Visualization aesthetics bias trust in science, news, and social media

Authors: Chujun Lin*, Mark Thornton

Affiliations:

Department of Psychological and Brain Sciences, Dartmouth College; Hanover, NH, USA.

* Corresponding author. Email: Chujun.Lin@Dartmouth.edu

Abstract: Scientists, policymakers, and the public rely on data visualizations to inform important decisions. The aesthetics of graphs may vary independent of data quality. Across five-preregistered studies (3,405 participants, 1,219 graphs), we comprehensively investigate whether the beauty of a graph influences how much people trust it. Results show that graph beauty predicts increased trust using real-world graphs across three important domains: social media, news, and scientific publications (Studies 1-3). This association had real-world correlates: graph beauty predicts increased citation and comment numbers. We next identified visual features that contribute to perceptions of graph beauty (Study 4). Using these features, we digitally manipulated graph beauty and showed that making a graph more beautiful causally increased trust (Study 5). This causal effect was driven by people's heuristic that beautiful graphs were made by competent people (Study 5b). These findings reveal a novel source of bias and indicate the importance of promoting visualization literacy.

Main

Data are reshaping every aspect of human life: consumers buy products guided by online reviews, drivers navigate to novel places relying on satellite data, and governments design cities based on local activity patterns. The explosion of big data calls for tools that digest and interpret data to better inform decision making. Among these tools are data visualizations, such as charts of COVID cases, weather forecasts, stock prices, and election outcomes. The popularity of data visualization has given rise to specialized platforms on social media (e.g., the “Data Is Beautiful” community on Reddit), revolutions in journalism (e.g., visual and data journalism), and new disciplines in universities (e.g., doctorates in data visualization). How might data visualizations influence whether people trust what they see on social media, the news, and in scientific articles?

Decisions about how to translate quantitative data into visualizations may influence people’s trust in the information. Such choices might include which statistics to visualize (e.g., mean, trend, correlation), what type of graph to use (e.g., bar, line, map), or how to aesthetically style the figure. Good choices make it easier for readers to understand the data¹⁻³, whereas bad choices lead to confusion and inaccessibility⁴⁻⁹. For instance, poor color choices such as pairing equal luminance red and green can hinder people with color-vision deficits from extracting useful information from graphs^{5,10}.

Aesthetic decisions influence not only the legibility of graphs, but also their visual appeal, or beauty. Choices such as color palette, spacing, font type, shape, and outlining can all influence whether a graph is perceived as beautiful or ugly, independent of the quality of the underlying data or how accurately the graph reflects those data. Recent advances in visualization methods have substantially enhanced people’s ability to produce beautiful graphs and lowered the technical barriers to doing so¹¹⁻¹³.

1 Many regard beauty as irrelevant to the interpretation and trust of graphs^{14–16}. However, others
2 have noted the power that aesthetically appealing figures can have, both for good and ill^{17–19}. For
3 instance, some have speculated that graph beauty may bias the scientific review process and
4 cause reviewers and editors to miss substantive flaws in research. There is good reason to believe
5 that this might be true: psychological research indicates that people are susceptible to a bias
6 known as the “halo effect” – and more specifically, the “beauty-is-good” stereotype²⁰. These
7 refer to people’s tendency to overgeneralize from one good feature of a person to other, unrelated
8 features. For example, people often believe that the individuals that they perceive as beautiful are
9 also trustworthy and intelligent, without any direct evidence^{21–24}. If this bias extends to data
10 visualizations, then we would expect that people will trust graphs more when they are more
11 beautiful, regardless of the quality of the data or whether the graph accurately portrays them.

12 Here we test the hypothesis that the beauty of data visualizations influences how much people
13 trust them. We first examined the correlation between perceived beauty and trust in graphs
14 (Studies 1-3). To maximize the generalizability and external validity of our findings, we
15 systematically sampled graphs (Fig. 1) of diverse types and topics (Fig. 2) from the real world.
16 These graphs spanned a wide range of domains, including social media (Study 1, 439 graphs),
17 news reports (Study 2, 110 graphs), and scientific publications (Study 3, 310 graphs). We asked
18 participants how beautiful they thought the graphs looked and how much they trusted the graphs.
19 We also measured how much participants felt that they understood the graphs, and found the
20 information interesting, surprising, and unpleasant, to control for the potential confounding
21 effects of these factors (Fig. 3A).

22 To understand whether visualization aesthetics have behavioral consequences in the real world,
23 we next examined the correlation between perceived beauty and the graphs’ real-world impact.

1 We measured real-world impact for all real-world graphs in Studies 1-3 using indices such as the
2 number of comments the graphs received on social media, and the number of citations the
3 graphs' associated papers had. We compared these measures of real-world impact to the beauty
4 ratings of the graphs given by our participants (Studies 1-3).

5 After establishing the association between graph beauty and trust in both our participants and in
6 naturally occurring data, we examined whether this association was indeed causal. We first
7 applied computational models to characterize the visual features that predicted how beautiful
8 people thought a graph looked (Study 4). Based on these features, we digitally manipulated the
9 beauty of the graphs that we generated using arbitrary data and measured how this manipulation
10 changed participants' trust in the graph (Studies 5a & b, 216 graphs). Finally, we examined a
11 potential explanation for why people trust beautiful graphs more: we measured how competent
12 people inferred the graph-maker to be upon viewing a graph, and tested whether this inferred
13 competence statistically mediated the causal effect of graph beauty on trust (Study 5b).

Figure 1. Systematic selections of real-world graphs. (A) Inclusions and exclusions of graphs from the Reddit community Data Is Beautiful in Study 1a. (B) Inclusions and exclusions of graphs from the Reddit community Data Is Beautiful in Study 1b. (C) Selections of news reports from online news outlets in Study 2. (D) Selections of academic journals, and inclusions and exclusions of their papers in Study 3. Some papers were associated with multiple subject areas.

Figure 2. Types and topics of real-world graphs in Studies 1-3. (A) Types and topics of the 439 graphs in Study 1 (left; including 139 graphs in Study 1a and 300 graphs in Study 1b) and two examples of the graphs (right). (B) Types and topics of the 110 graphs in Study 2 (left) and two examples of the graphs (right). (C) Types and subject areas of the 310 graphs in Study 3 (left). Most graphs included multiple panels; four examples of the panels are shown (right). Some of these examples (their papers) were associated with multiple subject areas; only one subject area is indicated here for simplicity.

Results

Beauty correlates with trust across graph types, topics, and origins

We found that participants' trust in graphs was associated with how beautiful participants thought the graphs looked for graphs across all three domains (Fig. 3B): social media posts on Reddit (Pearson's $r = 0.45$, $p = 4.15 \times 10^{-127}$ in Study 1a; $r = 0.41$, $p = 3.28 \times 10^{-231}$ in Study 1b), news reports ($r = 0.43$, $p = 1.14 \times 10^{-278}$ in Study 2), and scientific papers ($r = 0.41$, $p = 6.55 \times 10^{-234}$ in Study 3). These findings indicate that, across diverse contents and sources of the graphs, perceived beauty and trust in graphs are reliably correlated in the minds of perceivers.

The association between beauty and trust remained robust when controlling for factors that might influence both perceived beauty and trust, including how much participants felt that they understood the graphs, and found the information interesting, surprising, and unpleasant (linear mixed modeling: $b = 0.19$, standardized $\beta = 0.22$, $p = 1.05 \times 10^{-30}$ in Study 1a; $b = 0.14$, $\beta = 0.16$, $p = 8.81 \times 10^{-46}$ in Study 1b; $b = 0.14$, $\beta = 0.15$, $p = 5.35 \times 10^{-35}$ in Study 2; $b = 0.10$, $\beta = 0.12$, $p = 1.85 \times 10^{-25}$ in Study 3; see Fig. S1 for the coefficients of covariates). These findings indicate that beautiful visualizations predict increased trust even when controlling for confounding factors.

The association between beauty and trust remained robust when correcting for common method variance (i.e., a participant might use the beauty rating of a graph to inform trust rating, and vice

versa). How beautiful a participant rated a graph to be was significantly correlated with the aggregate trust rating given by the rest of the participants ($r = 0.14, p = 7.71 \times 10^{-13}$ in Study 1a; $r = 0.12, p = 4.22 \times 10^{-21}$ in Study 1b; $r = 0.13, p = 7.16 \times 10^{-25}$ in Study 2; $r = 0.05, p = 1.26 \times 10^{-4}$ in Study 3). These results indicate that beautiful visualizations predict increased trust even when controlling for within-participant measurement biases.

Figure 3. Correlations between beauty and trust in Studies 1-3. (A) Participants viewed each graph (top; an example from Study 3) and rated each graph on six aspects (bottom; the order was randomized). (B) The frequency of ratings (colored; presented with 2D kernel density) on the beauty and trust of the graphs in Studies 1a, 1b, 2, and 3 (from top to bottom), and univariate correlations between the two variables (line for linear regression, text for Pearson's correlation,

asterisks indicate statistical significance: *** for $p < 0.001$; N = 2,681 observations in Study 1a; N = 5,780 in Study 1b; N = 6,204 in Study 2; N = 6,030 in Study 3).

Beauty predicts real-world impact

We next moved beyond self-reported data and investigated the real-world behavioral consequences of the trust bias in beautiful graphs. We collected naturally occurring data about all graphs in Studies 1-3 (e.g., comments of the graphs on Reddit) and compared them to our participants' beauty ratings. We found that the real-world impact of the graphs was associated with how beautiful participants thought they were (Fig. 4).

The more beautiful graphs from Reddit were associated with higher numbers of comments in both Study 1a ($b = 0.04$, $\beta = 0.04$, $p = 0.011$) and Study 1b ($b = 0.11$, $\beta = 0.12$, $p = 2.84 \times 10^{-22}$).

The more beautiful graphs from scientific journals were associated with papers that had higher numbers of citations in Study 3 ($b = 0.07$, $\beta = 0.05$, $p = 0.001$; but not higher numbers of views, $b = 0.03$, $\beta = 0.02$, $p = 0.264$). The association between the perceived beauty of a paper's graphs and the paper's number of citations remained robust when controlling for publication date, how much participants felt that they understood the graphs, and found the information interesting, surprising, and unpleasant, and the random effects of participants and the papers' subject areas ($b = 0.05$, $\beta = 0.04$, $p = 0.005$). The association between graph beauty and paper citation was even more substantial when data were aggregated across papers per subject area (Fig. 4B). These findings suggest that people's bias in favor of trusting beautiful graphs has real-world behavioral consequences: beauty biases what information people attend to and endorse.

Figure 4. Correlations between beauty and real-world metrics. (A) Top: Pearson's correlation between the number of comments (log-transformed to correct for skewness) and the graph's average beauty rating by participants for 439 graphs in Studies 1a & b. Bottom: Pearson's correlation between the number of citations of the graph's associated paper (log-transformed to correct for skewness) and the graph's average beauty ratings by participants for 310 graphs in Study 3. Asterisks indicate statistical significance: *** for $p < 0.001$, ** for $p < 0.01$. (B) Pearson's correlation between the average number of citations and the average beauty ratings for 22 subject areas in Study 3. The line indicates linear regression, with the 95% confidence interval in shadow.

We did not find evidence that the more beautiful graphs from news reports were associated with higher real-world impact in Study 2 ($b = 0.04$, $\beta = 0.02$, $p = 0.074$). This null result may be due to the estimation of impact: we approximated the news reports' real-world impact using the counts of comments, shares, and likes about the news reports on four social media platforms (Facebook, Twitter, Pinterest, Reddit) tracked by a third-party app (<https://buzzsumo.com/>) because these data were not available on the news sources themselves. These counts were likely noisy and incomplete.

