Insta-learn:

Face Stereotypes Emerge and Persist through Insufficient Statistical Learning

Xuechunzi Bai¹, Stefan Uddenberg², Brandon P. Labbree, Alexander Todorov²

¹Department of Psychology, Princeton University

²The University of Chicago Booth School of Business

Word count: 11751

Correspondence to: xb2@princeton.edu

Abstract (218/250 words)

Social stereotypes are prevalent and consequential, yet sometimes inaccurate. How do people learn these inaccurate beliefs in the first place and why do these beliefs persist in the face of counter evidence? Building on past research on cognitive limitations and environmental sample biases, we propose an integrative perspective: Insufficient statistical learning (Insta-learn). Insta-learn posits that humans are active learners of the environment. Starting from a small sample, people are able to extract statistical patterns within the sample accurately and quickly. However, people do not continue sampling sufficiently. If they decide not to collect more samples once they are (prematurely) satisfied, inaccurate stereotypes can emerge even when more data would show otherwise. We investigated this hypothesis across six online experiments ($N = 1565$), using novel pairs of computer-generated faces and social behaviors. Fixing the population level statistics of face-behavior associations to zero and varying the initial sample statistics, we found that participants quickly learned the initial sample statistics (from as few as three examples) and persisted in using such spurious associations in their final decisions. Granting the sampling power to participants — samples were endogenously generated by participants and not defined by the experimenters — we found insufficient sampling caused spurious associations to persist. Insta-learn provides a domain-general framework for a mechanistic explanation of the emergence and persistence of social stereotypes.

Keywords: social stereotypes, face perception, over-generalization, statistical learning, endogenous sampling

All they have done for me is to make me a little more conscious of how hard it is to classify and to sample, how readily we spread a little butter over the whole universe.

- Walter Lippmann, 1922, Public Opinion, p. 104

Social stereotypes are prevalent. People readily associate groups of individuals' external features with internal traits (Allport, 1954; Hamilton & Sherman, 1994). External features include social categories such as gender and race, as well as appearance based on one's face or clothes. Internal traits include dominance or competence, trustworthiness or warmth, among many others associated with personality. Stereotypes such as "females are more likely to be friendly" or "feminine-looking faces are more trustworthy" often operate in people's daily judgments (Fiske, Cuddy, Glick, & Xu, 2002; Fiske & Bai, 2020; Oosterhof & Todorov, 2008; Todorov & Oh, 2020).

Social stereotypes are consequential. People not only form stereotypes about various individuals, but also use stereotypes in important decisions. For example, the stereotype associating women with lower productivity may lead to hiring discrimination in labor markets (Marianne & Mullainathan, 2004). The stereotype associating African descent with higher criminality may lead to fatal police shootings and severe sentencing in criminal justice (Eberhardt, Davies, Purdie-Vaughns, & Johnson, 2006). And the stereotype associating particular facial features with higher competence may lead to more votes in real political elections (Todorov, Mandisodza, Goren, & Hall, 2005).

Social stereotypes are beliefs about groups of individuals. Although these beliefs have behavioral consequences, they can often be incorrect. One reason stereotypes can be incorrect is when they are overgeneralized beliefs which do not represent the whole range of targeted

individuals (Allport, 1954; Fiske & Taylor, 1984/2020). Just as in the statistical definition, overgeneralization is a fallacy occurring when a statement is true only of a specific subset of the target group. That statement can be erroneous when the subset is not representative of the original population (generalization error; Hastie, Tibshirani, Friedman, 2009). Our question here is: If some beliefs are inaccurate, how do people learn them in the first place and why do these beliefs persist? This paper investigates the process of stereotype formation from the perspective of insufficient statistical learning.

Past Wisdom: Cognitive Limitations and Sample Biases

Past research suggests two possible views on the origins of how people form inaccurate stereotypes. One view, dating back to the cognitive revolution in the 1970s, asserts that false associations between groups and traits come from cognitive limitations (Bodenhausen & Morales, 2012; Fiske & Taylor, 1984/2020; Tversky & Kahneman, 1974). The environment is too complicated and information-dense to possibly encode accurately, so the mind makes do by allocating limited attention and encoding, which can lead to inaccurate impressions. For instance, the co-occurrence of two rare events, such as witnessing both a member of a minority group and a negative behavior, is salient information, which attracts attention. As a result, this rare co-occurrence looms large in the mind. In other words, the mind's processing system is biased, leading to the emergence of negative stereotypes about minority groups (illusory correlation; Hamilton & Gifford, 1976).

Another view, originating from ecological perspectives, argues that false associations between groups and traits arise from sample biases (Fiedler, 2000; Fiedler & Juslin, 2006). Even when the mind uses appropriate cognitive strategies to process data, the environmental inputs are

rarely representative. For instance, when more men (and therefore fewer women) hold leadership positions, people associate men more with leadership traits. This is an accurate inference based solely on the two variables given. However, it ignores causes that could explain the correlation between these two variables, such as systemic prejudice against women. In other words, the sample given is biased, which leads to the emergence of stereotypes linking gender and leadership ability (social role theory, Early & Wood, 1999; pseudocontingency; Fielder & Freytag, 2004).

New Perspective: Insufficient Statistical Learning

In this paper, we propose an integrative, third perspective in understanding the formation of inaccurate stereotypes: Insta-learn (*Insufficient Statistical Learning*). Aligned with statistical definitions of overgeneralization, Insta-learn defines stereotypes as associations learned from two or more variables (e.g., race and criminal behaviors, gender and leadership roles). These associations represent temporary patterns derived from small samples (e.g., media exposure, representations in one's childhood or organization) and may be spurious to the extent that they do not accurately reflect patterns in the larger population (e.g., the population of all ethnic minorities and their behaviors, or all men and women and their positions). Insta-learn proposes that although people learn accurately from small samples, because they stop sampling too early to collect adequate evidence, they may end up with spurious associations. As a result, people stay with such locally-accurate-but-globally-inaccurate associations, causing stereotypes to emerge and persist.

Although our idea is inspired by the sample bias literature (e.g., Fiedler, 2000), Insta-learn differs in at least one crucial respect: it assumes that samples are endogenously generated.

In other words, people have control over how many samples they will see before drawing conclusions. One key assumption of the Insta-learn framework is that people insufficiently sample information. Consistent with this assumption, across domains people have been shown to sample much less information than they originally thought they would (Klein & O'Brien, 2018). For example, when evaluating whether they liked or disliked a painting style, people originally thought they would choose to view 16 paintings on average before making up their minds — in fact, they only viewed three paintings. Similarly, when judging a person's intellect or character, people thought they would choose to view five pieces of information, but only chose to view three pieces. Likewise, recent impression formation research has shown that participant-determined samples are substantially smaller than experimenter-determined samples (Prager, Krueger, & Fiedler, 2018). When forming impressions, if the first few traits happened to convey a clear-cut impression (e.g., Bob is friendly, nice, trustworthy, and helpful), then people cut the sample short. Empirically, participants sampled about seven to eight traits from a pool of 36 before making judgments, especially when the initial traits contained negative or extreme signals. The consequence of this endogenous sampling is that people may end up viewing samples, which are not representative of the population data. In turn, these samples could lead to spurious associations. Because people have control over how many samples they choose to see, the process by which inaccurate stereotypes develop is not simply a result of the environmental statistics, but is shaped by individual cognition, namely insufficient sampling.

Insta-learn also differs from the cognitive limitation perspective (e.g., Hamilton & Gifford, 1976; Tversky & Kahneman, 1974), which assumes that environmental statistics are somewhat irrelevant because the mind eventually thin-slices the external world into incomplete

pictures. In contrast, Insta-learn predicts that as environmental samples change, so do the learned associations. If the sample indicates a clear statistical relationship, people should not ignore it but should be able to extract the pattern. This prediction is in line with statistical learning accounts within cognitive psychology (Saffran, Aslin, & Newport, 1996; Tenenbaum, Kemp, Griffiths, & Goodman, 2011). Ample evidence shows that not only adults, but even infants and preschoolers are able to extract statistical information from available data (Tenenbaum et al., 2011; Rhodes & Baron, 2019). For example, a recent experiment asked both adults and 4- and 5-year-olds to infer novel agents' group memberships and social preferences (Velez & Gweon, 2020). By only observing minimal statistical information about how preferences are distributed across social groups, participants readily inferred the preferences of new individuals when they had appropriate evidence to support these inferences (e.g., 4 out of 4 blue Gazorps like Kiki snack vs. 2 out of 4 blue Gazorps like Kiki snack). Similarly, in subtler tasks such as evaluating human faces, adult participants are able to extract underlying statistical distributions of novel faces and these distributions shape their evaluations accordingly. In this paradigm, novel faces were generated from various statistical distributions (e.g., Gaussian or uniform). Participants, who implicitly learned the face distribution, tended to evaluate faces closer to the central tendency more positively than faces toward the edges of the distribution (Dotsch, Hassin, & Todorov, 2016). The implication of this sample adaptive learning is that stereotypes are influenced by the immediate context. Because samples influence the way people see patterns, the process by which inaccurate stereotypes develop is not only a function of internal cognition but also a function of the external environment, namely sample statistics.

