments 4A and 4B, we explored whether the salience of minority group members also leads to prioritization for visual awareness, put simply—whether we are more likely to notice them. Participants in these experiments were students at the Hebrew University of Jerusalem, either Jewish Israelis ($n = 97$, Experiment 4A) or Palestinian Israelis ($n = 80$, Experiment 4B). Participants viewed faces of religious women wearing either a Muslim or a Jewish headscarf, masked using continuous flash suppression [CFS (52)]. In CFS, a target stimulus (here, a Muslim or Jewish woman) is presented to one eye and is initially suppressed from awareness by the simultaneous presentation of rapidly changing masks to the other eye. The dependent variable is the time it takes the target stimulus to "break" suppression (i.e., become available to visual awareness) and be detected by the participant (53). We also included a control condition in which target stimuli were diffeomorphic scrambles of the same faces (*SI Appendix*, Fig. 5) (54), which allows us to control for differences in low-level visual features between the images.

We predicted that, because they are members of the minority group in Israeli society (as well as on campus, see Experiment 1), faces of Muslim women would be more salient and, hence, detected faster. This would not be the case with the control, diffeomorphically scrambled faces.

The results indicate that the faces of Muslim women became available in visual awareness or were consciously noticed significantly faster than those of Jewish women (Experiment 4A: $F(1, 83) = 71.70$, $P < 0.001$, $\eta_p^2 = 0.46$; Experiment 4B: $F(1, 65) = 54.60$, $P < 0.001$, $\eta_p^2 = 0.46$). Importantly, the interaction between face type (Muslim versus Jewish) and image type (normal versus diffeomorphically scrambled) indicated that this effect did not solely originate from low-level visual differences between the face images (Experiment 4A: $F(1, 83) = 20.17$, $P < 0.001$, $\eta_p^2 = 0.20$; Experiment 4B: $F(1, 65) = 6.61$, $P = 0.012$, $\eta_p^2 = 0.09$; *SI Appendix*, Fig. 6). This effect was found among both Jewish and Palestinian participants alike ($F(1, 148) = 1.30$, $P > 0.25$).

Taken together, these results show that, as one would expect given the human mind's general tendency to attend to rare and unexpected events, individuals from a minority group (who are often rare and unexpected) are not only overrepresented in memory-based and perception-based judgments, but this salience starts as early as prioritization for visual awareness. Importantly, this effect occurred in a similar manner for participants who are themselves from the minority group.

Experiment 5 tests the role of social knowledge. We propose that, in addition to the mere statistical prevalence that guides expectations, our knowledge of the social structure of the environment leads us to expect specific social groups to be the minority (33–35, 55). Experiment 5 tests this notion: all experiments so far adopted an ecological perspective and therefore confounded social knowledge with statistical prevalence; Black Americans are the minority in the United States, and Palestinians are the minority in Israel and in all experiments we reported so far as well. However, to test the role of social knowledge, Experiment 5 isolates these two factors: in one condition, Black Americans are the majority of the faces presented in the study, and in the other, White Americans are the majority. If this social knowledge plays no role and only statistical prevalence matters, then we should expect to find symmetrical overestimations between faces of Black Americans and White Americans displayed in the same statistical prevalence. However, if this effect is dependent on our social knowledge, then we do not expect symmetry.

Experiment 5. Participants ($n = 100$, all White American) were randomly assigned to one of two conditions, the Black American minority condition and the White American minority condition. Identically to Experiment 2, the participants in this experiment viewed images of matrices of Black American and White American faces and were asked to estimate the overall prevalence from memory after viewing all images. In the Black American minority condition, 25% of the faces were Black American, while the other 75% were White American. The White American minority condition was the mirror image: 25% White American faces and 75% Black American faces.[#]

The results of this experiment indicate that social knowledge plays a role in the minority salience effect. While both minority overestimation and majority underestimation were significantly different from zero in both conditions, this effect was larger in the Black American minority condition (minority overestimation M = 17.17%, SD = 15.77, t (47) = 7.54, $P < 0.001$, d = 1.09; majority underestimation M = −19.44%, SD = 16.71, t (47) = −8.05, $P < 0.001$, d = 1.16) than in the White minority condition (minority overestimation M = 4.86%, SD = 13.96, t (49) = 2.46, $P = 0.017$, d = 0.35; majority underestimation M = −5.40%, SD = 12.43, t (49) = −3.07, $P = 0.003$, d = 0.43). Independent sample t tests indicated that both the moderation of minority overestimation and of majority underestimation by condition was significant (minority overestimation t (95) = 4.20, $P < 0.001$, $d = 0.85$; majority overestimation t (95) = 4.58, $P < 0.001$, $d = 0.93$; Fig. 3).

These results highlight the social underpinnings of the minority salience effect. We do not merely rely on the statistical prevalence in a given context. Rather, our social knowledge participates in the process and influences the saliency of minority groups. The same statistical prevalence is perceived differently according to the expectations created by our social knowledge.