Visual features predict beauty

We investigated what people perceived as “beautiful” in the context of data visualizations. We applied computational models in the Open Source Computer Vision Library (OpenCV)²⁵ and Google’s Tesseract-Optical Character Recognition Engine (Python-tesseract)²⁶ to extract from each graph its resolution, the distributions (across all pixels of the image) of hue, saturation, luminance, and contrast, and word counts (e.g., title, explanatory text). Hue, luminance, and saturation were the strongest predictors of perceived beauty, either with only their summary statistics of mean and standard deviation (Fig. 5A; adjusted $R^2 = 0.15$) or together with their full frequency distributions over all possible values (e.g., how frequently each value of luminance from 0 to 255 occurred in a given image; Fig. 5B; adjusted $R^2 = 0.39$). Graphs using vibrant colors were perceived as most beautiful (Fig. 5C); whereas, graphs using black, white, and red colors were perceived as least beautiful (Fig. 5D).

Figure 5. Predictors of graph beauty in Studies 1-3. (A) Coefficients from linear regression of beauty on features across $N = 859$ graphs in Studies 1-3. Thick lines indicate 90% confidence intervals; thin lines and density indicate 95% confidence intervals. (B) Unique explained variance of each type of features as assessed with adjusted R^2 in linear regressions. Hue, saturation, luminance, and contrast are characterized by the full frequency distribution across intensity values, as well as their mean and standard deviation. (C) Uniform manifold approximation and projection (UMAP) of the top 10,000 colors that were most positively correlated with beauty. Each color was specified by a combination of hue, saturation, and luminance values. Size indicates the magnitude of the correlation (Fisher's z-transformed). (D) UMAP of the top 10,000 colors that were most negatively correlated with beauty.

The causal effect of beauty on trust

After identifying what made a graph beautiful and establishing the association between graph beauty and trust in both self-reported data and naturally occurring data, we examined whether

1 this association was causal (Studies 5a & b). We experimentally manipulated graph beauty based
2 on beauty-related visual features discovered in Study 4. Specifically, we generated graphs of
3 diverse types (bar, line, map, scatter, pie, heat) and topics (politics, sports, science, economics,
4 health, and nature) using arbitrary data. For each graph of a certain type and topic, we
5 manipulated the graph beauty based on color, style, and resolution to generate two versions (Fig.
6 6A; Study 5a also manipulated the font type and size). This procedure generated 144 graphs in
7 Study 5a and 72 graphs in Study 5b. We collected beauty and trust ratings about these graphs
8 from participants in two separate blocks. Our beauty manipulation was successful: the more
9 beautiful version of a graph on average received about 1-point higher beauty rating (on a 7-point
10 scale) than the less beautiful version ($b = 0.98$, $\beta = 0.56$, $p = 3.77 \times 10^{-62}$ in Study 5a; $b = 1.27$,
11 $\beta = 0.70$, $p = 2.59 \times 10^{-84}$ in Study 5b).

12 Making a graph more beautiful caused people to trust it more (Fig. 6B): the more beautiful
13 version of a graph on average received half-a-point higher trust rating than the less beautiful
14 version ($b = 0.50$, $\beta = 0.32$, $p = 1.35 \times 10^{-18}$ in Study 5a), controlling for the random effects of
15 individual participants, graph types, and graph topics. This causal effect remained significant
16 when the beauty manipulation did not include font changes ($b = 0.28$, $\beta = 0.17$, $p = 1.08 \times 10^{-5}$
17 in Study 5b). These findings support a causal interpretation of the correlations observed in
18 Studies 1-3, indicating that graph beauty is indeed responsible for increasing trust.

19 We further characterized how the causal effect of graph beauty on trust might be moderated by
20 five factors: the objective misleadingness of the graphs (Study S1, 4 graphs), the familiarity of
21 the topics of the graphs (Study S2, 4 graphs), the between- versus within-subject design of the
22 study (Study S3, 24 graphs), the different components of beauty manipulation of the graphs
23 (Study S4, 40 graphs), and the statistical power of the study (Study S5, 72 graphs). The details of

these experiments are reported in the Supplemental Information. The effect of beauty on trust was positive in all of these five supplementary studies but was not statistically significant in four of them (except for Study S5; see Table S1). This is likely because these supplementary studies used considerably smaller sample sizes (in the number of graphs and participants) than Studies 5a-b and were thus underpowered (see Table S1). For instance, although the causal effect of beauty on trust in Study S4 was not statistically significant, its effect size ($b = 0.305$) was larger than those in Study 5b and Study S5 ($bs = 0.279$ and 0.241 ; see Table S1). A mega-analysis of all data in these supplementary studies indicated a significant effect of beauty on trust ($b = 0.16$, $p = 2.35 \times 10^{-4}$; see Fig. 6C), even without the inclusion of Studies 5a-b. Thus, these studies further support the conclusion that beauty causally influences trust in graphs.

We investigated why people might trust more beautiful graphs more (Study 5b). Based on the correspondent inference theory²⁷, people might use graph beauty as a cue to infer the graph maker's disposition (e.g., competence). Therefore, we asked participants, "How competent is the person who made the graph? That is, based on the graph, how efficient and capable do you think the graph maker is to do things in general (e.g., to conduct research)?" The causal effect of beauty on trust was fully statistically mediated by people's inferences about the general competence of the graph maker (Fig. 6D; average causal mediation effect $b = 0.27$, total effect $b = 0.28$, proportion mediated = 97.4%). Making a graph look more beautiful caused people to infer that the graph-maker was more competent ($b = 0.45$, $p = 2.18 \times 10^{-11}$), and this inference was associated with increased trust ($b = 0.61$, $p = 5.70 \times 10^{-210}$). After controlling for the effect of beauty through competence inferences, there was no remaining significant effect of beauty on trust ($b = 0.01$, $p = 0.899$). These findings suggest that more beautiful designs improve trust only when they make people think that the graph-maker is competent (e.g., a beautiful but childish-looking graph might not increase trust).

Figure 6. Causal effects of beauty on trust. (A) Manipulations of one example graph of a specific type and topic in the two experimental conditions. (B) Left: Manipulation check of beauty in Study 5a (purple) and Study 5b (magenta). The b s indicate the coefficients from linear mixed model regressions of participants' rated beauty (7-point Likert scale) on beauty manipulations (binary) across $N = 2,574$ observations in Study 5a and $N = 2,424$ observations in Study 5b. The random effects of participants, graph types, and graph topics are included in all models as random intercepts. The d s indicate the Cohen's D effect sizes across all observations. $***$ for $p < 0.001$. Right: Causal effects of beauty on trust in Study 5a (purple) and Study 5b (magenta). The b s indicate the coefficients from linear mixed model regressions of participants' reported trust (7-point Likert scale) on beauty manipulations (binary). (C) Left: Manipulation check of beauty in supplementary Studies S1-S5 (grey) and the mega-analysis pooling all data across the five supplementary studies (orange). The b indicates the coefficient from linear mixed model regressions of participants' rated beauty on beauty manipulations using $N = 4,386$ observations pooling across Study S1 ($N = 490$), Study S2 ($N = 139$), Study S3 ($N = 1,080$), Study S4 ($N = 715$), and Study S5 ($N = 1,962$). Right: Causal effect of beauty on trust in Studies S1-S5 (grey) and the mega-analysis (orange). (D) Mediation analysis on the causal effect of beauty manipulation on trust ratings ($N = 2,283$ complete observations in Study 5b).

No evidence that education reduces bias

We observed no evidence that greater educational attainment mitigates the graph beauty bias.

Rather, we observed weak evidence that more educated people may be even more susceptible to perceiving beautiful graphs as trustworthy. In Study 1b, participants' highest level of education positively moderated the association between trust in graphs and how much participants thought the graphs were beautiful ($\beta = 0.037, p = 0.013$). In the other correlational studies, the moderation effect of education was in the same direction, but not statistically significant ($\beta = 0.003, p = 0.884$ in Study 1a; $\beta = 0.019, p = 0.217$ in Study 2; not tested in Study 3 since all participants were required to be college graduates). Education did not significantly moderate the causal effect of beauty on trust either ($\beta = -0.069, p = 0.057$ in Study 5a; $\beta = 0.047, p = 0.220$ in Study 5b). These findings suggest that receiving more overall education does not alleviate people's trust bias in favor of beautiful graphs; worse yet, more education may even exacerbate this bias in some cases.

We also found no evidence that more specific education targeting data visualizations reduces bias either (Study 5b): these measures included how often participants saw data visualizations ($\beta = 0.057, p = 0.138$), made data visualizations ($\beta = 0.061, p = 0.110$), whether they received education on data visualizations ($\beta = 0.102, p = 0.228$), read books or tutorials ($\beta = -0.070, p = 0.367$), took classes ($\beta = 0.116, p = 0.133$), received training ($\beta = 0.118, p = 0.253$), and received related degrees ($\beta = -0.288, p = 0.200$). These findings suggest that experience and education with data visualizations do not automatically reduce the trust bias in more beautiful graphs.

We also tested the moderation effect of other factors, including political affiliation and personality (agreeableness and conscientiousness) for graphs from news reports in Study 2. None

of these factors had a significant moderation effect ($\beta = 0.01$, $p = 0.500$ for political affiliation; $\beta = -0.02$, $p = 0.181$ for agreeableness; $\beta = 0.03$, $p = 0.063$ for conscientiousness). These results indicate that there is no evidence that the extent to which an individual perceives more beautiful graphs as more trustworthy would be influenced by whether the individual is republican or democrat, agreeable or critical, and conscientious or careless.

Discussion

This investigation characterized the relationship between the beauty of data visualizations and how much people trust them. We found that the more beautiful people perceived a graph to be, the more they trusted it (Fig. 3B). This association was robust across graphs of diverse content, including entertaining topics on social media (Fig. 2A), news on business, politics, health, technology, and sports (Fig. 2B), and sophisticated scientific findings across a range of disciplines (Fig. 2C). The effect sizes were large and highly consistent across studies, explaining 16% to 20% of the variance in trust (Fig. 3B).

We showed that this association was causal: visualizing the same data in a more beautiful way caused people to trust it more (Fig. 6B). The causal effect of beauty on trust was fully mediated by inferred competence of the graph-makers (Fig., 6C). These findings indicate that decisions about visualization aesthetics influence people's interpretation of the data. People trust the information more when presented in a more beautiful way that conveyed the competence of the graph-maker.

Two mechanisms may underlie this bias. First, perceiving more beautiful graphs as more trustworthy may be an overgeneralization of the beauty-is-good stereotype. This stereotype was originally discovered in how people perceive other people²⁰. It describes the common belief that

1 physically attractive people also have other positive social attributes such as being more
2 trustworthy and friendlier^{21,28}. People are also influenced by this stereotype when evaluating
3 non-human objects. For instance, people associate more beautiful products, packaging, and web
4 pages with better performance^{29–31}. The trust bias in favor of more beautiful graphs found here
5 may be another manifestation of the beauty-is-good stereotype.

6 A second, non-mutually exclusive, explanation suggests that this apparent bias may be rooted in
7 rational thinking. More beautiful graphs may indicate that the data is of higher quality and that
8 the graph maker is more skillful. However, this reasoning is not necessarily accurate. It does not
9 require sophisticated techniques to make more beautiful graphs: simply changing the color of the
10 graph shifts people's perception of graph beauty (Fig. 5); we easily generated graphs of different
11 degrees of beauty using arbitrary data (Fig. 6A). Therefore, our findings highlight that the beauty
12 of a graph may not be an informative cue for its quality. Even if beauty was correlated with
13 actual data quality in the real world, this would be a dangerous and fallible heuristic to rely upon
14 for evaluating research and media.