As illustrated, a more comprehensive account of how inaccurate stereotypes emerge requires us to think of the process as an active interaction between what samples are fed to minds, how those minds process those samples, and whether those minds decide to collect more samples before making judgments. Insta-learn posits that humans are active learners of the environment. Starting from a small sample, people can extract statistical patterns within the sample accurately and quickly. However, people may prematurely decide that they are satisfied with what they have seen and fail to sample enough. In turn, such insufficient sampling could give rise to inaccurate stereotypes even when a fuller picture of the data would show otherwise.

Case Study: Face stereotypes

As a case study, we focus on face stereotypes, as they are both theoretically relevant and operationally rich. Theoretically, face stereotypes, defined as spurious associations between facial appearance and social behaviors, are thought to originate from overgeneralization (Zebrowitz, Fellous, Mignault, & Andreoletti, 2003; Zebrowitz & Montepare, 2008; Oosterhof & Todorov, 2008). For example, baby-faced adults are stereotyped as trustworthy and submissive. According to the overgeneralization account, this stereotype may emerge from an initially accurate generalization — namely, that most babies are not harmful or able to act on their individual agency. However, people then over-generalize those impressions to adults whose facial features resemble those of babies, leading to the assumption that these adults would have similar qualities. Although theoretically relevant, previous accounts of the origin of face stereotypes are highly domain-specific. Here we propose a more general mechanism. Operationally, features from human faces are on a higher-dimensional continuous space, as compared to race or gender which are, although not ideal, often categorically defined. This

continuous feature of human faces gives us flexibility and advantages to manipulate experimental stimuli which we discuss below.

Two main challenges arise in studying the Insta-learn hypothesis in the context of face stereotypes. The first challenge lies in disentangling whether a face stereotype is learned from sample or population statistics. In real world data, they tend to co-vary, perhaps due to self-fulfilling prophecy or explicit discrimination. To solve this problem, we created novel pairings of facial features and social behaviors, pairings for which participants could not possibly have any prior knowledge. This allowed us to experimentally manipulate both the population and sample statistics to test whether participants learn spurious associations from initial samples even though the population statistics would indicate an entirely different relationship. The second challenge lies in identifying whether the spurious association is driven by inadequate sampling on the part of participants or by other cognitive processes. To address this challenge, we made the sampling procedure flexible. We experimentally encouraged participants to flexibly sample more or fewer examples, so that we could compare their results with those of participants who were forced to view all the examples in the population.

To our knowledge, this is the first set of studies examining (a) whether people learn associations between the appearance of novel faces and social behaviors from small samples, (b) whether people update their learned associations with more samples, and (c) the mechanisms driving decisions not to sample further. We formalize this problem in a novel statistical learning paradigm to quantitatively delineate the emergence and persistence of novel face stereotypes.

Overview of empirical tests

To test the present hypothesis, participants were asked to make judgments on how generous a new face is based on the participants' past observations of other faces. In all experiments, we showed participants a small sample of people's faces varying on one underlying non-social appearance dimension and their social behaviors, namely, charitable donations. We then asked participants to either (a) continue viewing all remaining sets of faces and behaviors, or (b) stop viewing more faces at any time if they feel confident to make a decision. We finally asked participants to make judgments on new faces they never saw during any training phase, by asking how much they think this new person would donate. At the end of each experiment, we collected standard demographic information. See a conceptual paradigm in Figure 1.

Importantly, unbeknownst to participants, facial appearance and charitable behaviors were completely uncorrelated in the population dataset from which the samples were drawn. That is, any given face was equally likely to commit any given behavior. If participants sample all faces and learn the (lack of) covariation accurately during the training phase, they should therefore make random guesses in the testing phase. However, if participants sample only a few faces and jump to conclusions too soon, they will make decisions according to the initial samples they happened to see. Critically, the initial samples were chosen by the experimenters to show a linear relationship between facial appearance and charitable behavior—a fact also unknown to participants. Hence, if participants make decisions based solely on the initial sample statistics, their decisions should reveal a systematic linear relationship, and should not look like random guesses.

Figure 1. A conceptual illustration of the experimental paradigm.

Note. Experimental paradigm. Unbeknownst to participants, the systematic variation of face stimuli (a. Face stimuli example using one identity) and the amounts of charitable donation were uncorrelated (b. Population training dataset with three initial samples highlighted in dotted circle). First, we showed a positively/negatively/zero-correlated initial sample to participants (c. Initial sample learning as in the positive initial example). Next, we asked participants to make initial guesses on new testing faces (d. Manipulation check). Then, we prepared all remaining samples in the population dataset for participants to view. Here participants were assigned to either fixed or free sampling experimental conditions. According to the instructions, participants in the fixed sampling condition can ‘Press V to view more’ until the end, whereas participants in the free sampling condition can ‘Press V to view more or S to skip to the decision page’ (e. Sampling). Finally, we asked participants to make final guesses on the initial testing faces and more new testing faces (f. Testing). See details in the main text and try this experiment using the link in supplemental materials.

Experimental stimuli were designed with the following considerations. First, face stimuli drew on previous research in face perception and impression formation (Todorov et al., 2015). To avoid prior beliefs on what types of faces look trustworthy or generous, using FaceGen3.1, we generated novel faces that are orthogonal to the core social dimensions: attractiveness,

competence, dominance, extroversion, likability, threat, and trustworthiness. Visual inspection of the changing dimension reveals that it roughly reflects the width of a face. The manipulated dimension ranges from -10 SD to +10 SD (or -8 to +8 in Experiments 2, 3, 4, 5, and 6), in increments of 2 SD. To avoid memory biases, we generated 200 distinctive identities, so that participants never saw two identical faces during the task. Next, behavioral stimuli drew on trust games in behavioral economics and simple charitable giving scenarios. Participants read how much money a person (exemplified by a face) donated in a trust game or a hypothetical charitable event. The amount indicated levels of generosity of a stimulus face. The amount of money ranged from \$0 to \$100, in increments of \$10 (or \$25 in Experiments 2, 3, 4, 5, and 6). Within this paradigm, Experiment 1 manipulated the initial samples of faces and behaviors to be linearly and positively correlated. Half of the participants were given the opportunity to sample as many faces as they want (i.e., free sampling) and half were asked to sample all faces (i.e., fixed sampling). We hypothesized participants in the free sampling condition would report a more positive linear relationship than participants in the fixed sampling condition. Experiment 2 tested the robustness of the initial sample manipulations by including the following simple relationships: positive linear, negative linear, and zero correlation. We hypothesized a main effect of the initial samples and sampling condition differences as in Experiment 1. Experiment 3 tested the hypothesis in a more naturalistic way, by removing manipulation checks and fixed sampling conditions. We hypothesized the results would replicate Experiment 2's free sampling condition. Experiment 4 tested this paradigm with a more stringent design using stronger incentives for accuracy and aggressively fewer trials of initial samples. We hypothesized the results would replicate Experiment 3. Experiment 5 explored the effects of the strength of the initial sample

statistics. We hypothesized participants would persist in using spurious associations, replicating Experiment 4, in their final decisions regardless of a weaker initial signal. Finally, Experiment 6 explored the utility of additional, reliable information. Because throughout the experiments, the population statistics was set at zero, it was not obvious what participants should do in such a situation, besides randomly guessing. Therefore, we hypothesized that to the extent that people utilize additional and reliable information, they should be able to ignore the spurious associations inferred from the initial samples.

All hypotheses, analysis plans, and study materials are pre-registered at https://osf.io/syc6b/?view_only=555262392d9b4943b2145c0ed00efd23. See more details and a live experiment demo via supplemental materials. Experiments were programmed in JavaScript using the jsPsych library (de Leeuw, 2015). Analyses were performed in Python 3.7.1 and R 3.3.1. The protocol for these studies was approved by the Institutional Review Board for Human Subjects. Informed consent was obtained from all participants before the experiment and debrief was released after the experiment.

Experiment 1

The first study establishes a baseline paradigm on which other experiments build upon. We predicted that participants would learn linear associations between novel faces and social behaviors even from a population dataset which shows no such associations. These associations are formed in two steps. First, an initial sample, biased and small, allows people to learn a spurious association. We hypothesized participants would be able to learn the linear relationship from the initial sample. Second, a lack of sufficient sampling prevents them from unlearning the spurious association. We hypothesized participants in the free sampling condition would be less

likely to revise their prior spurious beliefs as they would sample fewer examples compared to those who were forced to sample the whole population dataset. The hypothesis would fail if participants do not demonstrate a linear relationship in their manipulation check decisions, or if responses in their final decisions between the two sampling conditions do not differ systematically.

Methods

Participants. We recruited 120 participants (7920 nested trials) from the university subject pool in order to obtain .80 statistical power to detect small-to-medium-sized effects in a multi-trial between-participants design with two conditions. We intentionally chose subject pool participants because we reasoned students might be more focused on long tasks than online workers.

Procedures. Participants received an online experiment link via email. They were instructed to complete the study within 24 hours and to provide a completion code to receive credit. The completion time was around 45 minutes. The experimental procedures can be divided into four phases: initial sample learning, manipulation check, sampling, and testing phases (Figure 1).