Taken together, our results indicate that we perceive the composition of our social environments in a biased fashion. Experiment 6 tests whether this bias—perceiving an environment as more diverse than it truly is—leads to decreased support for diversity-promoting policies. In this experiment, the participants learned the same prevalence of a minority group either from descriptive information provided to them directly or in an experiential manner (56), that is, by viewing the faces of individuals and estimating the prevalence of the minority group. The information provided is identical in both conditions. Our hypothesis is simple: participants in the experiential condition are likely to be affected by minority salience and overestimate the prevalence of the minority; this overestimation is likely to inform their support for diversity-promoting policies. Put simply, we predict that minority salience will lead to reduced support toward diversity-promoting policies.[##]

Experiment 6. Participants ($n = 100$) were informed that they would view information about two different college programs that included 2,000 students each. In the experiential condition, the participants were informed that they would see the faces of all the students in this college program and viewed 20 matrices with 100 faces each. The matrices paradigm was identical to the previous experiments, except that, in order to emulate more natural environments (e.g., refs. 57 and 58), the overall prevalence of Black American faces was set to 5%. In the descriptive condition, participants saw a vignette describing the college program, stating that 5% of the 2,000 students in the program were Black Americans. In each condition after viewing the matrices or reading about the program, participants were asked whether this college program should be more motivated to increase the diversity of the student population and adopt diversity-promoting policies (e.g., "In your opinion, how much effort should this program invest in increasing the racial diversity of students?"; not at all [0] to a great extent [100]; for a full list, see *SI Appendix*). The order of the conditions was counterbalanced between participants. We hypothesized that, in the experiential condition, the participants would overestimate the prevalence of Black Americans and therefore be less likely to support diversity-promoting policies.

Indeed, in the experiential condition, the participants overestimated the prevalence of Black Americans (mean estimate = 14.75%, SD = 9.19; t (94) = 10.34, $P < 0.001$, d = 1.06) and underestimated the prevalence of White Americans (mean estimate = 83.26%, SD = 10.31; t (94) = −11.73, $P < 0.001$, d = 1.14). Moreover, the results of a repeated measures ANOVA indicated that, in the experiential condition, the participants reported less support for diversity-promoting policies (M = 71.07, SD = 23.03) than in the descriptive condition (M = 74.50, SD = 23.10; $F(1, 94) = 7.67$, $P = 0.007$, $\eta_p^2 = 0.075$). This effect did not significantly interact with the order of conditions ($F < 1$), and a supplementary analysis shows that it was also evident in a between-participants comparison including only the first condition displayed for each participant

[#]Preregistration available at https://aspredicted.org/blind.php?x=/N78_6B8.

[##]Preregistration available at https://aspredicted.org/e2zk5.pdf.

PSYCHOLOGICAL AND COGNITIVE SCIENCES

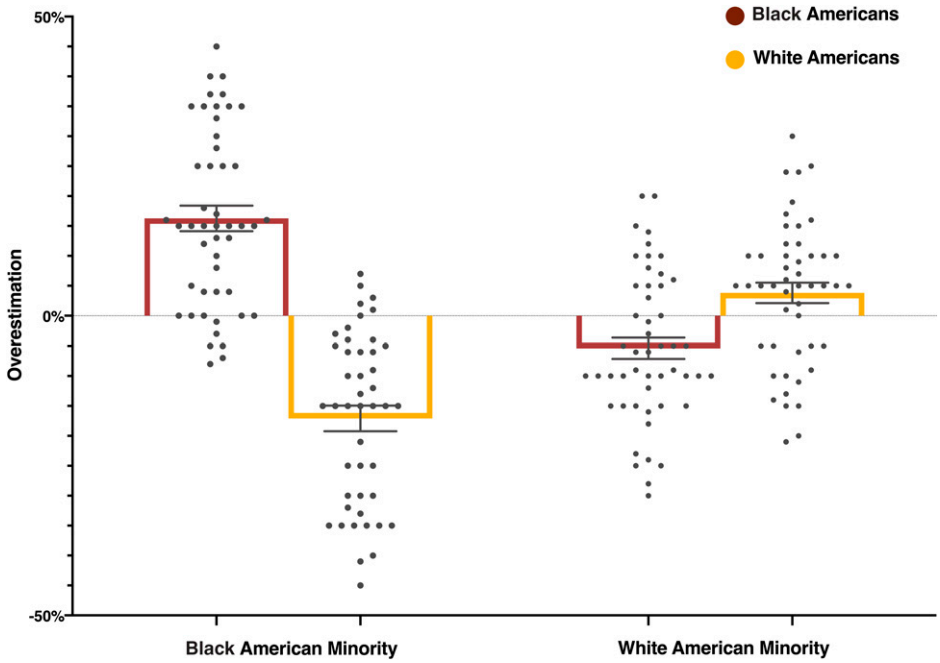

Fig. 3. The results of Experiment 5. Zero denotes accurate estimates, positive values denote overestimation, and negative values underestimation. In the Black American minority condition, the overall percentage of Black American faces was 25%. In the White American minority condition, the overall percentage of Black American faces was 75%. In both conditions, the participants were asked to estimate the prevalence of Black American faces (red) and White American faces (yellow). In both conditions, the participants overestimated the minority group and underestimated the majority group. However, the magnitude of the effects in the White American minority condition was significantly smaller, denoting the role of prior social knowledge.

(Experiential: M = 65.71; SD = 24.71; Descriptive: M = 78.02, SD = 19.95; $F(1, 93) = 7.18$, $P = 0.009$, $\eta_p^2 = 0.072$).

Moreover, the participants' overestimation significantly correlated with the difference between their support for diversity-promoting policies between conditions ($r = 0.34$, $P < 0.001$). The larger the overestimation, the larger the difference was in their attitudes between the two conditions. This same pattern was also found for the estimates of White Americans ($r = -0.30$, $P = 0.003$): the less the participants underestimated the prevalence of White Americans in the experiential condition, the more they supported diversity-promoting policies in the experiential condition compared to the descriptive condition.

These results indicate that the biased perception of an environment as more diverse than it truly is has consequences—it leads to decreased support for diversity-promoting policies.