15 Importantly, we showed that the trust bias in favor of beautiful graphs predicted behaviors in the
16 real world (Fig. 4). Internet users were more likely to comment on social media posts that
17 contained more beautiful graphs (Study 1). This behavioral tendency may accelerate the
18 circulation of less credible but more beautifully presented information on the internet, which may
19 in turn exerts a greater influence on public opinions. We also found that publications containing
20 more beautiful graphs were associated with a higher number of citations across a range of
21 disciplines (Study 3). Although we cannot verify whether papers with more beautiful graphs
22 were more scientifically sound, this finding suggests that researchers may be biased by appealing
23 data visualizations when evaluating the scientific merit of a study¹⁸.

1 These findings highlight the need to educate both the general public and the scientific
2 community on data visualizations³². Our findings suggest that current curricula may not be
3 addressing this need: we found no evidence that higher levels of education reduce people's trust
4 bias in favor of beautiful graphs, and in one study we instead found that education increased this
5 bias. These findings call for updates in course materials to foster better skills in interpreting data
6 visualizations. Some colleges and universities have already added courses and degrees in data
7 visualizations. However, given the ubiquity of data visualizations in everyday life, it will be
8 helpful to promote data literacy from an earlier stage, such as high school³³.

9 We maximized the generalizability of our findings by sampling diverse graphs from the real
10 world. However, three limitations constrain our conclusions. First, all participants were recruited
11 online in the United States. They were not representative of the U.S. population or humans in
12 general. That said, this limitation may have less of an impact on the generalizability of our
13 findings than on other social science research^{34,35} because the populations that were not included
14 in our samples (e.g., people without internet access) are by the same token also less likely to be
15 exposed to data visualizations in the course of everyday life. In other words, internet users are
16 not only convenient subjects – they are also a population of particular interest for this research.
17 Second, although our participants varied considerably in education level, we did not specifically
18 sample scientists or other data visualization experts. As a result, our present data cannot
19 determine whether professional training in data visualizations could override people's bias in
20 favor of beautiful graphs. That said, the presence of an effect of beauty on citations provides
21 some evidence that this bias may persist among those populations. Finally, in all studies, we
22 presented graphs with minimal contextual information – at most the titles and abstracts of
23 scientific articles. It remains an open question how much people may be biased by the beauty of
24 data visualizations when they know the sources of the graphs or read the full media posts and

research articles. For example, on Reddit, popular (and hence highly visible) comments often call out poor data visualization practices. Readers who check these comments may have their beauty bias somewhat deflated, relative to seeing the graph in isolation.

Altogether, the present investigation offers strong empirical evidence that people perceive more beautiful graphs as more trustworthy across a wide range of sources, topics, and graph types. These findings identify a common bias that may influence public trust and misinformation in science, news, and social media. The current study presents a systematic approach for estimating such biases with high internal and external validity. We hope that others can adapt this approach to characterize other ways in which people (mis)interpret graphs. By doing so, the field can develop a comprehensive understanding of the psychology of data visualizations and generate strategies that help people better navigate today's big data world.

Methods

All data collections in our main studies were pre-registered except for Study 1a. The registrations of methods and planned analyses can be accessed on the Open Science Framework (OSF):

https://osf.io/px4mt/?view_only=7641adc103104b15a81d3c46bdf84236 for Study 1b,

https://osf.io/p7fsm/?view_only=4f32b6ab90504c278077b5c9c8bcd561 for Study 2,

https://osf.io/rgves/?view_only=4735e9c6dfae4bc082afa45c83cb51f0 for Study 3,

https://osf.io/bekz8/?view_only=fc1517e5b85948d48811b93976d66f22 for Study 5a, and

https://osf.io/4ya2f/?view_only=ec862e8095514c0a81d966b5f4a1f39f for Study 5b. We did not

deviate from our planned data collection and analysis methods in these studies. We also

conducted and reported exploratory analyses – the Pearson correlations between beauty ratings

and trust ratings (Fig. 3B), the Pearson correlation between beauty ratings and trust ratings when

correcting for common method variance, the Pearson correlations between beauty ratings and

real-world measures (Fig. 4), and the moderation effects of education specific to data visualizations. Study 4 analyzed existing data collected in Studies 1-3 and thus was not preregistered. One of our supplementary studies (Study S1) was pre-registered (see Supplemental Information). Participants in all studies provided informed consent in a manner approved by the Committee for the Protection of Human Subjects of Dartmouth College. All materials, data, and code can be accessed on OSF:

https://osf.io/yutgx/?view_only=b679d0841fd440eb901a28cecf5908b9.

Study 1a

Participants

Participants were recruited from Amazon Mechanical Turk (MTurk) via Cloud Research (formerly known as TurkPrime) ³⁶. We recruited participants based on the following inclusion criteria. Participants were required to be i) aged 18 and older, with ii) normal or corrected-to-normal vision, iii) at least high school education, iv) a good performance history on Amazon Mechanical Turk (approval rate $\geq 99\%$, previous submissions ≥ 50), and v) located in the United States.

Study 1a was a pilot study. We recruited 141 participants in total (Table 1), with around 20 participants completing each of the seven experiment modules (in which participants rated a subset of around 20 graphs; see Materials below). Participant demographics are as follow: 58 women, 76 men, 2 non-binary; Age [range = 19 to 70, $M = 40$, $SD = 12$, ten participants refused to answer]; 107 White, 17 Black, 13 Asian, 4 others.

We processed the data according to the following exclusion criteria. A trial was excluded if i) the response time for examining the graph was $\leq 1s$; ii) the response time for rating the graph on all five aspects was $\leq 5s$; iii) the total response time for examining and rating the graph was $\leq 6s$ or

≥ 160 s. A participant was excluded if more than half of their trials were excluded. According to these criteria, $n = 166$ out of 2,800 trials were excluded, and $n = 6$ out of $N = 141$ participants were excluded. The final sample size was 135 participants (see Table 1 for demographics).

Table 1. Summary of sample size and demographics.

Studies	Power Analysis	Pre-registered Recruitment Plan	Total	Exclude	Final	Gender	Age	Race
Study 1a	N/A	N/A	141	4%	135	57 women 71 men 2 non-binary	19 to 70 M = 40 SD = 12	105 White 14 Black 12 Asian 4 Others
Study 1b	$N \geq 280$ across 15 modules	20 per module if exclude < 6%; otherwise, recruit until final size = 19	321	10%	289	103 women 185 men 1 non-binary	18 to 72 M = 39 SD = 12	234 White 19 Black 25 Asian 11 Others
Study 2	$N \geq 280$ across 5 modules	60 per module if exclude < 6%; otherwise, recruit until final size = 56	328	14%	282	106 women 175 men 1 non-binary	21 to 69 M = 39 SD = 11	231 White 29 Black 17 Asian 5 Others
Study 3	$N \geq 280$ across 15 modules	20 per module if exclude < 6%; otherwise, recruit until final size = 19	334	13%	292	120 women 171 men 1 non-binary	21 to 75 M = 38 SD = 11	230 White 34 Black 20 Asian 8 Others
Study 5a	$N \geq 400$	500 if exclude < 20%; otherwise, recruit until final size = 400	501	14%	429	175 women 253 men 1 non-binary	19 to 74 M = 39 SD = 11	335 White 48 Black 29 Asian 17 Others
Study 5b	$N \geq 400$	500 if exclude < 20%; otherwise, recruit until final size = 400	501	19%	404	233 women 166 men 4 non-binary	19 to 75 M = 42 SD = 13	320 White 36 Black 28 Asian 15 Others

Note: All survey questions were voluntary. Non-responses were omitted.

Materials

Study 1a used 139 graphs from the Reddit community of Data Is Beautiful (r/DataIsBeautiful) as stimuli. These graphs were randomly assigned to 7 modules, with six modules containing 20 graphs and one module containing 19 graphs. We selected the 139 graphs from a much larger set based on 9 criteria (Fig. 1A). We i) first included the 486 most popular posts in Data Is Beautiful from Sep-Oct, 2020 by applying the filters “Top” and “This Month”. From these posts, we excluded posts ii) that did not include valid media components, which indicated that no graphs were in these posts; iii) whose media components were not static images, since we only focused

on static graphs; and iv) that were not created originally by Reddit users. We extracted the graphs from the remaining posts, and excluded graphs v) that were too complex, such as long graphs that expanded multiple pages; vi) that were not in English; vii) whose text labels were not clear, such as displaying excessive text in illegible font sizes. These selection processes retained 291 graphs. Since we targeted around 100 graphs in this pilot study, we further viii) excluded graphs with lower popularity (“upvote” was not above the median), and ix) for graphs depicting the exact same content (e.g., all locations of Costco near the user), included only the most popular one (highest “upvote”). These processes selected the final set of 139 graphs.

Procedures

Participants were randomly assigned to one of the seven modules. In each module, participants viewed a subset of graphs and evaluated them on five aspects: i) beauty, ii) trust, iii) interest, iv) understanding, and v) surprise. In each trial, participants first viewed the title of the Reddit post (from which the graph was extracted). Participants clicked a button when they finished reading the title. Participants then viewed the graph that appeared below the title. Participants clicked a button when they finished examining the graph. Then the five questions appeared below the graph (as in Fig. 3A, except that the question about unpleasantness was not asked in the present study). The order of the five questions was randomized for each trial and each participant. Participants indicated their answers for each question by clicking the radio buttons on a 7-point Likert scale, anchored at 1 = not at all and 7 = extremely. Participants clicked a button to confirm that they had finished answering all questions, and moved on to the next trial. The order of trials (i.e., graphs) was randomized for each participant. After completing all trials, participants filled out a brief survey about their demographics (age, gender, highest level of education, race, and ethnicity).

Statistical analyses

I. Association between perceived beauty and trust

We tested whether participants' trust in a graph was associated with how much they thought the graph looked beautiful. We used Linear Mixed Modeling (LMM) to regress participants' trust ratings (7-point) on their beauty ratings of the graphs (7-point) across all participants and the graphs they evaluated ($N = 2,681$ observations), while controlling for the random effects of individual participants (random intercepts). A positive and statistically significant coefficient of beauty ratings in this model would indicate that the more beautiful graphs were associated with higher trust.

We further controlled for the factors that might influence both participants' trust in the graphs and how much they thought the graphs looked beautiful. These factors included how much participants felt that they understood the graphs, and found the information interesting and surprising. We included these three covariates (7-point) in the above LMM model. A positive and statistically significant coefficient of beauty ratings in this model would indicate that the more beautiful graphs were associated with higher trust even when controlling for potential influencing factors.

II. Moderation effect of education

We tested whether participants' education level (5-point: less than high school credential; high school credential; some post-high-school, no bachelor's degree; bachelor's degree; graduate degree) would moderate the association between perceived beauty and trust in graphs. We used LMM to regress participants' trust ratings on their beauty ratings (z-scored), education (z-scored), and the interaction between beauty ratings and education, while controlling for the three covariates and the random effects of individual participants ($N = 2,681$ observations). A positive

and statistically significant coefficient of the interaction term between beauty ratings and education would indicate that the more educated people were the more likely that their trust in graphs was associated with how beautiful they perceived the graphs to be.