Participants first agreed to the consent form and to participate all the way to the end of the experiment, via a procedure used to reduce subject attrition in online subject pools (Zhou & Fischbach, 2016). Participants who did not agree to these terms were directed out of the study.

Participants were then provided with the general task instructions. Their task was to decide how much money to share with new people. They were first taught the basic rules of a trust game (Berg et al., 1995). This game is often used to test cooperative motives. In the context of the game, participants typically share more Monetary Units (MUs) with partners they perceive as generous. Specifically, we paired our participants with hypothetical partners instantiated by

artificial faces. Participants were given 10MUs. They were told that they must send some amount of their MUs to an anonymous partner. They could send as few as 0 MU to as much as 10 MUs. They also learned that whatever amount they send to the partner would be quadrupled. The partner would then receive the money and make a similar choice, deciding whether to give some amount of the now-quadrupled money back to the participant. Participants were told that the goal is to earn as many MUs as possible, because these units would be converted to a real bonus at the end of the experiment (100 MU = \$1). To make the rules easier to understand, we explicitly mentioned that the optimal strategy would be to share more MUs with those who are more likely to send MUs back. We told participants that how much each partner would send back was pre-determined, and that they could learn about the group of potential partners by observing some of their behaviors first.

Next, participants entered the initial sample learning phase. They saw 11 pairs of faces and past behaviors. We told them that to help with their decisions, they will first see 11 randomly selected pictures of partners and their past donations. The 11 examples were chosen by the experimenters to create a positive linear relation between appearance of faces and behaviors. The varying facial features are rich and holistic, but roughly, narrower faces donated more ($r = 1.0$). Each stimulus pair of face-behavior was displayed sequentially, with one face image in the center of the screen and one behavioral statement (boldfaced text “This person shared x%”) under the face image. Below the behavioral statement was the instruction to “Press V to view next”. After pressing V on their keyboard, participants saw a catch trial with a centered black cross for 50 milliseconds before the next stimulus trial. The faces were presented according to the MUs they returned. For example, the face stimulus which returned 0% of their MUs was shown first, then

the face stimulus which returned 10%, and so on until the last face stimulus which returned 100%.

After the initial sample learning phase, participants entered the manipulation check phase. They were asked to share money with 33 new partners from the same group, using a slider ranging from 0 to 10 MUs. This was done to make sure participants learned the (intentionally obvious) statistical pattern between faces and their behaviors in the initial small sample. Face stimuli appeared at the center of the screen sequentially and in a random order. Under the face image, participants could move the slider to indicate their decisions. They could then advance to the next face.

After the manipulation check phase, participants entered the sampling phase. They were told that there were 100 more members in the group. To encourage accurate learning, the following text was displayed in bold: “Decisions based on small samples are not always accurate”. Participants could view more examples before making their final decisions.

Approximately half of the participants were randomly assigned to the free sampling condition ($N = 62$) while the remainder were assigned to the fixed sampling condition ($N = 58$). In the free sampling condition, participants had the option to view more examples by pressing the V-key, or to skip to the decision page by pressing the S-key. In the fixed sampling condition, participants could only view more examples (by pressing the V-key) until they viewed all faces (i.e., the population dataset). As mentioned above — and in contrast to the initial sample — the statistical pattern in the population dataset showed no linear relationship at all; faces with different widths had equal probability to donate any amount ($r = 0.0$).

Finally, participants entered the testing phase. They were asked to share MUs with 66 members of the group (i.e., stimulus faces). As in the manipulation check phase, they could indicate their decisions by moving the slider below the face image anywhere between 0 to 10 MUs. Among the 66 faces, 33 were from the manipulation check phase, and 33 were completely new faces that participants had never seen prior to this phase. And just as in the manipulation check phase, face stimuli appeared at the center of the screen sequentially and in a random order. Participants advanced to the next face by clicking a “next” button on the same page. In this testing phase, we tested whether and how participants (a) changed their prior decisions, and (b) generalized the learned association to new members.

Results

As predicted, data from the manipulation check phase showed participants learned the positive linear relationship from the initial sample. Using a multilevel model with initial decisions as the dependent variable and faces as the independent variable, with error clustered at the individual level, we observed that participants’ decisions were associated with our manipulated facial dimension: for each unit increase in the manipulated facial dimension, participants shared 1.88 more monetary units, $b = 1.88$, 95% $CI [1.72, 2.04]$, $p < .001$ (Figure 2a).

Moreover, data from the testing phase showed different associations were learned as a function of the sampling condition. Participants’ decisions in the free sampling condition were more likely to reflect the linear relationship between appearance and behaviors in the initial biased sample than participants’ decisions in the fixed sampling condition. In a multilevel model with final decisions as the dependent variable, faces and sampling conditions (0 for fixed and 1

for free) as independent variables, and error clustered at the individual level, we observed that sampling condition significantly interacts with appearance, interaction $b = .32$, 95% $CI [.09, .55]$, $p = .006$ (Figure 2e). Lastly, we found participants in the free sampling condition on average viewed less than one fifth (23/121) of the total population of faces (Figure 3a).

In addition to group level effects, we observed individual heterogeneity. Using trials of each participant as the unit of analysis, faces as the independent variable and decisions as the dependent variable, we obtained participants' decision slopes in the initial testing and final testing with 95% confidence intervals. Results showed that whereas 25 out of 58 participants in the fixed sampling condition used positive linear slopes, 34 out of 62 participants in the free sampling condition used positive linear slopes, indicating the persistence of the initial spurious relationship (see supplemental materials).

In sum, we found participants accurately learned the statistical relationship between novel faces and charitable behaviors from small samples. However, participants in the free sampling condition sampled very few faces and behaviors before making decisions. As a consequence, they were less likely to revise their initial judgments, leading to maintaining incorrect beliefs about the relationship between appearance and behavior.

Experiment 2

The second study aims to rule out an alternative interpretation of the results of Experiment 1. Participants could have applied some meta-priors such as un-manipulated signals in facial features that were accidentally consistent with the positive linearity, regardless of what the initial sample intended to show. For example, faces with greater width are perceived as less

trustworthy and less likely to reciprocate (Stirrat & Perrett, 2010), which is consistent with our previous results.

To further establish that participants do learn from the initial samples, we manipulated the statistical relationship in the initial sample. In addition to the positive linear relationship between appearance and behaviors in Experiment 1, we added one condition where the initial sample included the completely opposite relationship (i.e., a negative linear relationship) and another condition where the initial sample included no relationship at all (i.e., a correlation of zero). We hypothesized that manipulation check tests would reveal a difference in decisions across initial sample conditions. The hypothesis would fail if participants in all three conditions uniformly revealed a positive linear relationship in their manipulation check decisions.

Methods

Participants. We collected data from 320 participants (9600 nested trials) via Amazon Mechanical Turk Prime (i.e., Now Cloud Research -- a service making it easier to recruit and manage higher quality participants on Amazon Mechanical Turk) to ensure inclusion of at least 50 participants in each condition to detect small-to-medium-sized effects in a multi-trial between-participants design. We changed our participant pool to online workers due to COVID-19-related restrictions. All participants accessed the link and provided a completion code for monetary compensation. Our participants were 61% male, 75% White, with an average age of 37 and a standard deviation of 20, and 42% having earned a bachelor's degree. In addition, we asked for free comments at the end of the experiment: the majority of the participants were satisfied and engaged in the task (see supplemental materials).

Procedures. The procedures were identical to Experiment 1 with minor changes for simplification purposes. First, we reduced the number of stimuli to ensure response quality from online workers (Robinson et al., 2019). This task took about 10 minutes with fewer trials and simpler instructions. Second, we simplified the general instructions, replacing the relatively complex trust game with a simpler estimation task. Participants were introduced to a hypothetical charitable event. Their task was to guess how generous a new attendee is based on observations of past attendees. They read: “A large group of people attended a charity event. Each attendee was endowed with \$100 and donated some of this amount. Your task is to figure out whether appearance predicts charitable behavior; donating more money.” As in Experiment 1, there were four phases.

In the initial sample learning, participants were randomly assigned to see 5 examples (instead of 11) with either positive ($r = 1.0$, $N = 106$), negative ($r = -1.0$, $N = 100$), or near zero ($r = 0.1$, $N = 114$) correlations.

In the manipulation check phase, participants estimated the donations of 15 new examples (instead of 33), from \$0 to \$100 (instead of 0 to 10 MUs).

In the sampling phase, participants were reminded that they only saw 5 examples, which may or may not be representative of how the rest of the attendees donated. Participants in the free sampling condition saw boldfaced instructions saying “You can sample up to 20 additional samples. That is, you can continue viewing more examples (by pressing V-key), or you can stop at any time to make your decisions (by pressing S-key).”. In contrast, participants in the fixed sampling condition saw boldfaced instructions saying “You will see 20 additional examples. That

is, you will continue viewing more examples before making your decisions (by pressing V-key only)”).

In the testing phase, participants evaluated 30 new faces (instead of 66) with 15 old faces from the manipulation check and 15 new faces they had never seen before. At the end of the study, standard demographic information was collected.