A Note on Attitudes. In order to test whether minority salience is more or less likely to occur for people holding different attitudes or abiding with different ideologies, in all experiments, we also examined participants' social attitudes and attitudes toward each of the tested groups (i.e., Black and White Americans in American samples and Palestinian and Jewish Israelis in the Israeli samples; for a full list of measures, see *SI Appendix*). None of the tested attitudes consistently correlated with the estimates of minority and majority groups. For example, in a combined analysis of Experiments 2 and 3 ($n = 300$), the estimates did not correlate with explicit preference ($r = 0.01$), perceived threat ($r = 0.07$), or political ideology ($r = 0.08$). The only exception was the social dominance orientation (SDO) scale (59, 60), with a modest correlation of $r = 0.15$, $P = 0.006$ (for an elaboration on this, see *SI Appendix*).

These are all null effects on exploratory measures, and as such, they should be cautiously interpreted. However, they may suggest that social attitudes, at least as they are measured here, do not play a role in determining minority salience and overestimation. Further evidence for this interpretation comes from the finding of no differences between majority and minority participants in Experiments 1 through 4. Given that majorities and minorities typically hold different opinions about their groups (e.g., refs. 61 and 62), it suggests attitudes do not play a significant role in bringing about minority salience.

We propose that this may be the case because the effect is driven by the basic architecture of our cognitive system. Regardless of our own attitudes or our ingroup affiliation, unexpected stimuli are more salient, and minority groups are unexpected for all of us since we share the same social knowledge (63). Overall, 82.6% of the participants tested throughout these experiments overestimated the prevalence of minorities. Simply put, these results suggest that this effect occurs for almost all of us.

Seemingly contradictory, in Experiment 6, the overestimation of the minority group predicted less support for diversity-increasing measures in that specific social context (i.e., the college program). Likewise, previous studies on national polls established that a bias in perception of society as a whole (overestimating the size of a national minority group) was related to attitudes toward society as a whole (41–44, 64). The key difference between those two findings and the ones reported in this section is that, in this section, we did not test attitudes relating to the perceived social context but rather general attitudes about society.

Overall, then, we suggest that while minority salience is not linked to (or resulting from) general social attitudes, it does bring about attitudes that pertain to the diversity of the relevant social context.

To further test this suggestion, we ran a preregistered study** on 150 participants (see details in *SI Appendix*; Study SI1) testing the associations between minority salience and specific and general social attitudes. As in the experiential condition in Experiment 6, participants viewed 20 matrices that were composed of 5% Black American faces and were informed that

**Preregistration available at https://aspredicted.org/blind.php?x=/DR1_6CY.

these are faces of students from a college. We then asked participants about their attitudes toward diversity-promoting policies in this college and about their general social attitudes: SDO and liberal–conservative ideologies. Replicating our previous results, we found that general attitudes were not correlated with either an overestimation of Black Americans (SDO: $r = 0.05$, $P > 0.5$; Ideology: $r = 0.04$, $P > 0.6$) or underestimation of White Americans (SDO: $r = -0.07$, $P > 0.4$; Ideology: $r = -0.08$, $P > 0.3$). However, attitudes toward diversity-promoting policies in this college did correlate both with the overestimation of Black Americans ($r = -0.17$, $P = 0.046$) and the underestimation of White Americans ($r = 0.22$; $P = 0.007$). The more a participant overestimated the prevalence of Black Americans and underestimated the prevalence of White Americans, the less they supported diversity-promoting policies in this college. Moreover, in a regression model predicting support for diversity-promoting policies in this college, the effect of overestimation of Black Americans ($\beta = -0.33$, t (139) $= -2.28$, $P = 0.024$) was significant beyond the effects of SDO ($\beta = -0.12$, t (139) $= -0.69$, $P = 0.49$) and ideologies ($\beta = -0.66$, t (139) $= -7.70$, $P < 0.001$). Similarly, an underestimation of White Americans ($\beta = 0.35$, t (146) $= 2.80$, $P = 0.006$) significantly predicted support for diversity-promoting policies in this college beyond the effects of SDO ($\beta = -0.27$, t (146) $= -1.56$, $P = 0.12$) and ideologies ($\beta = -0.54$, t (146) $= -6.86$, $P < 0.001$).

These results show again that general social attitudes are not related to the biased perception of diversity of a social context. However, the degree to which a social context is incorrectly perceived as diverse is related to support for diversity-promoting policies in this context, over and above the influence of one's general social attitudes. These data are consistent with previous findings in the literature on national polls in which a bias in perception of society as a whole (the size of a national minority group) was related to attitudes toward society as a whole.

Taken together, our results thus paint a more complete picture. On the one hand, social attitudes do not seem to play a role in creating minority salience—people who hold different social attitudes do not perceive the environment differently. However, minority salience, a biased perception of an environment, leads to decreased support for diversity-promoting policies in this environment.

Discussion

Taken together, our results from 12 experiments and 942 participants indicate that minority salience and overestimation are robust phenomena. We consistently overestimate the prevalence of individuals from minority groups and underestimate the prevalence of members from the majority group, thus perceiving our social environments as more diverse than they truly are. Our experiments also indicate that this effect may be found at the level of priority for visual awareness and that it is social in nature: our social knowledge, our representation of the overall composition of our social environment, shapes this effect. Importantly, this illusion of diversity is consequential in that it leads to less support for measures to increase diversity.