III. Association between perceived beauty and real-world popularity

We tested whether the real-world popularity of the graphs was associated with how beautiful our participants thought the graphs looked. We approximated the real-world popularity of the graphs using the number of comments that their posts received on Reddit. We used LMM to regress the beauty ratings (7-point) on the number of comments of the graphs (log-transformed, ranging from 1.79 to 8.58) across all participants and the graphs they evaluated (N = 2,681 observations), while controlling for the random effects of individual participants. A positive and statistically significant coefficient of the number of comments would indicate that the more beautiful graphs were associated with the posts that had greater real-world popularity on Reddit.

Study 1b

Participants

Participants were recruited from MTurk as in Study 1a. Participants were recruited based on the same five inclusion criteria as in Study 1a.

We planned to recruit 300 participants in total. This sample size was determined based on formal power analyses (targeting $\alpha = .05$ and 95% power) and participant exclusion contingencies (Table 1). We performed power analysis using data from Study 1a for our main hypothesis: that the perceived beauty of the graphs would be associated with the real-world popularity of the graphs (the number of comments on Reddit), taking into account the random effects of individual participants. Data from Study 1a showed that the effect size was 0.04 (using R package *lme4* and function *lmer* to regress the beauty ratings on the log-transformed number of comments and the

random intercepts of participants). A simulation power analysis showed that to detect an effect size of 0.04, we would need around 280 participants (using *R* package *simr* and functions *extend* and *powerCurve* to simulate the effect as a function of participants).

Assuming an exclusion rate of around 6%, we determined to recruit 300 participants in total (20 participants per module). We also planned that, if the actual exclusion rate turned out to be over 6% in a module (more than one participant), we would recruit participants until the final sample size after exclusion (see criteria below) reached 19 participants per module. This contingent plan ensured that the total sample size across the 15 modules (see Materials below) would not be fewer than 280 participants, as required by the power analysis.

We processed the data according to the following exclusion criteria. A trial was excluded if i) the response time for examining the graph was < 1s; ii) the response time for rating the graph on all six aspects was < 7s; iii) the total response time for examining and rating the graph was < 8s or > 200s; iv) all six aspects received the same rating. A participant was excluded if > 20% of their trials were excluded. According to these criteria, $n = 453$ out of 6,420 trials were excluded, and $n = 32$ out of $N = 321$ participants were excluded (actual exclusion rate > 6% in some modules). The final sample size was 289 participants, with every module having 19 or 20 participants as planned (see Table 1 for demographics).

Materials

Study 1b used 300 graphs from the Reddit community of Data Is Beautiful. These 300 graphs were randomly assigned to 15 modules, each containing 20 graphs. We selected these 300 graphs from a much larger set based on 10 criteria (Fig. 1B). We i) first included the 807 most popular posts in Data Is Beautiful from Jan-Feb, 2021 by applying the filters “Top” and “This Month”. From these posts, we excluded posts ii) that did not include valid media components, which

1 indicated that no graphs were in these posts; iii) whose media components were not static
2 images, since we only focused on static graphs; iv) that were not created originally by Reddit
3 users; and v) for posts created by the same author, included only the most popular post (highest
4 “upvote”). We extracted the graphs from the remaining posts, and excluded graphs vi) that were
5 too complex, such as long graphs that expanded multiple pages; vii) that were not in English;
6 viii) whose text labels were not clear, such as displaying excessive text in illegible font sizes; ix)
7 that was a photo instead of data visualization; and x) for graphs depicting the exact same content
8 (e.g., the relationship between film length and average IMDB rating), included only the most
9 popular one (highest “upvote”). These processes selected the final set of 300 graphs.

10 Procedures

11 Participants were randomly assigned to one of the fifteen modules. The procedures were the
12 same as in Study 1a, except that participants evaluated the graphs on six instead of five aspects:
13 i) beauty, ii) trust, iii) interest, iv) understanding, v) surprise, and vi) valence. The valence
14 question asked, “How unpleasant is the information? That is, how much does the graph describe
15 something you find negative, or bad news?” We added the question about valence for its
16 potential impacts on both perceived beauty and trust in graphs.

17 Statistical analyses

18 *I. Association between perceived beauty and trust*

19 We tested whether participants’ trust in a graph was associated with how much they thought the
20 graph looked beautiful, using the same LMM model as in Study 1a ($N = 5,780$ observations).

21 We further controlled for four covariates that might influence both participants’ trust in the
22 graphs and how much they thought the graphs looked beautiful as in Study 1a. These four

covariates were how much participants felt that they understood the graphs, and found the information interesting, surprising, and unpleasant.

II. Moderation effect of education

We tested whether participants' education level would moderate the association between perceived beauty and trust in graphs, using a similar LMM model as in Study 1a but with four instead of three covariates. That is, we regressed participants' trust ratings on their beauty ratings (z-scored), education (z-scored), the interaction between beauty ratings and education, while controlling for the four covariates and the random effects of individual participants (N = 5,780 observations).

III. Association between perceived beauty and real-world popularity

We tested whether the real-world popularity of the graphs was associated with how beautiful our participants thought the graphs looked, using the same LMM model as in Study 1a (N = 5,780 observations; the log-transformed number of comments ranged from 0.69 to 8.16).

Study 2

Participants

Participants were recruited from MTurk as in Study 1a and Study 1b. Participants were recruited based on the same five inclusion criteria as in Study 1a and Study 1b.

We planned to recruit 300 participants in total. This sample size was determined based on formal power analyses (targeting $\alpha = .05$ and 95% power) and participant exclusion contingencies (Table 1). We performed a power analysis for our main hypothesis: that the perceived beauty of the graphs would be associated with the real-world popularity of the graphs across four popular social media platforms (Facebook, Twitter, Pinterest, and Reddit). We expect that this effect

would be similar to that in Study 1a. The same power analysis procedure as in Study 1b indicated that we would need around 280 participants.

Assuming an exclusion rate of around 6%, we determined to recruit 300 participants in total (60 participants per module). We also planned that, if the actual exclusion rate turned out to be over 6% in a module (more than three participants), we would recruit participants until the final sample size after exclusion (see criteria below) reached 56 participants per module. This contingent plan ensured that the total sample size across the 5 modules (see Materials below) would not be fewer than 280 participants, as required by the power analysis.

We processed the data according to the following exclusion criteria. A trial was excluded if i) the response time for examining the graph was < 1 s; ii) the response time for rating the graph on all six aspects was < 7 s; iii) all six aspects received the same rating. A participant was excluded if $> 20\%$ of their trials were excluded. According to these criteria, $n = 624$ out of 7,216 trials were excluded, and $n = 46$ out of $N = 328$ participants were excluded (actual exclusion rate $> 6\%$ in some modules). The final sample size was 282 participants, with every module having 56 or 57 participants as planned (see Table 1 for demographics).

Materials

Study 2 used 110 graphs from 11 online news sources. These 110 graphs were randomly assigned to 5 modules, each containing 22 graphs. We selected these 11 online news sources based on six criteria (Fig. 1C, left). We included sources that i) were relatively popular based on the number of circulation or subscriptions; ii) were relatively factual based on the study of media bias by a media watchdog organization (<https://www.adfontesmedia.com/>); iii) frequently showed graphs in their news reports; iv) provided free access to news reports that were several months old; v) provided options for readers to share the news reports on various social media, so that we could

1 assess the real-world popularity of these reports; and vi) whose real-world popularity data across
2 multiple social media platforms were available via a social enterprise for online content research
3 (<https://buzzsumo.com/>). Based on these six criteria, we targeted 11 news sources: Bloomberg,
4 FiveThirtyEight, NPR News, The Economist, Financial Times, The Guardian, The New York
5 Times, The Wall Street Journal, The Washington Post, USA Today, and Vox.

6 We selected 10 news reports from each of these 11 news sources. News reports were sampled
7 from the most recent to the older ones based on four criteria (Fig. 1C, right). We targeted news
8 reports that i) had a graph (if multiple, randomly selected one graph); ii) were from the past one
9 to six months (Feb-June, 2021), so that the real-world popularity data had meaningfully accrued
10 and the contents were still recent enough to be relevant; iii) covered one of the five topics that
11 were overlapping across the 11 news sources (economy and business, health, politics, science
12 and technology, sports); and iv) were in one of the six types that were most common based on
13 the Reddit community Data Is Beautiful (bar plot, line plot, map, scatter plot, heatmap, and pie
14 chart). These processes selected the final set of 110 graphs (Fig. 2B). Source information was
15 removed from all graphs to avoid the effects of participants' heterogeneous preferences for
16 different news sources.

17 Procedures

18 Participants were randomly assigned to one of the five modules. The procedures were the same
19 as in Study 1b. After completing all trials, participants completed a brief survey about their
20 demographics, news media engagement (5-point Likert scales on trust in news, and news
21 following frequency), political affiliation (5-point Republican-Democrat Likert scale), and
22 personality traits (Ten-Item Personality Inventory³⁷). We added questions beyond demographics

1 in the survey to understand those factors' potential moderation effects on the relationship
2 between beauty and trust in graphs.

3 Statistical analyses

4 *I. Association between perceived beauty and trust*

5 We tested whether participants' trust in a graph was associated with how much they thought the
6 graph looked beautiful, using the same LMM model as in Study 1 (N = 6,204 observations).

7 We further controlled for the four covariates that might influence both participants' trust in the
8 graphs and how much they thought the graphs looked beautiful as in Study 1b.

9 *II. Moderation effect of education, political affiliation, and personality*

10 We tested whether participants' education level would moderate the association between
11 perceived beauty and trust in graphs, using the same LMM model as in Study 1b (N = 6,204
12 observations).

13 We also tested whether participants' political affiliation (5-point) and personality traits (in
14 particular, agreeableness and conscientiousness; 7-point) would moderate the association
15 between perceived beauty and trust in graphs. We used the same LMM model as in testing the
16 moderation effect of education, but replacing education with these three different moderators one
17 at a time (z-scored).

18 *III. Association between perceived beauty and real-world popularity*

19 We tested whether the real-world popularity of the graphs was associated with how beautiful our
20 participants thought the graphs looked. We approximated the graphs' real-world popularity using
21 the counts of comments, shares, and likes of the news reports on multiple social media platforms.
22 These counts across platforms were highly correlated (mean $r = 0.72$). As pre-registered, we first

performed a principal component analysis (PCA) on these different counts (log-transformed, z-scored), and approximated the graphs' real-world popularity using the first PC (with varimax rotation). We used LMM to regress the beauty ratings on the scores of the first PC across all participants and the graphs they evaluated ($N = 6,204$ observations), while controlling for the random effects (random intercepts) of individual participants and news sources. A positive and statistically significant coefficient of the PC scores would indicate that the more beautiful graphs were associated with the news reports that had greater real-world popularity across multiple social media platforms.

Study 3

Participants

Participants were recruited from MTurk as in Studies 1a, 1b, and 2. Different from Studies 1a, 1b, and 2, here we required participants to be native English speakers and have at least Bachelor's degrees. We used these two criteria because participants in the present study viewed graphs from academic papers. We expected that these graphs were more difficult to comprehend than those from Reddit or the news. Expecting this difficulty, we only included papers whose title and abstract (which were also shown to participants to provide background information) had no acronym and were understood by people with college education (see Materials below).

We planned to recruit 300 participants in total. This sample size was determined based on formal power analyses (targeting $\alpha = .05$ and 95% power) and participant exclusion contingencies (Table 1). We performed power analysis for our main hypothesis: that the perceived beauty of the graphs would be associated with the real-world popularity of the papers (e.g., the number of citations). We expect that this effect would be similar to that in Study 1a. The same power analysis procedure as in Studies 1b and 2 indicated that we would need around 280 participants.