Results

As predicted (Figure 2b, 2c, 2d), data from the manipulation check phase showed that participants learned the respective linear relationship from the assigned initial sample: positive initial, $b = 2.20$, 95% $CI [1.68, 2.74]$, $p < .001$; negative initial $b = -1.51$, 95% $CI [-1.72, -1.29]$, $p < .001$; zero initial $b = .45$, 95% $CI [-.06, .95]$, $p < .001$. This main effect indicates that the observed relationship in Experiment 1 is caused by the statistics hidden in the manipulated initial samples rather than some meta-priors. It indicates that participants indeed learned the relationship between an arbitrary facial dimension and behavior from a small sample.

Replicating the sampling interaction in Experiment 1, data from the testing phase found participants in the free sampling condition were more likely to rely on the initial patterns for their final decisions than participants in the fixed sampling condition (Figure 2f, 2g, 2h). In both the positive initial condition, the interaction $b = .33$, 95% $CI [.01, .64]$, $p = .043$, and in the negative initial condition, the interaction $b = -1.01$, 95% $CI [-1.33, -.69]$, $p < .001$, were statistically significant. Although we expected to see lower variations in the zero initial condition, we observed a significant interaction $b = -.45$, 95% $CI [-.15, -.75]$, $p = .003$. Zero initial participants in the fixed sampling condition revealed a significantly negative linear relationship compared to participants in the free sampling condition, as if the more faces they saw, the more structure they

imagined. Lastly, free sampling participants viewed on average 11 out of 20 additional samples (Figure 3b).

Individual level analysis revealed heterogeneity. Among free sampling participants, 18 out of 55 in the positive condition and 18 out of 52 in the negative condition used statistically significant positive or negative slopes in their final decisions, respectively. Whereas among fixed sampling participants, 15 out of 51 in the positive and 7 out of 48 in the negative condition still used statistically significantly positive or negative slopes in their final decisions, respectively, even though they had seen the population dataset. In terms of the zero initial condition, some participants did indeed learn the zero relationship, but around one third of the participants (32 out of 114) still used statistically significant linear slopes to make their decisions (see supplemental materials).

In sum, we found that participants were able to learn statistical patterns from a small sample of faces and behaviors. The relationship learned can be positive, negative, or zero, which is a function of the statistical patterns encoded into the initial sample. Participants were less likely to revise their previously learned beliefs if they were free to sample faces and behaviors. The lack of sufficient observations impeded accurate learning. In fact, when participants were exposed to the full dataset, they were more likely to revise their priors to approach the ground truth, although not perfectly.

Figure 2. Decisions before and after the free versus fixed sampling conditions.

Note. Interaction effects of the sampling manipulation from Experiments 1 and 2. The horizontal axis indicates the intensity of the test face dimension, ranging from -8 SD to +8 SD (-10 SD to +10 SD in Experiment 1). The vertical axis indicates estimated amounts of charitable giving for a given face, ranging from \$0 to \$100 (0MUs to 10 MUs in Experiment 1). Figures are grouped by Experiment 1 (a, e), Experiment 2 positive initial (b, f), Experiment 2 zero initial (c, g), and Experiment 2 negative initial (d, h), with fixed (black) and free (red) sampling conditions. The plotted lines represent fitted linear regressions displaying the central tendency and 95% confidence intervals. The upper figures (a-d) represent interaction effects before the sampling manipulation. As expected, participants in both conditions saw the same examples, so there should be no differences. The lower figures (e-h) represent interaction effects after the sampling manipulation. As hypothesized, participants in the free sampling condition were more likely to use the correlations learned from the initial samples, i.e., steeper slopes (i.e., red) as compared to the fixed sampling condition. See details in the main text and supplemental materials.

Figure 3. Number of viewed samples in fixed versus free sampling conditions.

Note. Sample counts in Experiments 1 and 2. The vertical axis indicates the number (mean and 95% confidence intervals) of examples participants viewed in free (red) versus fixed (black) sampling conditions, visually grouped by Experiment 1 (a) and 2 (b) with positive, negative, and zero initial correlations. Dashed lines indicate all samples participants could have seen. As shown, participants in the free sampling condition consistently viewed significantly fewer examples than those in the fixed sampling condition. See details in the main text, and relations to final decisions in Figure 2.

Experiment 3

The third study aims to replicate the main results of the first two experiments, mimicking real-world decisions in a more naturalistic way. We used a subtle manipulation in this study to overcome two concerns. First, in most real-time decisions, people are not typically asked to reveal what they have learned before they have finished learning new things. Moreover, making evaluations before subsequent learning can cause an anchoring effect (Tversky & Kahneman, 1974) that biases the results toward initial pattern confirmation. We therefore removed the manipulation check phase. Second, in most real-time decisions, people are rarely forced to view all samples before making decisions (Klein & O'Brien, 2018). We therefore removed the fixed

sampling manipulation. We hypothesized that participants in this study should behave like those in the free sampling condition of Experiments 1 and 2 — that they would learn spurious associations, sample fewer faces, and be less likely to revise their beliefs. The hypothesis would fail if we do not observe linearity in participants' final decisions, especially for those who are assigned to positive or negative initial sample conditions. We also explored whether individual differences in tolerance toward ambiguity (Need for cognitive closure, 15-item short scale; Roets & Van Hiel, 2011) and other social demographics relate to participants' decisions.

Methods

Participants. We collected data from 300 new participants (9000 nested trials) via Amazon Mechanical Turk Prime Cloud Research (i.e., a platform with higher quality participants) for this study with the same power considerations as in the previous two studies. Our participants were 56% male, 73% White, with an average age of 40 and a standard deviation of 13, and 47% having earned a bachelor's degree. Free response comments again suggested satisfaction and engagement.

Procedures. As mentioned, the procedures were identical to Experiment 2 except that we removed manipulation check tests and the fixed sampling conditions. This task took less than 10 minutes.

Participants were introduced to the estimation task of generosity, as in Experiment 2. The instructions hinted that they could see up to 25 pictures of attendees and their donations. They saw 5 faces in the beginning and then they could continue viewing more faces or stop at any time. Afterwards, they made decisions about 30 new attendees. After reading the instructions, participants were randomly assigned to positive ($r = 1.0$, $N = 95$), negative ($r = -1.0$, $N = 100$),

and near zero ($r = 0.1$, $N = 105$) initial sample conditions. Right after viewing these initial samples, participants were prompted to freely sample, such that they can ‘Press V to view more, or S to skip to the decision page. Participants then provided their decisions on test faces, and completed questionnaires and standard demographic questions.

Results

First, as expected, we replicated the free sampling condition in Experiment 2 (Figure 4a). Data from the testing phase showed that participants who were assigned to the positive initial condition judged new faces based on a positive linear relationship, $b = 1.33$, 95% CI [1.39, 1.72], $p < .001$. Participants who were assigned to the negative initial condition judged new faces based on a negative linear relationship, $b = -.54$, 95% CI [-.70, -.38], $p < .001$. Participants who were assigned to the near zero initial condition judged new faces based on a slightly positive but closer to zero linear relationship, $b = .45$, 95% CI [.07, .83], $p < .001$. In addition to the initial 5 faces, participants on averaged viewed 9 more faces (out of 20) before making decisions (Figure 5a).

We explored which individual level characteristics correlate with levels of persisting on the initial spurious patterns. We estimated whether a participant persisted in using spurious associations by obtaining their decision responses with 95% confidence intervals. If their decision intervals included 0, this was coded as ‘no’, indicating their decisions did not systematically differ from zero and, thus, were more in line with the population statistics. If their decision intervals did not include 0, this was coded as ‘yes’, indicating their decisions systematically deviated from zero and, thus, were more in line with the initial sample. We ran logistic regressions with various individual predictors. Although the sign of the need for cognitive closure indicated that the more participants were intolerant toward ambiguity the more

likely they were to rely on prior knowledge of some linear relationship, the magnitude was not statistically significant ($b = -.25$, 95% $CI [-.55, .05]$, $p = .108$). Neither length of task time, total numbers of samples viewed, age, race, gender, nor education level predicted persistence (more details in supplemental materials).

In sum, we replicated the previous findings, using a more naturalistic and subtle design. Participants learned the statistical relationship between novel faces and social behaviors from very few samples — five examples. They were less likely to sample the complete dataset even though they knew it contained only 25 examples. As a result, their final decisions were heavily influenced by the initial samples.

Experiment 4

The fourth study aims to extend Experiment 3, using a more stringent and less artificial design. First, we added stronger monetary incentives for accuracy. Although we emphasized accuracy throughout the previous studies (e.g., hypothetical trust game earnings in Experiment 1; boldfaced text reminding participants that decision based on small samples can be misleading), this emphasis was more instructional than substantial. Moreover, given that online workers are paid by time cost, it is possible that their behaviors (e.g., sampled fewer faces) were driven by higher incentives for speed. To compensate, we added stronger incentives for accuracy with real money.

Second, we revised the instructions to make the task less contrived. Instructions in previous studies and the sorted order of the presentation of the initial stimuli might have mis-signalized that the initial samples were drawn randomly. These design artifacts could give the impression that the relationship in the initial sample was more or less representative of the

population (Tversky & Kahneman, 1974). We therefore revised the instructions to minimize these issues.