"Subjective probabilities play an important role in our lives … The decisions we make, the conclusions we reach, and the explanations we offer are usually based on our judgments of the likelihood of uncertain events" [Kahneman and Tversky (65)]. If our minds create an illusion of diversity, making our environments seem more diverse than they actually are, this illusion will be the basis for our decisions. Indeed, our experiments demonstrate that minority salience affects support for diversity-promoting policies. This suggests that the erroneous perception of minorities as more prevalent in our immediate environments, our campuses, workplaces, and neighborhoods leads to similar outcomes. Previous research also lends support to this idea at a national level: the perceived size of a national minority is associated with a host of adverse outcomes (41–44, 64). Given the importance of this topic, we believe future research should further examine how this bias shapes attitudes and policies and design interventions that can correct for it in order to inform public discussions and promote more diverse and inclusive environments.

Our findings indicate that both the statistical prevalence of minority groups in a given context and our understanding of the general structure of our society play a role in creating minority salience. However, beyond knowledge of our society as a whole, we also have knowledge of specific social contexts which should lead us to anticipate that certain groups will be less prevalent in certain contexts (e.g., Black Americans in higher education). In these contexts, the presence of members from the minority group who break this expectation may be even more salient, and hence, their prevalence may be overestimated even more. Future research should therefore examine whether the same actual prevalence of a minority group is perceived differently in different social contexts according to the expectations created by our social knowledge regarding these contexts. Notably, the groups we examined are not only minorities; they represent marked social identities [i.e., social identities that are perceived as atypical compared to unmarked "default" social identities (66–68)]. This marked status may exacerbate the salience of a minority group in addition to their statistical prevalence and our social knowledge. Future empirical work will be required to examine this suggestion.

In addition, our experiments tested the minority salience effect exclusively on ethnic minorities (Palestinian Israelis and Black Americans); however, it is important to note that the mechanism extends beyond ethnic minorities. Based on the same theoretical principles, we would predict that similar overestimation can be found in any social context in which one group is a statistical minority or in which social knowledge creates expectations that a group would be a minority (e.g., women in natural sciences, or technological professions). Likewise, minority salience may also operate at the level of public discourse, with minority opinions being more salient and thus having a larger proportional impact [e.g., the antivaccine movement (69)].

Finally, in line with previous research on decision-making (56, 70), our results suggest that the manner in which we obtain our knowledge about the prevalence of a minority group matters. In our case, when people obtain this knowledge experientially, they end up with biased perceptions. This highlights the importance of institutions providing descriptive knowledge—the true prevalence of minority groups—in order to allow individuals to form well-informed decisions and opinions. It seems likely that, often, decision-makers base their decisions on data available to them: the actual prevalence of social groups in the relevant environment. It is clear from our results that if they do not, they should. However, even so, it is important that they take the illusion of diversity into consideration—as it can automatically affect their less deliberate intuitions, decisions, and judgments. Research on bias correction (71) suggests that while these sorts of biases can be counteracted, one must first be aware of the bias and make an effort to correct it.

This conclusion applies even more for those of us who do not have access to the relevant data. When we think about how diverse our campuses are, or our scientific fields, work environments, or communities, we rarely know the actual data. The illusion of diversity therefore informs our intuitions, judgments, and decisions and may affect crucial aspects such as how we vote on by-laws and whether we support new policies.

Materials and Methods

All studies were approved by the Hebrew University Institutional Review Board, and all participants gave their consent to participate.

Preregistration Documents. The procedures, sample size, exclusion criteria, and data analysis plans for Experiments 1 to 3 and 5 to 6 were preregistered prior to data collection. Experiment 4 was not preregistered. However, our analysis protocol is identical to our previous published work (72).

Experiments 1A and 1B. A total of 32 Jewish Israeli students (69.7% female, mean age = 24.94) participated in Experiment 1A, and 33 (79.4% female, mean age = 21.82) Palestinian Israeli students participated in Experiment 1B. All participants were students at the Hebrew University of Jerusalem. Participants completed an online questionnaire, presented in Hebrew in Experiment 1A and in Arabic in Experiment 1B. The questionnaire directed students to remember walking up the main hallway of the Hebrew University campus (Mount Scopus) followed by questions about the typical experience (e.g., "Is there a long line for the cafeteria?" and "Out of all the people you encounter along the walk, how many are faculty members?"). Afterward, participants were asked to estimate "Out of all the people you encounter along the walk, how many are Arab" on a 0 to 100% scale. Participants then completed six filler questions (e.g., "What percent of the Hebrew university students are blond?") and were then asked to estimate the percentage of Arab students at the university followed by two additional filler questions and one attention check question. The actual percentage of Arab students at the university was obtained from the Hebrew University student enrollment offices. According to the preregistered plan, participant recruitment lasted for 2 wk or until 30 participants were recruited for each experiment. Participants who failed to answer the attention check (none in 1A; $n = 5$ in 1B) as well as estimates that were more than two SDs larger or smaller than the group mean (none in 1A; $n = 2$ in 1B) were excluded from analysis.[††]