Assuming an exclusion rate of around 6%, we determined to recruit 300 participants in total (20 participants per module). We also planned that, if the actual exclusion rate turned out to be over 6% in a module (more than one participant), we would recruit participants until the final sample size after exclusion (see criteria below) reached 19 participants per module. This contingent plan ensured that the total sample size across the 15 modules (see Materials below) would not be fewer than 280 participants, as required by the power analysis.

We processed the data according to the following exclusion criteria. A trial was excluded if i) the response time for examining the title and abstract of the paper was < 1 s; ii) the response time for examining a graph was < 1 s; iii) the response time for rating the graph on all six aspects was < 7 s; iv) all six aspects received the same rating. A participant was excluded if $> 20\%$ of their trials (i.e., graphs) were excluded. According to these criteria, $n = 524$ out of 6,902 trials were excluded, and $n = 42$ out of $N = 334$ participants were excluded (actual exclusion rate $> 6\%$ in some modules). The final sample size was 292 participants, with every module having 19, 20, or 22 (in two modules) participants as planned (see Table 1 for demographics).

Materials

Study 3 used 310 graphs from 155 scientific publications (2 graphs from each paper). Graphs from these 155 papers were randomly assigned to 15 modules, with 10 modules containing 20 graphs from 10 papers and 5 modules containing 22 graphs from 11 papers. We selected these 155 papers from two journals (Fig. 1D). These two journals were selected based on four criteria. We targeted journals that were i) domain general, so that we could include graphs of diverse subject areas (e.g., physical, biological, health sciences); ii) open access, so that we could freely download their information and show them to participants; iii) used similar display formats, to minimize potential biases from format differences between journals (e.g., whether or not the

figure and paper titles were displayed in the image file); and iv) substantially differed in their impact factors, so that we could compare graphs from papers with a wide range of real-world popularity. Based on these four criteria, we selected two journals: Scientific Reports, and Nature Communications.

From these two journals, we targeted papers that were published within a single year to minimize the effect of publication time differences on the real-world popularity metrics (e.g., the number of citations). We targeted the year of 2016, because it was old enough that the real-world popularity metrics had meaningfully accrued, and recent enough that both journals had started tagging papers for their subject areas. In the year of 2016, $N = 20,540$ papers were published in Scientific Reports, and $N = 3,526$ papers were published in Nature Communications.

From those papers, we selected a smaller subset for each journal based on seven criteria. We excluded papers that i) were without subject-area tags, since we aim to sample papers across a wide range of subject areas; ii) whose subject areas did not overlap with the common subject areas between the two journals; iii) whose subject areas had fewer than ten papers; and iv) had fewer than two graphs, since we aim to sample two graphs from each paper. We used automated web scraping to extract a paper's title and abstract, which were shown to participants to help them understand the graphs. We excluded papers whose titles and abstracts v) were not properly extracted (15%); vi) had any acronym (a word with more than two capital letters); and vii) were not understood by people below college education (readability > 16 , quantified with *R* package *qdap* and function *automated_readability_index*). These procedures retained $N = 1,028$ papers from Scientific Reports, and $N = 222$ papers from Nature Communications.

For the remaining papers, we identified the common subject areas between the two journals that had more than five papers ($n = 22$ subject areas; see Fig. 2C). We randomly selected five papers

from each of these subject areas for each journal by i) first randomly shuffling the 22 subject areas, and then ii) sequentially selecting 5 papers for each subject area randomly. The order of the subject areas made a difference to the paper selection because most papers had multiple subject-area tags. Selecting one paper would not only up the number of the papers for the current subject area, but potentially other (subsequent) subject areas. No more paper was selected for a subject area if there were already five or more papers associated with it when its turn came. These processes selected $N = 77$ papers from Scientific Reports and $N = 78$ papers from Nature Communications. We randomly selected two graphs from each paper, resulting in the final set of 310 graphs in total. The papers and their graphs that were randomly assigned to each module were balanced between the two journals (i.e., an equal number of papers from each journal).

Procedures

Participants were randomly assigned to one of the fifteen modules. In each module, participants viewed a subset of papers and their graphs. Participants evaluated the graphs on six aspects as in Studies 1b and 2: i) beauty, ii) trust, iii) interest, iv) understanding, v) surprise, and vi) valence. The order of the papers and the order of the two graphs from each paper were randomized for each participant.

In each trial, participants first viewed the title and abstract of the paper. Participants clicked a button when they finished reading those materials. Participants then viewed one of the graphs and its legend from that paper. Participants clicked a button when they finished examining the graph. Then the six questions appeared in random order below the graph (Fig. 3A). Participants indicated their answer for each question by clicking the radio buttons on a 7-point Likert scale as in Studies 1-2. Participants clicked a button to confirm when they finished answering all questions for this graph. Participants then viewed the other graph (and its legend) from this same

paper. Participants repeated the previous procedures and answered the six questions (in random order) for this second graph. After completing all questions for both graphs from the paper, participants moved on to the next paper. After evaluating all graphs from all papers in the module, participants filled out a brief survey about their demographics.

Statistical analyses

I. Association between perceived beauty and trust

We tested whether participants' trust in a graph was associated with how much they thought the graph looked beautiful, using the same LMM model as in Studies 1-2 with an additional random effect: the graphs' associated papers (N = 6,030 observations).

We further controlled for the four covariates that might influence both participants' trust in the graphs and how much they thought the graphs looked beautiful as in Studies 1b and 2.

II. Association between perceived beauty and real-world popularity

We tested whether the real-world popularity of the scientific papers that used the graphs was associated with how beautiful our participants thought the graphs looked. We approximated the real-world popularity of the papers using two different measures: i) the number of citations, and ii) the number of views the papers received that were provided on the journal website. We averaged the ratings on each aspect (beauty, trust, interest, understanding, surprise, valence) across the two graphs from each paper for each participant. We fit two LMM models, one for each type of the popularity measures. In each model, we regressed the average beauty of the graphs per paper on the paper's popularity (log-transformed) across all participants and the papers they evaluated (N = 3,015 observations), while controlling for the random effects of individual participants and the papers' subject areas. Since most papers were associated with multiple subject areas, we used a random selection procedure to determine the subject area for

each paper in the LMM model. Specifically, we fit the LMM model for 1,000 iterations; in each iteration, we randomly selected one subject area for each paper from their associated subject areas; finally, we averaged the model estimates (arithmetic mean of coefficients, harmonic mean of p-values) across all iterations. A positive and statistically significant coefficient of popularity would indicate that the more beautiful graphs were associated with papers that had greater real-world popularity.

We further controlled for the factors that might influence both the papers' real-world popularity and how much participants thought the graphs looked beautiful. These factors included the number of days since the paper was published, and participants' ratings on how interesting, understandable, surprising, and unpleasant the information was. We added these covariates to the above LMM model for each popularity measure, and repeated the analysis procedures.

Study 4

Statistical analyses

I. Visual features predict graph beauty

We used computational models provided in Open Source Computer Vision Library (OpenCV)²⁵ and Google's Tesseract-OCR Engine (Python-tesseract)²⁶ to characterize visual features of the graphs, including their resolution, word counts, and the distribution of their hue, saturation, luminance, and contrast. We used linear regression to regress perceived beauty on the summary statistics of these features across all real-world graphs from Studies 1-3 ($N = 859$).

II. Unique explained variance of visual features

To understand the unique contribution of each type of features, we performed variance partition analysis. The features for hue, saturation, luminance, and contrast included both their summary statistics and their frequency distribution across intensity values. For each type of features, we fit

three linear regression models to all real-world graphs from Studies 1-3 ($N = 859$): (i) one regressing beauty on the given type of features, (ii) the second regressing beauty on the rest of features, and (iii) the third regressing beauty on all features. The variance explained by each model was assessed with adjusted R^2 . The unique variance explained by a given type of features equals the explained variance of model (iii) minus that of model (ii).

III. Univariate correlations between visual features and graph beauty

For the top three predictive features of beauty (hue, luminance, and saturation), we further investigated how the frequency at each intensity value of these features predicted beauty. For each feature type (e.g., luminance) and each intensity value of that feature (e.g., luminance ranged from 0 to 255), we computed the univariate Pearson correlation between beauty and the frequency of that intensity value across all real-world graphs from Studies 1-3 ($N = 859$). To visualize what colors – as specified by the combinations of hue, luminance, and saturation values – were perceived as the most and least beautiful, we averaged the univariate correlations (Fisher's z-transformed) across hue, luminance, and saturation for each combination of their intensity values. The 10,000 most positive and most negative average univariate correlations were respectively projected onto a two-dimensional space via uniform manifold approximation and projection (UMAP) for visualization.

Study 5a

Participants

Participants were recruited from MTurk as in Studies 1a, 1b, and 2. The only difference was that here we did not set any requirement for participants' education. We intended to include participants of all education levels in the present study to further investigate how education may moderate the causal effect of beauty on trust in the general public.

We planned to recruit 500 participants in total. This sample size was determined based on formal power analyses (targeting $\alpha = .05$ and 95% power) and participant exclusion contingencies (Table 1). We performed a power analysis based on our main hypothesis: that individual participants would trust the more beautiful graphs more than the less beautiful graphs. We targeted an effect size of $d = 0.1807$, a small effect size in pre-registered publications using within-subject designs according to a recent survey on the empirical effect sizes in psychological research³⁸. A two-sided paired t-test power analysis indicated that to detect such an effect size, we would need at least 400 participants (using *R* package *pwr* and function *pwr.t.test*).

Assuming an exclusion rate of around 20%, we determined to recruit 500 participants in total. We also planned that, if the actual exclusion rate turned out to be over 20%, we would recruit participants until the final sample size after exclusion (see criteria below) reached 400 participants. This contingent plan ensured that our final sample size would not be fewer than what the power analysis required.

We processed the data according to the following exclusion criteria. A trial was excluded if the response time was < 1500 milliseconds. A participant was excluded if i) more than 3 out of their 12 trials were excluded; or ii) they gave the same rating to all graphs on an aspect (e.g., beauty). According to these criteria, $n = 667$ out of 3,006 trials were excluded, and $n = 72$ out of $N = 501$ participants were excluded. The final sample size was 429 participants, reaching the sample size required by the power analysis as planned (see Table 1 for demographics).

Materials

We used random data to generate a total of 144 graphs (6 types x 6 topics x 4 experiment conditions). These graphs were in six types: bar plot, line plot, map, scatter plot, heatmap, and pie chart. These were the six graph types that were most common in the Reddit community Data

1 Is Beautiful (Study 1) and the same as those from the news (Study 2). These graphs covered six
2 different topics that were popular in the news: politics (lobbying), sports (soccer), science and
3 technology (clean energy), economic and business (money laundering), health (stress), and
4 nature (climate risk).

5 We used the same set of random data to generate all graphs of a specific type, regardless of their
6 topics. For a graph of a specific type and a specific topic, we digitally manipulated it to derive
7 four versions (Fig. 3A): i) more beautiful and less misleading, ii) less beautiful and less
8 misleading, iii) more beautiful and more misleading, and iv) less beautiful and more misleading
9 (see results on misleadingness manipulation in the Supplementary Information).