Third, we reduced the number of the initial samples to three. The fact that participants in previous studies learned the hidden statistical relationships from as few as five examples was already impressive, but it might be rare for people to encounter five examples with perfect linearity in rapid succession in real life. Three examples are therefore more plausible than five. We do not know if such a small number of examples is sufficient for people to form any meaningful beliefs, so we tested this here.

Lastly, we added one between-subject condition to explore how categorical information influences this process. This new categorical information cue mimics real-world decisions when people are primed with social categories, regardless of how the categories are related to behaviors. Specifically, we told participants that one of the attributes on which the faces vary is their width.

We hypothesized that participants in this study should replicate the behaviors observed in Experiment 3 (and in Experiments 1 and 2, in the free sampling condition) — that they would learn spurious associations and sample fewer faces, and as a result would not revise their learned inaccurate beliefs. Furthermore, participants who were cued that faces may vary in width could show stronger effects than those in the no cue condition. In both conditions, the hypothesis would fail if we do not observe linearity in participants' decisions who are assigned to see the initial three examples containing positive or negative linear associations between faces and behaviors.

Methods

Participants. We collected data from 320 new participants (9600 nested trials) via Amazon Mechanical Turk Prime Cloud Research for this study with the same power considerations as in the previous studies. Our participants were 52% male, 75% White, with an average age of 41 and a standard deviation of 13, and 66% having earned a bachelor's degree or higher. Free comments again suggested satisfaction. This task took less than 10 minutes.

Procedures. As mentioned, the procedures were identical to Experiment 3 with the following modifications: First, before entering the game, participants now saw a bonus instruction page: “Try your best to make accurate decisions. For each correct guess, you earn a \$0.05 extra bonus.” The research team granted corresponding amounts of bonus to participants afterwards. Note that if participants estimated correctly for all trials, they would increase total payoffs by 150%, which makes this accuracy incentive nontrivial.

Second, participants now learned that there will be 25 pictures (same as before), and they will first see three selected (instead of “five” “randomly” selected) pictures. Participants were then randomly assigned to positive, negative, and near zero initial sample conditions, where the three faces were presented in a random order (instead of an ascending order of donation size).

Third, in the newly added categorical information cue condition, participants read a hint before entering the game: “Attendees vary by many attributes, the width of their faces is one example.” In the no categorical cue condition, participants did not read this hint (same as before).

Results

First, data from the no categorical cue condition replicated Experiment 3 — even after adopting these stringent modifications, participants learned spurious associations and persisted with them in their decisions (Figure 4b). Data from the testing phase showed that participants

who were assigned to the positive initial condition ($N = 55$) judged new faces based on a positive linear relationship, $b = 1.19$, 95% CI [.97, 1.42], $p < .001$. Participants who were assigned to the negative initial condition ($N = 54$) judged new faces based on a negative linear relationship, $b = -0.46$, 95% CI [-.69, -.24], $p < .001$. Participants who were assigned to the near zero initial condition ($N = 51$) judged new faces based on a slightly positive but closer to zero linear relationship, $b = .28$, 95% CI [.04, .51], $p = .02$. In addition to the initial 3 faces, participants on average viewed 10 more faces (out of 22) before making decisions (Figure 5b).

Next, data from the categorical cue condition also replicated Experiment 3 (Figure 4c): Participants in the positive initial condition ($N = 57$) judged new faces based on a positive linear relationship, $b = .91$, 95% CI [.68, 1.13], $p < .001$; participants in the negative initial condition ($N = 53$) used a negative linear relationship, $b = -1.25$, 95% CI [-1.49, -1.02], $p < .001$; and participants in the near zero initial condition ($N = 50$) used a slightly negative but closer to zero linear relationship, $b = -.34$, 95% CI [-.58, -.10], $p = .005$. Participants in this condition also on average viewed 10 more faces before making decisions (Figure 5c).

As expected, the effects were stronger in the categorical cue condition than the no cue condition, but only in the negative and zero initial conditions: A significantly steeper negative trend in the negative (interaction $b = -.79$, 95% CI [-1.12, -.47], $p < .001$), and a significant change from slight positive to slight negative in the zero (interaction $b = -.62$, 95% CI [-.96, -.29], $p < .001$) initial conditions. There was no difference in the number of samples being viewed between the cue and no cue conditions (interaction $b = -.57$, 95% CI [-3.41, 2.27], $p = .694$).

Individual level analysis collapsing across the cue and no cue conditions replicated previous findings: The tendency to persist in using spurious correlations was not related to the need for cognitive closure, total number of trials viewed, total time spending in the task, age, gender, race or levels of education, suggesting that the mechanism is general across demographic groups, at least as measured by our paradigm (see supplemental materials).

In sum, participants learned spurious associations fast — from as few as 3 examples. They sampled fewer new faces even when they were incentivized to value accuracy rather than speed. As a result of this fast learning and insufficient sampling, when making new predictions, we observed participants persisted in using spurious beliefs that are consistent with the initial small samples but inconsistent with the larger dataset. This finding supports the Insta-learn hypothesis that inaccurate stereotypes are locally-accurate-but-globally-inaccurate associations, which can originate from insufficient statistical sampling. Categorical cues exacerbate this process, but even without any cues, we observed a robust effect.

Experiment 5

What do people do when the initial evidence is weak? So far, we manipulated the initial sample statistics to be either a perfect positive or a perfect negative linearity with the same three example faces. However, it is rare for people in everyday life to encounter examples with strong signals such as perfect linearity. It is also uncommon for people to encounter three identical targets with identical faces. In some situations, the strength of the statistical pattern may vary from person to person; the identity being encountered may also vary from person to person.

To introduce a more ecologically valid manipulation, we changed the group-level treatment to an individual-level treatment. Specifically, although each participant saw three

initial examples as in the previous study, the three examples were randomly chosen from the whole dataset. This change makes both the strength of the statistical pattern and the identity of the initial examples variable. In other words, both changes make the appearance and behaviors less transparent than previous experiments, particularly given different identities introduce variation beyond the variation on our manipulated dimension.

We hypothesized that participants would still be more likely to use spurious associations, and to sample fewer examples. However, we speculated participants who have encountered weaker signals in the initial three examples might continue to sample more. If so, we should observe their final decisions to correlate more with their own sampled statistics rather than the three initial examples or the whole population dataset. The hypothesis would fail if we observe participants sample almost all examples from the whole dataset, or if the associations in final decisions do not correlate with any patterns from the initial samples.

Methods

Participants. We collected data from 201 new participants (6030 nested trials) via Amazon Mechanical Turk Prime Cloud Research for this study. Power simulation used empirical results from 31 pilot trials, which suggest a multi-level model of 150 samples should give us a .81 effect size at .05 alpha level. Our participants were 51% male, 75% Whites, with an average age of 41 and a standard deviation of 14, and 57% having earned a bachelor's degree or higher. Free comments again suggested satisfaction and instructions clarity. In contrast with the previous studies, some participants indicated that this task was hard. This task took less than 10 minutes.

Procedures. The procedures and instructions were largely identical to Experiment 4's no cue condition, with one modification in the initial sampling phase: Participants saw three initial pairs

of examples, which were randomly drawn from the whole dataset of 25 total pairs (instead of fixed pairs). Note that this manipulation yields a total of 13,800 possible treatments (draw three pairs from 25 pairs), and we did not impose any constraints on which three pairs would be shown.

Results

First, consistent with previous experiments, participants on average viewed fewer examples than the total set before making decisions: 13 out of 25 examples in this study (Figure 5d). Second, consistent with previous experiments, not many participants gave random guesses in their final decisions. Roughly and descriptively, we found 36 out of 201 participants revealed a relationship between -0.1 to + 0.1 as their final decisions – an indication of random guessing.

More precisely, a group-level analysis categorized the statistical pattern in the initial three examples into three bins (as pre-registered): a Pearson correlation between faces and donations in the intervals $[-1, -0.1]$, $(-0.1, +0.1)$, $(+0.1, +1]$ was recoded as negative ($N=103$), zero ($N=17$), and positive ($N=81$) conditions respectively. We found participants in all three conditions tended to judge new faces based on a positive linear relationship, $b = .32$, 95% $CI [.15, .49]$, $p < .001$ (Figure 4d). We did not find any statistically significant interactions. In other words, regardless of what participants saw in their initial samples (positive, negative, or zero with varying strength), participants judged a narrower face to be more likely to donate. We replicated this finding with 5 fine-grained categories (see pre-registration and supplemental materials for details). Hence, the positive initial condition supported our group-level hypothesis, but the negative or the zero initial conditions did not.

Next, an individual-level analysis calculated Pearson correlation between faces and donations in the initial three examples, as well as Pearson correlation in the 15 testing trials. We hypothesized a positive rank correlation, which was not supported from the data (A Spearman rank-order correlation $r = .04$, $p = .591$). Hence, our individual-level hypothesis was not supported, since participants' final decisions were unrelated to the initial three examples.