Experiments 2A and 2B. Experiments 2A and 2B were identical, except for the participant pool. A total of 100 White American participants (65.0% female, mean age = 37.40) participated in experiment 2A, and 100 Black American participants (41.4% female, mean age = 31.45) participated in Experiment 2B. Participants were informed that they would be shown photos of American college students. Participants viewed matrices of the faces of Black American and White American individuals (Fig. 1). Each matrix included 100 faces. The matrices were generated using a Python 3 script randomly selecting from 330 images from the Chicago Face Database (73); all pictures were rated by above 90% of the participants in a norms study (73) as belonging to an Black American or a White American, and all displayed neutral facial expressions. The overall percentage of Black American faces was set to 25% in one condition and 45% in the second condition. In the 25% condition, the percentage of Black American faces in each matrix was either 10, 20, 30, or 40% in equal probability. In the 45% condition, the percentages in each matrix were either 30, 40, 50, or 60% with equal probability. The locations of the faces in each matrix were randomly distributed. Participants viewed 20 matrices, each displayed for 2 s, and were asked to indicate for each matrix whether it included mostly women's or men's faces (all matrices included 50% women faces). The images were displayed, and participant responses were recorded using the Gorilla online experiments framework (74). After viewing all matrices, participants were asked to estimate the percentage of African American faces they had viewed throughout the study and the percentage of White American faces. After completing the study, the participants were asked a number of questions about their social attitudes including explicit attitudes, group thermometer, perceived group threat, and SDO (*SI Appendix*). According to our preregistered analysis plan, estimates that are two SDs larger or smaller than the group mean were excluded from analysis (2A: two estimates of Black Americans and three of White Americans; 2B: three estimates of Black Americans, four of White Americans).

Experiments 3A, 3B, and 3C. In experiments 3A and 3B, participants were asked to provide estimates immediately after viewing each matrix. Participants (all based in the United States) were recruited through Prolific; there were 50 White American participants in Experiment 3A (70% female, mean age 34.88) and 50 Black American participants in Experiment 3B (42% female, mean age 31.56). Each participant viewed 100 matrices that included either 10, 20, 30, 40, or 50% Black American faces with equal probability for an overall percentage of 30% Black American faces. Unlike Experiments 2A and 2B, participants were asked to provide an estimate immediately after each matrix was displayed for 2 s. To prevent participants from trying to count or use other strategies (e.g., count the first row then multiply by 10), 50% of the trials were filler trials; after viewing each matrix, participants were asked to estimate one of four categories: Black Americans, White Americans, women (filler), or men (filler). The trials on which participants estimated the percentage of women/men were discarded from analysis. Therefore, for each participant, there were 50 critical trials: 25 estimates of Black Americans (five of each actual percentage) and 25 estimates of White Americans (five of each actual percentage condition). For each matrix, we calculated the difference between a participants' estimation and the actual percentage. According to our preregistered analysis plan, mean estimates that deviated by more than two SDs from the group's average for the actual percentage condition were discarded from analysis (5% of estimates in Experiment 1A; 7% in Experiment 1B). Identically to Experiments 2A and 2B, after all matrices were displayed, the participants were asked for their estimates regarding the overall percentage of Black American and White American faces displayed throughout the experiment as well as the attitude measures. According to our preregistered analysis plan, estimates that deviated by more than two SDs from the group's average were discarded from analysis (3A: two estimates of Black Americans and two of White Americans; 3B: three estimates of Black Americans, two of White Americans).

Experiment 3C employed the same overall design as Experiments 3A and 3B with the following modifications. Only nine images were included as stimuli. These images were photos of naturalistic scenes which included large crowds (at least 15 people) with both Black American individuals (between 10 to 65% in each image, mean 26.49%) and White American individuals. After viewing each image for 2 s, participants were asked to estimate only the percentage of Black American individuals in the image. After viewing all the images, participants were also asked to estimate the overall percentage of Black American individuals in the entire set of images from memory. According to our preregistration plan, six participants were excluded for failing attention checks, no estimates deviated over two SD from the group mean.

Experiments 4A and 4B. A total of 97 Jewish Israeli participants took part in Experiment 4A. A total of 80 Palestinian Israeli participants took part in Experiment 4B. A total of 17 faces of Muslim women (identified by wearing a Muslim headscarf) and 17 faces of Jewish women (identified by wearing a Jewish headscarf) were used, all in grayscale. To control for possible low-level visual differences in the stimuli, we also included diffeomorphic scrambles (54) of all 34 images. Both experiments were conducted in the Hebrew University of Jerusalem. Experiment 1A ran in Hebrew, Experiment 1B ran in Arabic. On each of the CFS trials, a fixation cross was presented binocularly at the center of each eye's visual field. The target stimuli (faces or diffeomorphic scrambles) were presented monocularly and gradually ramped up in contrast (from 0 to 50%) during the first second of presentation. The masks were patterns of randomly assigned colored squares, changing randomly at a rate of 10 Hz (*SI Appendix*, Fig. 5). On each trial, the mask was randomly presented to one eye, and the target stimulus was presented to the other eye. The target stimuli appeared either below or above fixation (probability = 0.5). The participants' task was to indicate whether they appeared above or below fixation by pressing the appropriate key (for similar procedure, see ref. 75). Participants were instructed to respond as quickly as they could. The time between the start of each trial and the participant's response (reaction time) was measured. The stimuli were displayed in interchangeable blocks of either faces or scrambles, and the order of blocks was counterbalanced between participants. Each participant saw six blocks of faces and six blocks of scrambles. The final design being a within-subject 2 (face: Muslim/Jewish) × 2 (image type: regular/scramble). Only data from trials in which the response was correct were included in the analysis (92.66 and 93.90% of trials in Experiments 4A and 4B, respectively); in addition, trials on which the reaction time deviated from the participant's mean reaction time by more than three SDs were discarded from analysis (2.1 and 2.2% in Experiments 4A and 4B, respectively). Finally, the data from participants who responded correctly on less than 90% of the trials (13 participants in each of Experiments 4A and 4B, respectively) as well as participants whose average reaction time deviated from the group mean by more than three SDs (0 and 1 participant in Experiments 4A and 4B, respectively) were excluded from analysis (72). At the end of the experiment, participants completed explicit and implicit measures testing their attitudes toward Muslims and Jews (see *SI Appendix* for an elaboration on this).