10 The manipulation of graph beauty was based on previous research on the aesthetics of data
11 visualization. Specifically, we manipulated the graph beauty in four aspects: image resolution,
12 font size, font type, and color saturation. The more beautiful versions of the graphs were plotted
13 with high image resolution, legible font size, Sans font type, and saturated color. The less
14 beautiful version of the graphs were plotted with low image resolution, smaller font size, Comic
15 Sans MS font type, and desaturated color (grayscale for line plots and scatter plots; 50%
16 saturation for other graph types).

17 Procedures

18 Participants were told that this was a study about how people perceived graphs. They were
19 instructed that the graphs used in the study were generated by the researchers themselves, and
20 that those graphs might or might not be accurate. Each participant viewed and evaluated six
21 graphs in this study. These six graphs were randomly selected for each participant with three
22 constraints: i) they were in six different types; ii) they were of six different topics; and that iii)

they sampled all four experiment conditions, with each experiment condition repeating no more than twice.

Participants completed two sessions of the study. In one session, participants evaluated how beautiful each graph looked. In the other session, participants reported how much they trusted each graph to be accurate. The order of the two sessions was randomized across participants. In each session, participants viewed the six graphs one by one in random order. For each graph, participants first saw the corresponding question of the session (e.g., “Do you trust this graph?”) and the graph shown below the question. After a brief time lag (1s), participants saw a 7-point Likert scale below the graph (e.g., anchored at 1 = Do Not Trust, 4 = Neutral, 7 = Trust). Participants could enter their answers as soon as the Likert scale appeared. Participants indicated their answers by pressing the number keys 1, 2, 3, 4, 5, 6, or 7 on the top row of their computer keyboard. There was no time limit for entering an answer. After evaluating all six graphs for the current session, participants moved on to the second session. After completing both sessions, participants filled out a brief survey about their demographics.

Statistical analyses

I. Manipulation check of graph beauty

We manipulated a graph in the same type on the same topic to be more and less beautiful and misleading. We measured the effect of these manipulations on how much participants thought the graph was beautiful and how much they trusted it. We first checked whether our manipulation of graph beauty was successful. We used LMM to regress the beauty ratings (7-point) on the beauty manipulation (binary) and misleadingness manipulation (binary) across all participants and the graphs they evaluated ($N = 2,574$ observations), while controlling for the random effects (random intercepts) of individual participants, graph types, and graph topics. A

positive and statistically significant coefficient of beauty manipulation in this model would indicate that our beauty manipulation was successful.

II. Causal effects of beauty manipulation

We tested whether making a graph look more beautiful would increase people's trust in it. We used LMM to regress participants' trust ratings (7-point) on the beauty manipulation of the graphs, while controlling for the misleadingness manipulation of the graphs and random effects of individual participants, graph types, and graph topics (N = 2,574 observations). A positive and statistically significant coefficient of beauty manipulation in this model would indicate that making a graph more beautiful caused people to trust it more.

III. Moderation effect of education

We tested whether participants' education level would moderate the causal effect of beauty manipulation on trust in graphs. We used LMM to regress participants' trust ratings on the beauty manipulation, the misleadingness manipulation, participants' education, the interaction between beauty manipulation and participants' education, and the interaction between misleadingness manipulation and participants' education, while controlling for the random effects of individual participants, graph types, and graph topics (N = 2,574 observations). A positive and statistically significant coefficient of the interaction term between beauty manipulation and education would indicate that the more educated people were, the more likely their trust in graphs would be biased by the beauty of the graphs.

Study 5b

Participants

Participants were recruited from MTurk as in Studies 1a, 1b, and 2. We followed the same sample size plan and exclusion criteria as in Study 5a. Two-hundred and seventy-five out of

2,424 trials were excluded, and $n = 97$ out of $N = 501$ participants were excluded. The final sample size was 404 participants, reaching the sample size required by the power analysis as planned (see Table 1 for demographics).

Materials

The same graphs as in Study 5a were used, except that a) only two experimental conditions were included: more and less beautiful (a total of 72 graphs); and b) that the beauty manipulation did not rely on the font.

Procedures

The same experimental procedures as in Study 5a were used, except that a) the six graphs each participant viewed were balanced sampled across both experimental conditions (three graphs in each condition), and that b) participants completed three sessions: besides reporting how beautiful each graph looked and how much they trusted each graph, participants also reported how competent they inferred the graph maker to do things in general (the order of the three sessions was randomized).

Statistical analyses

I. Manipulation check of graph beauty

We used LMM to regress the beauty ratings (7-point) on the beauty manipulation (binary) across all participants and graphs ($N = 2,424$ observations), while controlling for the random effects (random intercepts) of individual participants, graph types, and graph topics.

II. Causal effects of beauty manipulations

We tested whether making a graph look more beautiful without changing the font would still increase people's trust in it. We used LMM to regress participants' trust ratings (7-point) on the

1 beauty manipulation of the graphs, while controlling for the random effects of individual
2 participants, graph types, and graph topics (N = 2,424 observations).

3 *III. Moderation effect of education*

4 We tested whether participants' education level would moderate the causal effect of beauty
5 manipulation on trust in graphs. We used LMM to regress participants' trust ratings on the
6 beauty manipulation, participants' education, the interaction between beauty manipulation and
7 participants' education, while controlling for the random effects of individual participants, graph
8 types, and graph topics (N = 2,424 observations).

9 **References**

- 10 1. Franconeri, S. L., Padilla, L. M., Shah, P., Zacks, J. M. & Hullman, J. The Science of Visual
11 Data Communication: What Works. *Psychol Sci Public Interest* **22**, 110–161 (2021).
- 12 2. Hehman, E. & Xie, S. Y. Doing Better Data Visualization. *Advances in Methods and*
13 *Practices in Psychological Science* **4**, 1–18 (2021).
- 14 3. Dragicevic, P. & Jansen, Y. Blinded with Science or Informed by Charts? A Replication
15 Study. *IEEE Transactions on Visualization and Computer Graphics* **24**, 781–790 (2018).
- 16 4. Allen, E. A., Erhardt, E. B. & Calhoun, V. D. Data Visualization in the Neurosciences:
17 Overcoming the Curse of Dimensionality. *Neuron* **74**, 603–608 (2012).
- 18 5. Crameri, F., Shephard, G. E. & Heron, P. J. The misuse of colour in science communication.
19 *Nat Commun* **11**, 5444 (2020).
- 20 6. Jones, G. *How to lie with charts*. (LaPuerta, 2006).
- 21 7. Lauer, C. & O'Brien, S. How People Are Influenced by Deceptive Tactics in Everyday
22 Charts and Graphs. *IEEE Transactions on Professional Communication* **63**, 327–340 (2020).

8. Pandey, A. V., Rall, K., Satterthwaite, M. L., Nov, O. & Bertini, E. How Deceptive are Deceptive Visualizations?: An Empirical Analysis of Common Distortion Techniques. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* 1469–1478 (2015) doi:10.1145/2702123.2702608.
9. Yang, B. W., Vargas Restrepo, C., Stanley, M. L. & Marsh, E. J. Truncating Bar Graphs Persistently Misleads Viewers. *Journal of Applied Research in Memory and Cognition* **10**, 298–311 (2021).
10. Geissbuehler, M. & Lasser, T. How to display data by color schemes compatible with red-green color perception deficiencies. *Opt. Express, OE* **21**, 9862–9874 (2013).
11. Tareen, A. & Kinney, J. B. Logomaker: beautiful sequence logos in Python. *Bioinformatics* **36**, 2272–2274 (2020).
12. Barter, R. L. & Yu, B. Superheat: An R package for creating beautiful and extendable heatmaps for visualizing complex data. *J Comput Graph Stat* **27**, 910–922 (2018).
13. Theys, K., Lemey, P., Vandamme, A.-M. & Baele, G. Advances in Visualization Tools for Phylogenomic and Phylodynamic Studies of Viral Diseases. *Frontiers in Public Health* **7**, 208 (2019).
14. Lau, A. & Vande Moere, A. Towards a Model of Information Aesthetics in Information Visualization. *2007 11th International Conference Information Visualization (IV '07)* 87–92 (2007) doi:10.1109/IV.2007.114.
15. Judelman, G. Aesthetics and inspiration for visualization design: bridging the gap between art and science. *Proceedings. Eighth International Conference on Information Visualisation, 2004. IV 2004.* 245–250 (2004) doi:10.1109/IV.2004.1320152.

16. Harrison, L., Reinecke, K. & Chang, R. Infographic Aesthetics: Designing for the First Impression. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* 1187–1190 (2015) doi:10.1145/2702123.2702545.
17. Ertug, G., Gruber, M., Nyberg, A. & Steensma, H. K. From the Editors—A Brief Primer on Data Visualization Opportunities in Management Research. *AMJ* **61**, 1613–1625 (2018).
18. Seifirad, S. & Haghpanah, V. Letter to the editor: Clothes Do Not Make the Man: Well-favored Figures are Game-changers in the Biomedical Publication. *Journal of Korean Medical Science* **30**, 1713–1713 (2015).
19. Mayr, E., Hynek, N., Salisu, S. & Windhager, F. Trust in Information Visualization. *TrustVis@EuroVis* 25–29 (2019).
20. Dion, K., Berscheid, E. & Walster, E. What is beautiful is good. *Journal of personality and social psychology* **24**, 285 (1972).
21. Bascandziev, I. & Harris, P. L. In beauty we trust: Children prefer information from more attractive informants. *British Journal of Developmental Psychology* **32**, 94–99 (2014).
22. Darby, B. W. & Jeffers, D. THE EFFECTS OF DEFENDANT AND JUROR ATTRACTIVENESS ON SIMULATED COURTROOM TRIAL DECISIONS. *soc behav pers* **16**, 39–50 (1988).
23. Moore, F. R., Filippou, D. & Perrett, D. I. Intelligence and attractiveness in the face: Beyond the attractiveness halo effect. *Journal of Evolutionary Psychology* **9**, 205–217 (2011).
24. Zebrowitz, L. A. & Montepare, J. M. Social Psychological Face Perception: Why Appearance Matters. *Soc Personal Psychol Compass* **2**, 1497 (2008).
25. Bradski, G. The OpenCV Library. *Dr. Dobbs's Journal of Software Tools* (2000).
26. Hoffstaetter, S. pytesseract: Python-tesseract is a python wrapper for Google's Tesseract-OCR.

- 1 27. Jones, E. & Davis, K. From Acts To Dispositions The Attribution Process In Person
2 Perception - ScienceDirect. *Advances in Experimental Social Psychology* **2**, 219–266 (1965).
- 3 28. Zhao, N., Zhou, M., Shi, Y. & Zhang, J. Face Attractiveness in Building Trust: Evidence
4 from Measurement of Implicit and Explicit Responses. *soc behav pers* **43**, 855–866 (2015).
- 5 29. Sundar, A., Cao, E. S. & Machleit, K. A. How product aesthetics cues efficacy beliefs of
6 product performance. *Psychology & Marketing* **37**, 1246–1262 (2020).
- 7 30. Li, Y.-M. & Yeh, Y.-S. Increasing trust in mobile commerce through design aesthetics.
8 *Computers in Human Behavior* **26**, 673–684 (2010).
- 9 31. Crolic, C., Zheng, Y., Hoegg, J. & Alba, J. W. The Influence of Product Aesthetics on
10 Consumer Inference Making. *Journal of the Association for Consumer Research* **4**, 398–408
11 (2019).
- 12 32. Franconeri, S. L. Three Perceptual Tools for Seeing and Understanding Visualized Data.
13 *Curr Dir Psychol Sci* **30**, 367–375 (2021).
- 14 33. Börner, K., Bueckle, A. & Ginda, M. Data visualization literacy: Definitions, conceptual
15 frameworks, exercises, and assessments. *Proc Natl Acad Sci USA* **116**, 1857–1864 (2019).
- 16 34. Jolly, E. & Chang, L. J. The Flatland Fallacy: Moving Beyond Low–Dimensional Thinking.
17 *Topics in Cognitive Science* **11**, 433–454 (2019).
- 18 35. Nastase, S. A., Goldstein, A. & Hasson, U. Keep It Real: Rethinking the Primacy of
19 Experimental Control in Cognitive Neuroscience. *NeuroImage* **222**, 117254 (2020).
- 20 36. Litman, L., Robinson, J. & Abberbock, T. TurkPrime.com: A Versatile Crowdsourcing Data
21 Acquisition Platform for the Behavioral Sciences. *Behav Res* **49**, 433–442 (2017).
- 22 37. Gosling, S. D., Rentfrow, P. J. & Swann, W. B. A very brief measure of the Big-Five
23 personality domains. *Journal of Research in Personality* **37**, 504–528 (2003).