In addition to the fixed three initial samples, we explored how participant-generated samples influence their decisions (as pre-registered). Each participant chose the number of samples they wanted to see before moving to the decision phase. We therefore calculated a Pearson correlation between faces and donations for each participant using their own samples, and categorized the correlation into negative ($N=68$), zero ($N=81$), and positive ($N=52$) conditions. We found participants in the self-generated positive condition judged new faces based on a positive linear relationship, $b = 1.13$, 95% CI [.89, 1.38], $p < .001$. This is statistically significantly different from participants in the self-generated negative (interaction $b = -1.12$, 95% CI [-1.44, -.80], $p < .001$) and zero (interaction $b = -.93$, 95% CI [-1.23, -.62], $p < .001$) conditions. Similarly, we replicated this finding with 5 fine-grained categories (see supplemental materials for details). Hence, participants' behaviors in the self-generated positive and zero conditions were consistent with our hypothesis, but participants' behaviors in the self-generated negative condition were inconsistent (Figure 5e and 4e).

In terms of individual-level results, we calculated Pearson correlation in participant-generated samples and their testing trials. We observed a statistically significant positive rank correlation (A Spearman rank-order correlation $r = .21$, $p = .002$), indicating participants final decisions were correlated with the relationship they observed in their self-generated sample.

In sum, we partially replicated our main findings with individually heterogeneous treatments. Each participant saw different pairs of initial samples, randomly drawn from the dataset. This manipulation tested how the strength of the initial statistical patterns influences final decisions. With varying identities and varying strength of the statistical patterns across participants, we found participants still tended to under-sample the data and to use spurious associations in their final decisions. Yet, inconsistent with previous experiments, we found participants' final decisions were not related to the initial three examples, but were related to their self-generated samples. One key difference between this and the previous experiments is that the initial relationship was weaker (average correlation among the three examples in the positive initial condition was $r = 1.0$ in previous experiments vs. $r = 0.6$; in the negative initial condition, it was $r = -1.0$ vs. $r = -0.7$) simply by virtue of the design. The current finding suggests that when the initial relationship is weak, people are willing to sample more to extract a meaningful, stronger relationship. Nonetheless, our participants sampled insufficiently and as a result, their final decisions were related to the patterns from the samples that they themselves generated. In other words, participants may have settled on a relationship from their own samples, stopped sampling early, and thus persisted in using that relationship they learned earlier.

Experiment 6

What do people do when the actual relationship is zero? Throughout our experiments, we designed the population level statistics to be zero. However, it is not obvious what people should do in such a situation. If there is no other information available, it is possible that people assume that there must be some relationship and that they settle on the first relationship (positive or negative) they observe in their self-generated samples. This could explain the persistent reliance

on spurious associations in our previous experiments. But if people are provided with additional and reliable information, they need not rely solely on the information gained from observing the face and behavior pairs. Moreover, to the extent that they discover that the latter information is unreliable (e.g., a zero correlation), they can completely disregard this information and instead rely on the additional, reliable information. Hence, this final study aims to investigate to what extent people utilize additional and reliable information.

Specifically, we provided an opportunity to seek more information during the test phase. Before making each decision, participants were asked whether they want to see more information. If they opted to see, they were provided with information about what the person donated the last time. Unbeknownst to participants, this information was perfectly predictive of what the person donated. Across trials, this additional information also contained the ground truth value of a zero correlation between facial appearance and behaviors at the population level.

We made the following predictions: First, participants in the zero initial condition would be more likely to request more information than participants in the positive or negative linearity conditions because of uncertainty. Second, participants who requested more information would use the additional information to make decisions. As a result, participants in the zero initial condition who ask for more additional information would be less likely to reveal spurious linear relationships, as they have access to more information that they can rely on, when compared to participants in the positive or negative initial conditions who would ask for less additional information. The hypothesis would fail if we do not observe differences in the number of information requests, correlations between additional information and actual decisions, or decision behaviors between the initial zero, initial positive, or initial negative conditions,

especially among participants who requested less additional information than participants who requested more.

Methods

Participants. We collected data from 304 new participants (9120 nested trials) via Amazon Mechanical Turk Prime Cloud Research for this study with the same power considerations as in the previous studies. Our participants were 46% male, 71% White, with an average age of 40 and a standard deviation of 14, and 60% having earned a bachelor's degree or higher. Free comments again suggested engagement and attention. This task took about 10 minutes.

Procedures. As described, this study closely replicates Experiment 4's no cue condition, but with one change in the test phase. To set up the background, in the testing instructions, participants read: "Hint: For each correct guess, you earn 5 cents extra bonus. Your choices will be counted as accurate with a margin of +/- \$5. You can also choose to pay 2 cents to see extra information. The extra information is the charitable behavior of the person in a past situation. Press any key to begin."

Next, instead of making decisions directly, participants were asked to first decide "Do you want more information? Y for yes and N for no". If the Y-key was pressed, participants moved on to the decision trial and saw "Extra information: This person donated x% last time". They made decisions on a slider bar from 0 to 100 above the question "Out of \$100, how much do you think this person donated?" If participants pressed the N-key, they did not see the extra information. This was the same as in Experiment 4, where participants made decisions directly. By design, a random guess without requesting any extra information would result in, on average, 15 cents bonus. In contrast, an informed guess with requesting extra information on all trials

would result in, on average, 90 cents bonus. Thus, relying on additional information should be helpful and in the interest of participants.

Results

First, replicating our previous results, participants in the positive initial condition ($N = 110$) judged new faces based on a positive linear relationship, $b = .69$, 95% CI [.52, .86], $p < .001$. Participants in the negative initial condition ($N = 91$) judged new faces based on a negative linear relationship, $b = -.41$, 95% CI [-.60, -.22], $p < .001$. Participants in the zero initial condition ($N = 103$) judged new faces based on a slightly positive linear relationship, $b = .46$, 95% CI [.28, .63], $p < .001$. Similar to previous experiments, in addition to the initial three faces, participants on average viewed 11 more faces before making decisions. Next, we examined the utility of the additional information.

Our first hypothesis was not supported (5f). We hypothesized that participants in the initial zero condition should be more likely to request additional information than participants in the initial positive or negative conditions, but results revealed no differences between conditions (zero initial $b = 9.05$, 95% CI [3.23, 14.86], negative initial $b = 9.56$, 95% CI [7.11, 12.01], positive initial $b = 10.76$, 95% CI [5.00, 16.53], all pairwise comparisons *ns*). More descriptively, about two-thirds of the participants in all three conditions requested information fewer than 10 times and about one-third requested more than 20 times (out of 30 total).

Our second hypothesis was supported. As predicted, there was a strong positive linear correlation between hints and actual decisions: Multilevel regression with hints as the independent variable and actual decision as the dependent variable clustered within each participant revealed statistically significant correlations: $b = 0.99$, 95% CI [0.97, 1.00], $p < .001$.

Regardless of the initial sample statistics, as long as participants requested additional information, they relied heavily on that new information to make decisions.

Our third hypothesis was supported as well. Among participants who searched for more hints (i.e., requested more than 20 hints out of 30 total; pre-registered), there were no initial condition differences in their final decisions: Initial negative $b = -0.12$, 95% $CI [-0.50, .26]$, $p = .55$; which differs, but not significantly, from the initial positive by $b = .44$, 95% $CI [-0.07, .95]$, $p = .09$; and also not statistically significantly from the initial zero by $b = .29$, 95% $CI [-0.25, .83]$, $p = .30$. In other words, most of them used the information given and ignored the face information, resulting in decisions similar to the population statistics.

In contrast, participants who searched for fewer hints (i.e., requested fewer than 10 hints; pre-registered) showed systematic persistence of the initial sample statistics in their final decisions (Figure 4f): Initial negative $b = -.57$, 95% $CI [-0.80, -.34]$, $p < .001$; which significantly differs from the initial positive by $b = 1.39$, 95% $CI [1.07, 1.70]$, $p < .001$; and from the initial zero by $b = 1.08$, 95% $CI [0.77, 1.40]$, $p < .001$. In other words, participants who had seen initial examples containing negative (positive or zero) relationship predicted new behaviors accordingly, without realizing those were spurious associations.

In sum, additional reliable information does make a difference, although not for all participants. Participants who asked for more direct information made more accurate predictions than participants who did not ask for more information. Nonetheless, not asking for more information seems to be the default as indicated by two-thirds of our participants regardless of what they initially saw. Asking for more information, although it increases accuracy, seems to be the exception. We observed, again in the testing phase, insufficient sampling.

Figure 4. Main effects of initial sample on participants' final decisions.

Note. Main effects of the initial sample on participants' final decisions from Experiments 3, 4, 5, and 6. The horizontal axis indicates the variation of the test faces, ranging from $-8SD$ to $+8SD$. The vertical axis indicates estimated amounts of charitable giving for the faces, ranging from \$0 to \$100. Figures are grouped by Experiment 3 (a), Experiment 4 without cues (b), Experiment 4 with cues (c), Experiment 5 with fixed 3 initial samples (d), Experiment 5 with participant-generated samples (e), and Experiment 6 when participants requested fewer hints (f). Each figure plotted positive (black), negative (blue), and zero (gray) initial conditions with fitted linear regression lines displaying central tendency and 95% confidence intervals. Spurious associations emerged from small samples across all variant conditions, except in d. See details in the main text and supplemental materials.