Experiment 5. Identically to Experiment 2, the participants ($n = 100$, all White Americans, recruited through Prolific, 62% female, mean age 34.39) in this experiment viewed matrices of Black American and White American faces and were asked to estimate the overall prevalence of Black American and White

[††]This exclusion rule was not preregistered but was implemented once data were available as upon visual inspection there were clear outliers in the data. Analyses of the full dataset without exclusion indicated a stronger overestimation effect.

American faces from memory after viewing all images. Participants were randomly assigned to one of two conditions—the Black American minority condition and the White American minority condition. In the Black American minority condition, 25% of the faces were Black American, while the other 75% were White American (matrices had either 10, 20, 30, or 40% Black American faces with equal probability). The White American minority condition was the mirror image; 25% White American faces and 75% Black American faces (matrices had either 60, 70, 80, or 90% Black American faces with equal probability). After viewing all images, participants were asked to estimate the overall percentage of the minority group, the majority group, and answer the social attitudes questionnaire used in previous experiments.

Experiment 6. In experiment 6, 100 participants, all based in the United States (85.3% female, mean age 34.31), were recruited through Prolific. Each participant completed two within-participant conditions. In the experiential condition, the participants were instructed that they would see the faces of all the students in one college program and then viewed 20 matrices with 100 faces in each matrix. The matrices paradigm was identical to that of Experiment 2, except with a 5% overall prevalence of Black American faces (half the matrices included 4% Black American face, the other half 6%). After viewing all matrices, the participants were asked to estimate the prevalence of Black American faces and of White American faces. Afterward, they were asked four questions about their attitudes on whether this college program should be more motivated to increase the racial diversity of the student population and adopt diversity-promoting policies. The responses to these questions were averaged (alpha Cronbach > 0.88) and served as the dependent variable. In addition, the participants were asked three filler questions about policies to promote gender diversity. In the descriptive condition, the participants saw a vignette describing a different college program. The vignette described the population of students as 50% women, 5% Black American, and 95% White American. This description was identical to the prevalence of the social categories displayed in the matrices. After reading the vignette, the participants were asked the same questions as in the experiential condition. The viewing order of the conditions was counterbalanced between participants. In line with our preregistered analysis plan, five participants whose estimates of Black Americans or White Americans fell two SDs outside the group mean were disregarded from the analysis.

Data Availability. Anonymized behavioral data have been deposited in Open Science Framework (OSF) at https://osf.io/e6c7h/?view_only=7027350b69a84706a7381a97bcd6707e.