1 38. Schäfer, T. & Schwarz, M. A. The Meaningfulness of Effect Sizes in Psychological
2 Research: Differences Between Sub-Disciplines and the Impact of Potential Biases. *Front.*
3 *Psychol.* **10**, (2019).

Acknowledgments: We thank Sasha Brietzke, Tessa Rusch, Lindsey Tepfer, Landry Bulls, and Caroline Conway for their assistance.

Author contributions:

Conceptualization: CL, MT

Study Design: CL, MT

Material Preparation: CL

Data Collection: CL

Data Analysis: CL

Supervision: MT

Funding acquisition: MT

Writing – original draft: CL

Writing – review & editing: CL, MT

Competing interests: Authors declare that they have no competing interests.

Data and materials availability: All data, code, and materials used in the analysis are available on Open Science Framework:

https://osf.io/yutgx/?view_only=b679d0841fd440eb901a28cecf5908b9.

Supplementary Information for

Visualization aesthetics bias trust in science, news, and social media

Chujun Lin*, Mark Thornton

Correspondence to: Chujun.Lin@Dartmouth.edu

Fig. S1. Coefficients of beauty and four covariates in LMM across Studies 1-3. Each column plots the results from one study. For each study, we regressed participants' trust in the graphs on participants' ratings regarding five aspects of the graphs (panels; negativity was not measured in Study 1a), using linear mixed effect modeling. Individual participants were treated as random effects. Dots indicate the coefficients and whiskers indicate the standard errors. The number of observations were N = 2,681 in Study 1a; N = 5,780 in Study 1b; N = 6,204 in Study 2; N = 6,030 in Study 3).

Fig. S2. Examples of the least and most beautiful graphs. (A) The least beautiful graph in Study 1a based on participants' ratings (rated 2.65 on a scale from 1 to 7, averaged across the participants who rated this graph, $n = 20$). (B-C) The least beautiful graphs in Study 1b (both rated 2.00). (D) The least beautiful graph in Study 2 (rated 3.35). (E) The least beautiful graph in Study 3 (rated 2.63). (F-G) The most beautiful graphs in Study 1a (rated 5.55 and 5.58). (H-I) The most beautiful graphs in Study 1b (both rated 6.00). (J) The most beautiful graph in Study 2 (rated 5.67). (K) The most beautiful graph in Study 3 (rated 6.11).

Supplementary Text

I. Misleadingness manipulation in Study 5a

Besides the manipulation of graph beauty, we also added the manipulation of misleadingness of the graphs in Study 5a. We added the misleadingness manipulation to investigate whether beauty would still have an effect on trust if the graphs are misleading and to contextualize the causal effect size of beauty in comparison to misleadingness. The manipulation of graph misleadingness was based on previous research on best practices for data visualization ¹⁻⁶. Specifically, the more misleading version of the bar plot was truncated (starting from above zero) to exaggerate the relative group differences. The more misleading version of the line plot connected the available data points in a time series using piecewise polynomial interpolation (spline interpolation) instead of straight lines. The more misleading version of the map excluded New Zealand and Antarctica from the world map. The more misleading version of the scatter plot used uneven scales between the x- and y-axes to exaggerate the correlation in the data. The more misleading version of the pie chart omitted the baseline group to exaggerate the proportions of the other groups. The more misleading version of the heatmap used a non-uniform color scale (Jet) instead of a uniform color scale (Viridis).

We compared the causal effects of beauty manipulation versus misleadingness manipulation on the trust in graphs. We used bootstrap resampling to test whether the coefficient of beauty manipulation in the LMM model in Study 5a was significantly greater than the coefficient of misleadingness manipulation. In each resampling iteration, we randomly sampled participants with replacement, and used this new set of participants' data to refit the LMM model. We computed the difference between the absolute value of the coefficient for beauty manipulation and the absolute value of the coefficient for misleadingness manipulation. Over 1,000 iterations,

we generated an empirical distribution of the differences between the two coefficients. We obtained the 95% confidence interval of this difference by taking the .025 and .975 quantiles of the empirical distribution. A positive lower bound of the 95% confidence interval would indicate that the causal effect of beauty manipulation was significantly greater than the causal effect of misleadingness manipulation.

We found that people were insensitive to the actual misleadingness of the graph, showing no significant change in trust due to this manipulation ($b = -0.07$, $\beta = -0.02$, $p = 0.233$). The magnitude of the causal effect of graph beauty on trust was significantly greater than the magnitude of the effect of graph misleadingness ($\Delta|b| = 0.43$, 95% CI = [0.27, 0.58]; $\Delta|\beta| = 0.14$, 95% CI = [0.09, 0.19]). These findings indicate that people's trust in a graph is more heavily influenced by the beauty of the graph than the objective misleadingness of the graph.

II. Supplementary studies

We conducted five supplementary studies (Studies S1-S5) to further characterize how the causal effect of graph beauty on trust might be moderated by five factors: the objective misleadingness of the graphs (Study S1), the familiarity of the topics of the graphs (Study S2), the between- versus within-subject design of the study (Study S3), the different components of beauty manipulation of the graphs (Study S4), and the statistical power of the study (Study S5) (see below for the details on each study).

We did not find significant effects of beauty on trust in four of the studies. This is likely because those studies are underpowered, as they used a between-subject design and only one graph type of one topic (Studies S1-S2), or/and much smaller samples in terms of the number of observations (Studies S1-S4) than the other studies where the significant causal effect of beauty on trust was detected (Study 5a, Study 5b, and Study S5).

This possibility was supported by our mega-analysis (Fig. 6C). We pooled all data across the five supplementary studies (N = 4,386 observations) and performed a linear mixed effect analysis where we regressed participants' reported trust on beauty manipulation while controlling for the random effects of individuals, graph types, graph topics, and studies. We successfully detected a significant effect of beauty on trust in this larger sample ($b = 0.16$, $p = 2.35 \times 10^{-4}$), confirming that beauty does causally influence trust in graphs.

Table S1. Summary of study designs and causal effects of graph beauty on trust.

Study	Number of participants	Number of graphs	Number of graph topics	Number of graph types	Number of observations	Effect size of beauty on trust
Study 5a	429	144	6	6	2,574	0.50 ***
Study 5b	404	72	6	6	2,424	0.28 ***
Study S1	490	4	1	1	490	0.001
Study S2	139	4	1	1	139	0.135
Study S3	180	24	6	6	1,080	0.132
Study S4	143	40	5	5	715	0.305
Study S5	327	72	6	6	1,962	0.241 ***

Note: Effect size is the coefficient of beauty manipulation (b) in linear regressions of trust ratings. The random effects of individual participants, graph topics, and graph types are controlled for in linear mixed models when repeated measurements of these variables are present. Statistical significance: *** $p < 0.001$.

Participants in all supplementary studies provided informed consent in a manner approved by the Committee for Protection of Human Subjects of Dartmouth College. One of the supplementary studies, Study S1, was preregistered. The pre-registration of this study can be accessed on OSF: https://osf.io/ue34c/?view_only=05d4e595d83a4011a296af3544c43f00. All materials and data can be accessed on OSF: https://osf.io/yutgx/?view_only=b679d0841fd440eb901a28cecf5908b9.

Table 2S. Summary of sample size and demographics.

Study	Number of participants	Gender	Age	Race
Study S1	490	220 women, 261 men, 5 non-binary	Range from 20 to 78 M = 40, SD = 12	410 White, 27 Black, 34 Asian, 5 Others
Study S2	139	58 women, 80 men	Range from 22 to 71 M = 40, SD = 12	116 White, 5 Black, 14 Asian, 3 Others
Study S3	180	72 women, 106 men, 2 non-binary	Range from 20 to 76 M = 42, SD = 13	143 White, 19 Black, 13 Asian, 4 Others

Study S4	143	49 women, 92 men, 1 non-binary	Range from 19 to 78 M = 38, SD = 11	113 White, 14 Black, 13 Asian, 3 Others
Study S5	327	169 women, 158 men	Range from 18 to 78 M = 41, SD = 13	256 White, 30 Black, 29 Asian, 10 Others

Note: All survey questions were voluntary. Non-responses were omitted.

Supplementary Study S1

Materials

Study S1 intends to address the concern that the misleadingness manipulation we applied in Study 5a was in fact not misleading⁷. In this study, we manipulated the misleadingness of one graph (a bar plot) on one topic (how much different apps collect private data from users) using a less controversial misleading tactic: we changed the increments on the y-axis half way to exaggerate the difference between apps. We also manipulated the beauty of this graph using the same method as in Study 5b.

Procedures

Participants were randomly assigned one of the four graphs. Participants answered six questions about this graph in random order, including how beautiful the graph looked, how much they trusted the graph, understood the graphs, and found the information interesting, surprising, and unpleasant. Participants then answered two questions about how much they thought they would change their behavior based on the graph, one question about how much they thought the graph make was competent, one question about their current mood, and two questions about their interpretation of the graph. Participants entered their answers by pressing the number keys 1 to 7 on their computer keyboard. After completing all trials, participants filled out a brief survey about their demographics and experiences with data visualizations.

Analysis

We confirmed that our beauty manipulation was successful ($b = 1.74, p = 8.37 \times 10^{-33}; n = 490$ observations; this linear regression model also took into account the effect of misleadingness manipulation). The effect of beauty on trust was not significant ($b = 0.001, p = 0.994$; this linear regression model also took into account the effect of misleadingness manipulation and the interaction between beauty and misleadingness manipulations on trust).

Even though we did not observe a significant effect of beauty, we carried out all other pre-registered analyses as planned and report their results here. The causal effect of beauty on trust remained insignificant when controlling for the covariates of people's understanding of the graph, interest in the information, perceived surprisingness of the information, and perceived unpleasantness of the information ($b = -0.18, p = 0.222$). Besides trust, we also asked participants how much they would change their behavior based on the graph. We did not observe a significant effect of beauty on the behavior change tendency for either "start/continue using Facebook Messenger" ($b = 0.19, p = 0.295; b = 0.24, p = 0.187$ when controlling for the above four covariates) or "updating privacy setting on apps" ($b = 0.21, p = 0.207; b = 0.07, p = 0.675$ with covariates).