Figure 5. Number of viewed samples in consecutive decisions.

Note. Sample counts in Experiments 3, 4, 5, and 6. The vertical axis indicates the number (mean and 95% confidence intervals) of examples participants viewed in positive (black), negative (blue), and zero (gray) initial correlations, visually grouped by Experiment 3 (a), 4 without cues (b), 4 with cues (c), 5 with fixed 3 initial samples (d), 5 with participant-generated samples (e), and additional hint conditions (f). Dashed lines indicate all samples participants could have seen. The only difference between 5d and 5e is whether the same participant and their selected samples is categorized into positive, negative, or zero condition on the horizontal axis. As the number of viewed samples increases, so as the statistical relationship. Overall, participants in all experiments sampled fewer examples than the whole dataset provided regardless of various incentives. See details in the main text, and relations to final decisions in Figure 4.

General Discussion

Across six experiments, we found consistent evidence suggesting an Insta-learn process underlying the emergence and persistence of face stereotypes. First, participants learn quickly. They were sensitive to, and capable of learning, statistical patterns between arbitrary novel faces and social behaviors with samples consisting of as few as three examples. Second, participants do not sample enough. They were less likely to sample the complete dataset if they were given the opportunity to freely decide how many more examples they want to see. Due to the combination of quick learning and insufficient sampling, or insufficient statistical learning, participants in the free sampling condition were less likely to revise the previously learned spurious associations than those in the fixed sampling condition. Moreover, participants in a more naturalistic experiment replicated behaviors in the free sampling condition, suggesting the insufficient statistical learning occurs naturally. This effect holds even when participants knew they would earn bonuses with accurate decisions, suggesting that accuracy/speed tradeoffs do not mitigate this process. Even when the initial statistical patterns were weaker and noisier, participants still did not sample enough. They quickly settled on an initial relationship that persisted. When provided with additional direct and reliable information, participants were able to revise their inaccurate beliefs as long as they asked for this information. Yet, most participants rarely did ask for this information and, as a result, their spurious beliefs persisted. This mechanism of insufficient statistical learning is general across demographic groups such as age, gender, ethnicity, or social class, as well as individual differences in the need for cognitive closure. Taken together, face stereotypes — overgeneralization of spurious associations between

appearance of faces and behaviors — emerge from small biased samples and persist because of insufficient sampling.

The results of the current set of experiments suggest Insta-learn could be another, perhaps more integrative, way of looking at how inaccurate stereotypes emerge and persist. On the one hand, past research on cognitive limitations would predict that due to limited information processing capacities (e.g., Bodenhausen & Morales, 2012), participants cannot accurately learn about the initial sample statistics, or those who have seen the population dataset cannot learn the underlying statistical patterns. However, the current results (manipulation check and final decisions from the fixed sampling conditions in Experiments 1 and 2) invalidate such predictions. Instead, consistent with the environmental sample and statistical learning perspective, these results show that participants are in fact able to extract hidden statistics, even from high-dimensional stimuli of human faces and social behaviors of charitable giving. Consistent with prior works (e.g., Dotsch, Hassin, & Todorov, 2016; Velez & Gweon, 2020), our current results indicate that people do seem to employ basic statistical learning processes even in social domains.

On the other hand, past research on sample bias would predict that as long as the environmental data show no associations between variables, participants should be able to learn that fact. However, the current results (free sampling in Experiments 1 and 2, and Experiments 3, 4, 5, and 6) invalidate such predictions. Across all experiments, the prepared population dataset was designed to exhibit zero linear associations between faces and behaviors, but participants rarely learned that fact. The novel manipulation of granting the power of seeing more or fewer samples to participants enabled us to identify the underlying mechanism. Consistent with past

work (e.g., Klein and O'Brien, 2018; Prager, Krueger, & Fiedler, 2018), we found that participant-generated samples are much smaller than the samples needed to gain an accurate picture of the data. As a result, such insufficient sampling makes spurious beliefs sticky, and resistant to updating even when faced with the complete environmental data.

Why did participants stop sampling so early? The attitude asymmetry theory such as the BeanFest (Fazio, Eiser, & Shook, 2004) and impression formation theory such as the hot stove effect (Denrell, 2005) hypothesize that people stop sampling because they want to avoid negative outcomes. However, participants in our experiments did not receive any negative outcomes. Theories on outcome dependent sampling may suggest that unjustified confidence in decisions because of lack of disconfirming feedback may discourage sampling (illusion of validity in Einhorn & Hogarth, 1978; outcome-selective sampling in Fiedler, 2000). Yet, the question remains as why do people feel confident even though they are not seeing enough evidence? This is especially true for our experiments, which used only three examples. Taking a motivated tactician lens, the continuum model of impression formation may suggest that participants' motivation is insufficient, and that is why people do not bother to continue sampling (Fiske & Neuberg, 1991). But why do we need create sufficient motivation to sampling? From a rational learning perspective, the default should be to keep sampling until accurate decisions are warranted. Our work opens these theoretical questions for future investigations.

Importantly, the insufficient statistical learning perspective provides general and mechanistic explanations on why various types of stereotypes may prevail and persist even though those stereotypes could be inaccurate. For example, Insta-learn offers a domain-general explanation of the overgeneralization hypothesis in face stereotypes (baby-face

overgeneralization, emotion overgeneralization, etc., Zebrowitz, Fellous, Mignault, & Andreoletti, 2003; Zebrowitz & Montepare, 2008; Oosterhof & Todorov, 2008; Said, Sebe, & Todorov, 2008). Take the baby-face stereotype — i.e., baby-faced adults are perceived as submissive — as an example. According to our Insta-learn framework, this stereotype emerges because at first people learn the association between facial features and behaviors within the sample of babies. They learn that faces carrying baby-like features are associated with submissive qualities, and that this association is relatively stable across many babies. This parallels the initial learning phase in our experiment. Next, when seeing new baby-faced adults, people may stop sampling early; and not carefully observe those baby-faced adults' actual behaviors. This parallels the free sampling phase in our experiment. If people do not patiently observe and study a large number of baby-faced adults' behaviors, they will be less likely to revise the baby-face stereotypes. This parallels the final decision phase in our experiment. A concrete prediction from Insta-learn would be that there is individual heterogeneity in baby-face overgeneralization and this heterogeneity depends on how many baby-face adults one has actually observed or interacted with. If a person intensively interacts with many baby-faced adults, Insta-learn would predict that this person would not overgeneralize inappropriately, as compared to another person who does not interact with baby-faced adults at all. More interactions bring less biased impressions. Given that intergroup contact naturally brings more interactions, this barebone mechanism of Insta-learn might be one cognitive process that makes the contact hypothesis works (Allport, 1954).

Likewise, gender stereotypes of females being communal (Eagly & Wood, 2010) could originate from people quickly learning from observations of a small subset of females doing well

in assistant jobs, and then persist by failing to actively sample more females' and their actual behaviors. Similarly, racial stereotypes of Asians being competent but unfriendly (Fiske et al., 2002) could originate from people quickly learning from very small samples — as few as three Asian kids being smart but introverted in their own childhood — and then persist due to not actively sampling more Asians (18.64 million in the US and 4.3 billion in the world) and their actual behaviors. Even though competence and warmth may be completely orthogonal to race, people may be less likely to learn that because of insufficient statistical learning.

In addition to providing a general mechanistic explanation for stereotyping in social psychology, Insta-learn examines a boundary condition of statistical learning in traditional cognitive psychology. The question being asked here is when there is a mismatch between sample statistics and population statistics, what do humans learn? Humans make sense of the world from limited amounts of data (Tenenbaum et al., 2011). It is an incredible ability, but one that can have detrimental consequences under certain circumstances. In classic cognitive tasks, samples are often unbiased and resemble real-world statistics, such as object names or grammars (Xu & Tenenbaum, 2007; Saffran & Newport, 1999). However, in many social domains, samples do not resemble real-world statistics. Small samples can come from, for example, watching negative portraits of minority groups on news media, or hearing anecdotal stories about certain groups from families or friends. As intelligent humans, people learn the association between particular traits and minority groups quickly and accurately. However, as over-confident humans, they may stop sampling for information too early. If people are less likely to interact with those groups offline, or less likely to search about those groups online (additional information manipulation in Experiment 6), inaccurate stereotypes would remain. As shown in our studies, if

the sample statistics and the population statistics do not match, the same cognitive ability that makes humans so intelligent actually fuels the emergence of inaccurate stereotypes.