1. M. A. Craig, J. M. Rucker, J. A. Richeson, The pitfalls and promise of increasing racial diversity: Threat, contact, and race relations in the 21st century. *Curr. Dir. Psychol. Sci.* **27**, 188–193 (2018).
2. X. Bai, M. R. Ramos, S. T. Fiske, As diversity increases, people paradoxically perceive social groups as more similar. *Proc. Natl. Acad. Sci. U.S.A.* **117**, 12741–12749 (2020).
3. S. S. Levine *et al.*, Ethnic diversity deflates price bubbles. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 18524–18529 (2014).
4. A. D. Galinsky *et al.*, Maximizing the Gains and Minimizing the Pains of Diversity: A Policy Perspective. *Perspect. Psychol. Sci.* **10**, 742–748 (2015).
5. C. Herring, Does diversity pay?: Race, gender, and the business case for diversity. *Am. Sociol. Rev.* **74**, 208–224 (2009).
6. C. Wacongne, J. P. Changeux, S. Dehaene, A neuronal model of predictive coding accounting for the mismatch negativity. *J. Neurosci.* **32**, 3665–3678 (2012).
7. A. Treisman, "Consciousness and perceptual binding" *in* The Unity of Consciousness: Binding, Integration, and Dissociation, A. Cleeremans, C. Frith, Eds. (Oxford University Press, 2003), pp. 95–113. 10.1093/acprof:oso/9780198508571.003.0005.
8. L. Itti, P. Baldi, Bayesian surprise attracts human attention. *Vision Res.* **49**, 1295–1306 (2009).
9. M. R. P. LaPointe, B. Milliken, Semantically incongruent objects attract eye gaze when viewing scenes for change. *Vis. Cogn.* **24**, 1–15 (2016).
10. E. S. Spelke, "Preferential-looking methods as tools for the study of cognition in infancy" in *Measurement of Audition and Vision in the First Year of Postnatal Life: A Methodological Overview, G. Gottlieb, N. A. Krasnegor, Eds.* (Ablex, Norwood, NJ, 1985), pp. 323–363.
11. A. Treisman, Features and objects in visual processing. *Sci. Am.* **255**, 114B–125 (1986).
12. B. McElree, M. Carrasco, The temporal dynamics of visual search: evidence for parallel processing in feature and conjunction searches. *J. Exp. Psychol. Hum. Percept. Perform.* **25**, 1517–1539 (1999).
13. C. Stangor, D. McMillan, Memory for expectancy-congruent and expectancy-incongruent information: A review of the social and social developmental literatures. *Psychol. Bull.* **111**, 42–61 (1992).
14. J. Schomaker, M. Meeter, Short- and long-lasting consequences of novelty, deviance and surprise on brain and cognition. *Neurosci. Biobehav. Rev.* **55**, 268–279 (2015).
15. K. Friston, The free-energy principle: A unified brain theory? *Nat. Rev. Neurosci.* **11**, 127–138 (2010).
16. S. Kanaya, M. J. Hayashi, D. Whitney, Exaggerated groups: Amplification in ensemble coding of temporal and spatial features. *Proc. Biol. Sci.* **285**, 20172770 (2018).
17. M. W. Khaw, R. Kranton, S. Huettel, Oversampling of minority categories drives misperceptions of group compositions. *Cognition* **214**, 104756 (2021).
18. D. Landy, B. Guay, T. Marghetis, Bias and ignorance in demographic perception. *Psychon. Bull. Rev.* **25**, 1606–1618 (2018).
19. D. L. Hamilton, R. K. Gifford, Illusory correlation in interpersonal perception: A cognitive basis of stereotypic judgments. *J. Exp. Soc. Psychol.* **12**, 392–407 (1976).
20. A. Tversky, D. Kahneman, Availability: A heuristic for judging frequency and probability. *Cogn. Psychol.* **5**, 207–232 (1973).
21. A. Tversky, D. Kahneman, Advances in prospect theory: Cumulative representation of uncertainty. *J. Risk Uncertain.* **5**, 297–323 (1992).
22. J. Santeau, When does a response error become a judgmental bias? Commentary on "Judged Frequency of Lethal Events". *J. Exp. Psychol. Hum. Learn.* **4**, 579–581 (1978).
23. S. E. Taylor, S. T. Fiske, Salience, attention, and attribution: Top of the head phenomena. *Adv. Exp. Soc. Psychol.* **11**, 249–288 (1978).
24. S. E. Taylor, S. T. Fiske, Point of view and perceptions of causality. *J. Pers. Soc. Psychol.* **32**, 439–445 (1975).
25. S. E. Taylor, S. T. Fiske, *Social Cognition: From Brains to Culture* (Ref. Res. B. News, ed. 2, 2013).
26. G. D. Lassiter, Illusory causation in the courtroom. *Curr. Dir. Psychol. Sci.* **11**, 204–208 (2002).
27. P. Slovic, Perception of risk. *Science* **236**, 280–285 (1987).
28. S. V. Eisen, L. Z. McArthur, Evaluating and sentencing a defendant as a function of his salience and the perceiver's set. *Pers. Soc. Psychol. Bull.* **5**, 48–52 (1979).
29. S. T. Fiske, Attention and weight in person perception: The impact of negative and extreme behavior. *J. Pers. Soc. Psychol.* **38**, 889–906 (1980).
30. D. L. Hamilton, S. J. Sherman, "Illusory correlations: Implications for stereotype theory and research" in *Stereotyping and Prejudice, D. Bar-Tal, C. F. Grauman., A. W. Kruglanski, W. Stroebe, Eds.*(Springer, New York, NY, 1989), pp. 59–82.
31. C. Hoffman, N. Hurst, Gender stereotypes: Perception or rationalization? *J. Pers. Soc. Psychol.* **58**, 197–209 (1990).
32. D. L. Hamilton, T. L. Rose, Illusory correlation and the maintenance of stereotypic beliefs. *J. Pers. Soc. Psychol.* **39**, 832–845 (1980).
33. S. T. Fiske, S. L. Neuberg, A continuum of impression formation, from category-based to individuating processes: Influences of information and motivation on attention and interpretation. *Adv. Exp. Soc. Psychol.* **23**, 1–74 (1990).
34. D. I. Tamir, M. A. Thornton, Modeling the predictive social mind. *Trends Cogn. Sci.* **22**, 201–212 (2018).
35. M. Otten, A. K. Seth, Y. Pinto, A social Bayesian brain: How social knowledge can shape visual perception. *Brain Cogn.* **112**, 69–77 (2017).
36. R. Nadeau, R. G. Niemi, J. Levine, Innumeracy about minority populations. *Public Opin. Q.* **57**, 332–347 (1993).
37. C. J. Wong, "Little" and "Big" pictures in our heads. Race, local context, and innumeracy about racial groups in the United States. *Public Opin. Q.* **71**, 392–412 (2007).
38. J. Citrin, J. Sides, Immigration and the imagined community in Europe and the United States. *Polit. Stud.* **56**, 33–56 (2008).
39. H. Daniel, How many immigrants? Foreign-born population innumeracy in Europe. *Public Opin. Q.* **74**, 674–695 (2010).
40. I. Mori, Perils of Perception 2015 (2016). https://www.ipsos.com/sites/default/files/2016-12/Perils-of-perception-2016.pdf (Accessed 25 September 2018)
41. R. Alba, R. G. Rumbaut, K. Marotz, A distorted nation: Perceptions of racial/ethnic group sizes and attitudes toward immigrants and other minorities. *Soc. Forces* **84**, 901–919 (2005).
42. E. Schlueter, P. Scheepers, The relationship between outgroup size and anti-outgroup attitudes: A theoretical synthesis and empirical test of group threat- and intergroup contact theory. *Soc. Sci. Res.* **39**, 285–295 (2010).
43. M. A. Craig, J. A. Richeson, More diverse yet less tolerant? How the increasingly diverse racial landscape affects white Americans' racial attitudes. *Pers. Soc. Psychol. Bull.* **40**, 750–761 (2014).
44. M. A. Craig, J. M. Rucker, J. A. Richeson, Racial and political dynamics of an approaching "majority-minority" United States. *Ann. Am. Acad. Pol. Soc. Sci.* **677**, 204–214 (2018).
45. M. A. Craig, J. A. Richeson, On the precipice of a "majority-minority" America: Perceived status threat from the racial demographic shift affects White Americans' political ideology. *Psychol. Sci.* **25**, 1189–1197 (2014).
46. M. A. Craig, J. A. Richeson, Hispanic population growth engenders conservative shift among non-Hispanic racial minorities. *Soc. Psychol. Personal. Sci.* **9**, 383–392 (2018).
47. M. Hooghe, T. de Vroome, The perception of ethnic diversity and anti-immigrant sentiments: a multilevel analysis of local communities in Belgium. *Ethn. Racial Stud.* **38**, 38–56 (2015).
48. Z. Strabac, It is the eyes and not the size that matter. *Eur. Soc.* **13**, 559–582 (2011).
49. M. Semyonov, R. Raijman, A. Y. Tov, P. Schmidt, Population size, perceived threat, and exclusion: A multiple-indicators analysis of attitudes toward foreigners in Germany. *Soc. Sci. Res.* **33**, 681–701 (2004).