We did not find evidence that the effect of increasing graph beauty on increasing trust was greater than that of reducing graph misleadingness on increasing trust ($\Delta|b| = -0.01, 95\% \text{ CI} = [-0.28, 0.27]$, confidence interval estimated from 1,000 bootstrap resampling; $\Delta|b| = -0.01, 95\% \text{ CI} = [-0.29, 0.26]$ with covariates). We did not find evidence that the effect of graph beauty on behavioral change was greater than that of graph misleadingness on behavioral change either (for "start/continue using Facebook Messenger", $\Delta|b| = 0.06, 95\% \text{ CI} = [-0.36, 0.46]; \Delta|b| = 0.13, 95\% \text{ CI} = [-0.24, 0.49]$ with covariates; for "updating privacy setting on apps", $\Delta|b| = 0.06, 95\% \text{ CI} = [-0.36, 0.45]; \Delta|b| = -0.15, 95\% \text{ CI} = [-0.50, 0.26]$ with covariates).

We found no evidence that participants' inferred competence of the graph-maker statistically mediated the causal effect of beauty on trust ($b = 0.027, p = 0.700$). We found weak evidence that participants' current mood statistically mediated the causal effect of beauty on trust ($b = 0.028, p = 0.046$). However, since the total effect of beauty on trust was not significant in this study, we refrain from drawing conclusions about the mediation effects.

We found no evidence that general education level moderated the graph beauty bias ($\beta = -0.08, p = 0.381$; $\beta = -0.08, p = 0.360$ with covariates). We found no evidence that data-visualization-specific education moderated the causal effect of beauty on trust either, including how often participants saw graphs ($\beta = 0.06, p = 0.536$; $\beta = 0.07, p = 0.438$ with covariates), made graphs ($\beta = 0.07, p = 0.424$; $\beta = 0.05, p = 0.532$ with covariates), whether they received education on data visualizations ($\beta = 0.01, p = 0.968$; $\beta = -0.02, p = 0.925$ with covariates), read books or tutorials ($\beta = -0.08, p = 0.682$; $\beta = -0.03, p = 0.882$ with covariates), took classes ($\beta = -0.08, p = 0.694$; $\beta = -0.20, p = 0.253$ with covariates), received training ($\beta = 0.56, p = 0.024$; $\beta = 0.40, p = 0.077$ with covariates), and received related degrees ($\beta = -0.09, p = 0.889$; $\beta = 0.39, p = 0.533$ with covariates).

Finally, we found no evidence that beauty manipulation moderated the effect of misleadingness manipulation on the accuracy of graph interpretation ($\beta = 0.03, p = 0.639$; $\beta = 0.03, p = 0.619$ with covariates).

Supplementary Study S2

Materials

We conducted supplementary Study S2 to examine the sensitivity of the causal effect of beauty on trust to the topic of the graph. It is possible that people's trust in the graph vary much less with the beauty of the graph when they have a strong belief about the true statistics related to a

familiar topic (Study S1) versus when they have no prior about the trust statistics related to nonfamiliar topics (Study 5). We generated 14 topics that could possibly fit the description of the graph in Study S1 (e.g., the amount of money each app spent on user experience design in 2021; the number of users each app has around the globe; the percentage of traffic driven by each app daily across the internet, etc.). We asked participants over the internet how likely the graph (axis labels masked) described each of the 14 topics (N = 50; 14 women, 34 men, 1 non-binary; age range from 23 to 64, M = 38, SD = 11; 34 White, 10 Black, 2 Asian, 4 others). We selected the topic that was most ambiguous in terms of whether it could be described by the graph (with an average rating closest to the mid-point: the percentage of earnings of each app that came from advertising in 2021). We used the same graph in the four different conditions (more and less beautiful x more and less misleading) as in Study S1 but with this new topic/title. This study was otherwise administered identically to Study S1.

Analysis

We confirmed that our beauty manipulation was successful ($b = 2.54, p = 2.53 \times 10^{-20}; n = 139$ observations; this linear regression model also took into account the effect of misleadingness manipulation). The effect of beauty on trust was much larger after switching to an ambiguous topic in the present study than in Study S1 but it was still not significant ($b = 0.135, p = 0.719$; this linear regression model also took into account the effect of misleadingness manipulation and the interaction between beauty and misleadingness manipulations on trust). We did not find evidence that increasing misleadingness reduced trust ($b = -0.219, p = 0.572$), nor that the effect of increasing graph beauty on increasing trust was greater than that of reducing graph misleadingness on increasing trust ($\Delta|b| = -0.04, 95\% \text{ CI} = [-0.69, 0.62]$).

Supplementary Study S3

Materials

We conducted supplementary Study S3 to examine the sensitivity of the causal effect of beauty on trust to the between- versus within-subject design of the study. It is possible that people's trust in the graph vary much less with the beauty of the graph when they only see one graph in one beauty condition (Study S1) versus when they see multiple different graphs in two different beauty conditions (Study 5). We substituted the bar plot in Study 5b with the bar plot in Study S1 and subsampled only one topic for the rest of the graph types in Study 5b, resulting in a total of 24 graphs, with six different types, each type corresponding to one different topic, in four different experimental conditions (more and less beautiful x more and less misleading).

Procedures

Similar to the within-subject design in Studies 5a&b, each participant viewed and evaluated six graphs. These six graphs were randomly selected for each participant with two constraints: i) they were in six different types (and thus six different topics since one different topic was selected for each type); and that ii) they sampled all four experiment conditions, with each experiment condition repeating no more than twice. Participants completed two sessions of the study, rating the graphs on beauty and trust, as in Study 5a.

Analysis

We confirmed that our beauty manipulation was successful ($b = 1.40$, $p = 2.19 \times 10^{-54}$; $n = 1080$ observations; this linear regression model also took into account the effect of misleadingness manipulation, as well as the random effects of participants and graph types). The effect of beauty on trust was much larger after using a within-subject design in the present study than in Study S1 but it was still not significant ($b = 0.132$, $p = 0.116$; this linear regression model also took into account the effect of misleadingness manipulation and the random effects of participants and

graph types). We did not find evidence that increasing misleadingness reduced trust ($b = 0.019$, $p = 0.817$), nor that the effect of increasing graph beauty on increasing trust was greater than that of reducing graph misleadingness on increasing trust ($\Delta|b| = 0.07$, 95% CI = [-0.11, 0.25]).

Supplementary Study S4

Materials

We conducted supplementary Study S4 to examine the sensitivity of the causal effect of beauty on trust to the different components of beauty manipulation (in particular, the component of font change). It is possible that people's trust in the graph vary much less with the beauty of the graph when the font was kept consistent across beauty conditions (Study S1) versus when the font in the less beautiful version was less professional and less legible (Study 5a). We used the graphs from Study S3 except for the graph type of bar plot (the graph from Study S1), and in addition to the four experimental conditions in Study S3 (more and less beautiful x more and less misleading), we generated four more versions for each graph type where the beauty manipulation also included changes in the font type and font size as in Study 5a. That is, the present study used a total of 40 graphs, with five different graph types, each type of a different topic, in 8 different experimental conditions (more and less beautiful without font change x more and less misleading and more and less beautiful with font change x more and less misleading).

Procedures

Each participant viewed and evaluated five graphs. These five graphs were randomly selected for each participant with two constraints: i) they were in five different types (and thus five different topics since one different topic was selected for each type); and that ii) they sampled five of the eight different experiment conditions. Participants first completed a session where they rated the graph one by one on how much they trusted the graph, using a 7-point Likert scale. They then

completed a second session where they wrote down using their own words why they trusted each graph to the degree they reported in the previous session. Finally, participants completed a third session where they viewed the graphs one by one again and evaluated how beautiful the graph was, how much they understood the graph, how much they were interested in the information, and how surprising the information conveyed by the graph was. Participants entered their answers by pressing the number keys 1 to 7 on their computer keyboard. After completing all trials, participants filled out a brief survey about their demographics.

Analysis

We confirmed that our beauty manipulation was successful ($b = 1.03$, $p = 3.63 \times 10^{-18}$; $n = 715$ observations; this linear regression model also took into account the effect of text manipulation, misleadingness manipulation, as well as the random effects of participants and graph types). Text manipulation alone did not significantly influence perceived beauty ($b = 0.193$, $p = 0.096$). There was a negative trend in the interaction between beauty manipulation and text manipulation on trust, suggesting that removing the font change might reduce the effect of beauty on trust ($b = -0.209$, $p = 0.623$); however, since the effect was not significant, we refrain from drawing strong conclusions about the effect of font change (this linear regression model also took into account the effect of misleadingness manipulation and the random effects of participants and graph types). The effect of beauty on trust was positive and of a decent size as expected but it was still not significant ($b = 0.305$, $p = 0.311$). We did not find evidence that increasing misleadingness reduced trust ($b = -0.10$, $p = 0.646$), nor that the effect of increasing graph beauty on increasing trust was greater than that of reducing graph misleadingness on increasing trust ($\Delta|b| = 0.08$, 95% CI = [-0.17, 0.32]).

Supplementary Study S5

Materials

We conducted supplementary Study S5 to examine the sensitivity of the causal effect of beauty on trust to the statistical power of the study. Study S5 also served as a pilot study for our pre-registered Study 5b. We used the graphs from Study 5a except that the beauty manipulation did not include font change and that only two experiment conditions were examined (more and less beautiful). That is, we used a total of 72 graphs in the present study, with six different topics and six different types, each in two experiment conditions, which were the same as in Study 5b.

Procedures

Each participant viewed and evaluated six graphs. These six graphs were randomly selected for each participant with three constraints: i) they were in six different types; ii) they were of six different topics; and that iii) they were balanced across the two experiment conditions.

Participants completed three different sessions in random orders, in which they viewed the graphs one by one and answered how beautiful the graph was, how much they trusted the graph, and how competent the graph-maker was, respectively. Participants entered their answers by pressing the number keys 1 to 7 on their computer keyboard. After completing all trials, participants filled out a brief survey about their demographics.

Analysis

We confirmed that our beauty manipulation was successful ($b = 1.27$, $p = 1.64 \times 10^{-70}$; $n = 1962$ observations; this linear regression also took into account the random effects of participants, graph types, and graph topics). We observed a positive causal effect of beauty on trust that was of a comparable size to those in Studies S2-S4, and with almost doubled the sample size in the present study, this causal effect was significantly detected ($b = 0.241$, $p = 4.04 \times 10^{-4}$). These

1 findings confirmed that the consistent observation of positive but insignificant casual effects of
2 beauty on trust in Studies S1-S4 was likely due to the insufficient power of those studies.

3

Supplementary References

1. Allen, E. A., Erhardt, E. B. & Calhoun, V. D. Data Visualization in the Neurosciences: Overcoming the Curse of Dimensionality. *Neuron* **74**, 603–608 (2012).
2. Crameri, F., Shephard, G. E. & Heron, P. J. The misuse of colour in science communication. *Nat Commun* **11**, 5444 (2020).
3. Jones, G. *How to lie with charts*. (LaPuerta, 2006).
4. Lauer, C. & O'Brien, S. How People Are Influenced by Deceptive Tactics in Everyday Charts and Graphs. *IEEE Transactions on Professional Communication* **63**, 327–340 (2020).
5. Pandey, A. V., Rall, K., Satterthwaite, M. L., Nov, O. & Bertini, E. How Deceptive are Deceptive Visualizations?: An Empirical Analysis of Common Distortion Techniques. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* 1469–1478 (2015) doi:10.1145/2702123.2702608.
6. Yang, B. W., Vargas Restrepo, C., Stanley, M. L. & Marsh, E. J. Truncating Bar Graphs Persistently Misleads Viewers. *Journal of Applied Research in Memory and Cognition* **10**, 298–311 (2021).
7. Correll, M., Bertini, E. & Franconeri, S. Truncating the Y-Axis: Threat or Menace? in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* 1–12 (ACM, 2020). doi:10.1145/3313831.3376222.