Finally, uncovering the emergence and persistence of stereotypes from small sample based fast learning and insufficient subsequent sampling perspective can offer policy insights. It is often assumed that prejudiced personality or malicious motivations cause stereotyping and discrimination. Policies thus are designed to train the individuals to be not biased (e.g., diversity training, Bezrukova, Spell, Perry, & Jehn, 2016). We agree with this possibility. However, here, we suggest an even more insidious possibility — even among people who are well-intentioned and well-motivated, inaccurate stereotypes can still emerge and persist. This perspective recommends different policies in addressing intergroup conflicts. Rather than targeting individuals — increasing intergroup knowledge or regulating prejudices (e.g., Chang et al., 2019) — our results indicate another effective intervention could aim to change the initial sample structure. For example, in a corporation, if new hires see Whites, Blacks, Hispanics, and Asians occupy higher status jobs, they will be less likely to form race-occupation stereotypes early in their careers. Moreover, there should be more incentives to motivate people to continue sampling more, even though they might think they already know enough. For example, inviting more members of underrepresented groups in STEM field workshops, symposiums, and talk series can be one way to expose audiences to more diverse samples. Diverse representation and active exploration matters (Bai, Ramos, & Fiske, 2020; Eagly & Koenig, 2021).

Our work can benefit from further replications to test the generalizability of the current hypothesis. First, our participants were all English-speaking adults. It would be important to test

non-English speaking participants and younger participants. Granted we have fewer theories or data on whether people living in different cultures endorse distinct types of statistical learning, cross-cultural investigations are needed to move forward. Second, our experiments were all conducted online. It would be critical to test offline decisions as well. For example, it is important to know whether participants would behave similarly if it were an in-person fundraising event where people get to mingle with each other, and the potential payoffs are high. Third, our experiments only tested one social behavior which is charitable donation vis-à-vis generosity. It is helpful to know whether the same mechanism applies to other social evaluative domains such as agency or trustworthiness. We would expect as the stakes become higher, we should be more likely to see Insta-learn processes in place — stop sampling early, but more direct tests are needed.

Conclusion

It is often celebrated that humans learn *so* much from *so* little. It is an incredible ability which allows human societies to thrive. However, maybe humans also learn *too* much from *too* little. The ability to learn quickly from small numbers of examples and insufficiently sample more examples can cause inaccurate stereotypes to emerge and persist.

References

- Allport, G. W. (1954). *The nature of prejudice*. Addison-Wesley.
- Bai, X., Ramos, M. R., & Fiske, S. T. (2020). As diversity increases, people paradoxically perceive social groups as more similar. *Proceedings of the National Academy of Sciences*, 117(23), 12741-12749.
- Berg, J., Dickhaut, J., & McCabe, K. (1995). Trust, reciprocity, and social history. *Games and economic behavior*, 10(1), 122-142.
- Bezrukova, K., Spell, C. S., Perry, J. L., & Jehn, K. A. (2016). A meta-analytical integration of over 40 years of research on diversity training evaluation. *Psychological Bulletin*, 142(11), 1227.
- Bodenhausen, G. V., & Morales, J. R. (2012). Social cognition and perception. *Handbook of Psychology, Second Edition*, 5.
- Chang, E. H., Milkman, K. L., Gromet, D. M., Rebele, R. W., Massey, C., Duckworth, A. L., & Grant, A. M. (2019). The mixed effects of online diversity training. *Proceedings of the National Academy of Sciences*, 116(16), 7778-7783.
- de Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a web browser. *Behavior Research Methods*, 47(1), 1-12.
- Dotsch, R., Hassin, R. R., & Todorov, A. (2016). Statistical learning shapes face evaluation. *Nature Human Behaviour*, 1(1), 1-6.
- Eagly, A. H., & Koenig, A. M. (2021). The Vicious Cycle Linking Stereotypes and Social Roles. *Current Directions in Psychological Science*, 09637214211013775.

Eagly, A. H., & Wood, W. (1999). The origins of sex differences in human behavior: Evolved dispositions versus social roles. *American psychologist*, 54(6), 408.

Einhorn, H. J., & Hogarth, R. M. (1978). Confidence in judgment: Persistence of the illusion of validity. *Psychological review*, 85(5), 395.

Eberhardt, J. L., Davies, P. G., Purdie-Vaughns, V. J., & Johnson, S. L. (2006). Looking deathworthy: Perceived stereotypicality of Black defendants predicts capital-sentencing outcomes. *Psychological science*, 17(5), 383-386.

Fazio, R. H., Eiser, J. R., & Shook, N. J. (2004). Attitude formation through exploration: Valence asymmetries. *Journal of personality and social psychology*, 87(3), 293.

Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to judgment biases. *Psychological review*, 107(4), 659.

Fiedler, K., & Freytag, P. (2004). Pseudocontingencies. *Journal of Personality and Social Psychology*, 87(4), 453.

Fiedler, K., & Juslin, P. (Eds.). (2006). *Information sampling and adaptive cognition*. Cambridge University Press.

Fiske, S. T., & Bai, X. (2020). Vertical and horizontal inequality are status and power differences: applications to stereotyping by competence and warmth. *Current Opinion in Psychology*, 33, 216-221.

Fiske, S. T., Cuddy, A. J., Glick, P., & Xu, J. (2002). A model of (often mixed) stereotype content: competence and warmth respectively follow from perceived status and competition. *Journal of personality and social psychology*, 82(6), 878.

- Fiske, S. T., & Neuberg, S. L. (1990). A continuum of impression formation, from category-based to individuating processes: Influences of information and motivation on attention and interpretation. *Advances in experimental social psychology*, 23, 1-74.
- Fiske, S. T., & Taylor, S. E. (1984/2020). *Social cognition*. Reading, Mass: Addison-Wesley Pub. Co.
- Hamilton, D. L., & Gifford, R. K. (1976). Illusory correlation in interpersonal perception: A cognitive basis of stereotypic judgments. *Journal of Experimental Social Psychology*, 12(4), 392-407.
- Hamilton, D. L., & Sherman, J. W. (1994). *Stereotypes*. In R. S. Wyer, Jr. & T. K. Srull (Eds.), *Handbook of social cognition: Basic processes; Applications* (p. 1-68). Lawrence Erlbaum Associates, Inc.
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). Model assessment and selection. In *The elements of statistical learning* (pp. 485-585). Springer, New York, NY.
- Kalish, M. L., Griffiths, T. L., & Lewandowsky, S. (2007). Iterated learning: Intergenerational knowledge transmission reveals inductive biases. *Psychonomic Bulletin & Review*, 14(2), 288-294.
- Klein, N., & O'Brien, E. (2018). People use less information than they think to make up their minds. *Proceedings of the National Academy of Sciences*, 115(52), 13222-13227.
- Bertrand, M., & Mullainathan, S. (2004). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American economic review*, 94(4), 991-1013.

- Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences*, 105(32), 11087-11092.
- Prager, J., Krueger, J. I., & Fiedler, K. (2018). Towards a deeper understanding of impression formation—New insights gained from a cognitive-ecological perspective. *Journal of personality and social psychology*, 115(3), 379.
- Lick, D.J., Alter, A.L., & Freeman, J.B. (2017). Superior pattern detectors efficiently learn, Activate, Apply, and Update social stereotypes. *Journal of Experimental Psychology: General*. 147 (2), 209-227.
- Litman, L., Robinson, J., & Abberbock, T. (2017). TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods*, 49(2), 433-442
- Roets, A., & Van Hiel, A. (2011). Allport's prejudiced personality today: Need for closure as the motivated cognitive basis of prejudice. *Current Directions in Psychological Science*, 20(6), 349-354.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274(5294), 1926-1928.
- Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition*, 70(1), 27-52.
- Said, C. P., Sebe, N., & Todorov, A. (2009). Structural resemblance to emotional expressions predicts evaluation of emotionally neutral faces. *Emotion*, 9(2), 260.
- Stirrat, M., & Perrett, D. I. (2010). Valid facial cues to cooperation and trust: Male facial width and trustworthiness. *Psychological science*, 21(3), 349-354.

- Tajfel, H., Turner, J. C., Austin, W. G., & Worchel, S. (1979). An integrative theory of intergroup conflict. *Organizational identity: A reader*, 56, 65.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *science*, 331(6022), 1279-1285.
- Todorov A, Oh. D. (2021) The structure and perceptual basis of social judgments from faces. *Advances in experimental social psychology*, Vol 63.
- Todorov, A., Mandisodza, A. N., Goren, A., & Hall, C. C. (2005). Inferences of competence from faces predict election outcomes. *Science*, 308(5728), 1623-1626.
- Todorov, A., Olivola, C. Y., Dotsch, R., & Mende-Siedlecki, P. (2015). Social attributions from faces: Determinants, consequences, accuracy, and functional significance. *Annual review of psychology*, 66.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *science*, 185(4157), 1124-1131.
- Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological review*, 114(2), 245.
- Vélez, N., & Gweon, H. (2020). Learning from other minds: An optimistic critique of reinforcement learning models of social learning. *CogSci Proceedings*.
- Zebrowitz, L. A., Fellous, J. M., Mignault, A., & Andreoletti, C. (2003). Trait impressions as overgeneralized responses to adaptively significant facial qualities: Evidence from connectionist modeling. *Personality and social psychology review*, 7(3), 194-215.
- Zebrowitz, L. A., & Montepare, J. M. (2008). Social psychological face perception: Why appearance matters. *Social and personality psychology compass*, 2(3), 1497-1517.

Zhou, H., & Fishbach, A. (2016). The pitfall of experimenting on the web: How unattended selective attrition leads to surprising (yet false) research conclusions. *Journal of Personality and Social Psychology*, *111*(4), 493-504.