50. M. A. Craig, J. A. Richeson, Majority no more? The influence of neighborhood racial diversity and salient national population changes on Whites' perceptions of racial discrimination. *RSF* **4**, 141–157 (2018).
51. T. D. Snyder, C. de Brey, S. A. Dillow, *Digest of Education Statistics 2018 (NCES 2020-009)* (National Center for Education Statistics, 2019).
52. N. Tsuchiya, C. Koch, Continuous flash suppression reduces negative afterimages. *Nat. Neurosci.* **8**, 1096–1101 (2005).
53. T. Stein, M. N. Hebart, P. Sterzer, Breaking continuous flash suppression: A new measure of unconscious processing during interocular suppression? *Front. Hum. Neurosci.* **5**, 167 (2011).
54. B. Stojanoski, R. Cusack, Time to wave good-bye to phase scrambling: Creating controlled scrambled images using diffeomorphic transformations. *J. Vis.* **14**, 6 (2014).
55. N. Reggev, A. Chowdhary, J. P. Mitchell, Confirmation of interpersonal expectations is intrinsically rewarding. *Soc. Cogn. Affect. Neurosci.* **16**, 1276–1287 (2021).
56. R. Hertwig, G. Barron, E. U. Weber, I. Erev, Decisions from experience and the effect of rare events in risky choice. *Psychol. Sci.* **15**, 534–539 (2004).
57. Yale University Office of Institutional Research, "Yale University Common Data Set" (2021). https://oir.yale.edu/sites/default/files/factsheet_2021_vf_04062021.pdf. Accessed 21 February 2022.
58. Office of Planning and Institutional Research, Columbia University, "Headcount Enrollment by School, Race/Ethnicity, and Citizenship, Fall 2020" (2020) https://opir.columbia.edu/sites/default/files/content/Statistical Abstract/opir_enrollment_ethnicity.pdf. Accessed 21 February 2022.
59. F. Pratto, J. Sidanius, L. M. Stallworth, B. F. Malle, Social dominance orientation: A personality variable predicting social and political attitudes. *J. Pers. Soc. Psychol.* **67**, 741–763 (1994).
60. A. K. Ho *et al.*, The nature of social dominance orientation: Theorizing and measuring preferences for intergroup inequality using the new SDO_7 scale. *J. Pers. Soc. Psychol.* **109**, 1003–1028 (2015).
61. B. A. Nosek, M. R. Banaji, A. G. Greenwald, Harvesting implicit group attitudes and beliefs from a demonstration web site. *Gr. Dyn.* **6**, 101–115 (2002).
62. S. Smooha, *Index of Arab-Jewish Relations in Israel, 2003-2009* (Jewish-Arab Center, Haifa University, Haifa, 2010).
63. P. G. Devine, Stereotypes and prejudice: Their automatic and controlled components. *J. Pers. Soc. Psychol.* **56**, 5–18 (1989).
64. J. T. Pickett, T. Chiricos, K. M. Golden, M. Gertz, Reconsidering the relationship between perceived neighborhood racial composition and whites' perceptions of victimization risk: Do racial stereotypes matter? *Criminology* **50**, 145–186 (2012).
65. D. Kahneman, A. Tversky, Subjective probability: A judgment of representativeness. *Cogn. Psychol.* **3**, 430–454 (1972).
66. W. Brekhus, A sociology of the unmarked: Redirecting our focus. *Sociol. Theory* **16**, 34–51 (1998).
67. W. Brekhus, Social marking and the mental coloring of identity: Sexual identity construction and maintenance in the United States. *Sociol. Forum* **11**, 497–522 (1996).
68. B. Urciuoli, The irony of diversity numbers. *Signs Soc. (Chic.)* **6**, 88–110 (2018).
69. P. Ball, Anti-vaccine movement could undermine efforts to end coronavirus pandemic, researchers warn. *Nature* **581**, 251 (2020).
70. G. Barron, I. Erev, Small feedback-based decisions and their limited correspondence to description-based decisions. *J. Behav. Decis. Mak.* **16**, 215–233 (2003).
71. T. D. Wilson, N. Brekke, Mental contamination and mental correction: Unwanted influences on judgments and evaluations. *Psychol. Bull.* **116**, 117–142 (1994).
72. A. Y. Sklar *et al.*, Reading and doing arithmetic nonconsciously. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 19614–19619 (2012).
73. D. S. Ma, J. Correll, B. Wittenbrink, The Chicago face database: A free stimulus set of faces and norming data. *Behav. Res. Methods* **47**, 1122–1135 (2015).
74. A. L. Anwyl-Irvine, J. Massonnié, A. Flitton, N. Kirkham, J. K. Evershed, Gorilla in our midst: An online behavioral experiment builder. *Behav. Res. Methods* **52**, 388–407 (2020).
75. A. Y. Sklar *et al.*, Did you see it? Robust individual differences in the speed with which meaningful visual stimuli break suppression. *Cognition* **211** (2021). 10.1016/j.cognition.2021.104638.