
Is the Dunning-Kruger effect real? Or are unskilled people more rational than it seems?

Independent research report by Belen Cobeta and Spencer Greenberg

Abstract

The Dunning-Kruger effect is a cognitive bias that purportedly makes unskilled people at a given task unaware of their lack of competence because they lack the means to evaluate their own performance. It is often discussed alongside a “false-consensus effect” whereby high performers underestimate their skill level. There is growing literature challenging the existence of the Dunning-Kruger effect. This article draws on a study of 1,817 U.S. participants who completed intelligence tasks and each estimated their own IQ percentile. It uses simulations to test competing explanations for apparent Dunning-Kruger effects. The simulations show that “fake” Dunning-Kruger effects can be generated by floor and ceiling effects in a test, as well as measurement noise (even among perfectly rational, unbiased self-assessors). The simulations also show that rational actors, acting under uncertainty, can experience a genuine “Closer-to-the-Average Effect,” which pulls people’s estimates of their own skill toward the mean. When combined with a “Better-Than-Average” effect, results match the empirical data associated with a Dunning-Kruger effect. The article concludes that a real Dunning-Kruger probably exists, but it may (at least partly) reflect rational responses to uncertainty, rather than pure irrationality.

Keywords

Dunning-Kruger effect, self-assessment, cognitive bias, measurement error, Bayesian reasoning, Better-Than-Average effect, overconfidence, metacognition

Suggested citation: Belen Cobeta and Spencer Greenberg. (2024). Is the Dunning-Kruger effect real? Or are unskilled people more rational than it seems? Version 1.0. Clearer Thinking Research Reports.

© 2026 Clearer Thinking. Licensed CC BY-NC-ND 4.0.

You may have heard the claim that people who are unskilled in a particular domain overestimate their ability in that domain. They don't even know enough to know how bad they are. The formal scientific study of this idea originated in 1999, when Justin Kruger and David Dunning introduced the concept of the Dunning-Kruger effect, a cognitive bias that makes unskilled people at a given task unaware of their lack of competence because

they lack the means to evaluate their own performance.¹ Their studies also showed that high performers underestimate their skill level, which they suggest is caused by the false-consensus effect - where individuals believe their own traits, beliefs, and behaviors are more common in the general population than they actually are. Over the years, numerous studies have demonstrated empirical results similar to the findings that initially led to Kruger and Dunning's conclusions. However, there is growing literature challenging the existence of the Dunning-Kruger effect. How can anyone deny something that's been so widely replicated? Well, buckle up, as we're going to tell you a data-driven story that delves into:

- cognitive biases beyond the Dunning-Kruger effect
- how tricky math can be
- how irrational-seeming behavior isn't necessarily irrational

In their original work, Kruger and Dunning ran four different studies where participants had to complete a test on different domains (humor, logic, and grammar) and estimate how well they did compared to other people similar to them. Participants who were less skilled tended to substantially overestimate their performance - the famous finding that now bears the authors' names.

Recently, we ran our own large study on a wide variety of intelligence tasks that enabled us to test the Dunning-Kruger effect. Since performance on intelligence tasks is, itself, a skill (summarized by a person's "IQ score"), we were able to check whether the Dunning-Kruger effect existed in our data.

Our Study


We recruited 1,817 participants in the United States using our Positly study participant platform², gave them intelligence related tasks to assess their IQ scores, and then gathered their self-assessed IQ percentile by asking them:

"Out of 100 random people in your country who are of your own age and gender, how many of those 100 people do you think you would do better than on an intelligence test designed to accurately measure IQ (if you all took the same test)?"

These self-assessed IQ percentiles can easily be converted to self-assessed IQ scores, since IQ is known to have a distribution that is approximately a bell curve (i.e., "normal distribution"). This graph shows the relationship between measured IQ and self-assessed IQ in our study, with the blue line reflecting the trend in this relationship:

¹Kruger, Justin, and David Dunning. "Unskilled and Unaware of It: How Difficulties in Recognizing One's Own Incompetence Lead to Inflated Self-Assessments." *Journal of Personality and Social Psychology* 77, no. 6 (1999): 1121–1134. <https://doi.org/10.1037/0022-3514.77.6.1121>.

²Positly. "Positly: Find Verified Survey and Study Participants." <https://www.positly.com/>.



We can see that people with higher measured IQs believed themselves to have higher IQs - but not by much!

If both measured IQ and self-assessed IQ were completely accurate (i.e., the test was a perfect test, and people were perfect at evaluating their own scores), all of the data points would fall on the gray dashed line. However, we see that this is not the case: participants with lower measured IQ (i.e., those to the left on the chart) overestimated their IQ, showing a behavior consistent with the Dunning-Kruger effect. This is encapsulated by the fact that, on the left side of the chart, the blue trend line lies above the gray dashed line showing that lower IQ participants overestimated their ability. Likewise, we see that higher IQ participants tended to underestimate their IQs, as encapsulated by the fact that toward the right of the chart the blue line lies below the gray dashed line.

Kruger and Dunning generated a slightly different type of graph to show the results of their studies. For each study, they grouped participants into quartiles of performance and, for each group, they plotted the mean test score and mean perceived ability, both in percentile units.




Figure 1. Perceived ability to recognize humor as a function of actual test performance (Study 1).





Figure 2. Perceived logical reasoning ability and test performance as a function of actual test performance (Study 2).

When we group our data in a similar way we produce a chart that closely resembles Kruger and Dunning's:



Remarkably, there isn't that much of a difference in self-estimated percentile between the 25% of participants who performed the worst (who estimated themselves to be at the 63rd percentile on average) compared to those who were in the top 25% of performers (who estimated themselves as being at the 70th percentile on average).

As in other Dunning-Kruger effect studies, we found that:

- All four quartiles reported a self-assessed IQ percentile greater than 50
- The lowest performers overestimated their performance significantly more than the highest performers underestimated theirs. This can be seen in the chart above by observing that the brown line begins far below the orange line (representing worse performers overestimating their ability) but the brown line ends up above the orange line (representing the top performers underestimating their performance).

So our results seem to have reproduced the Dunning-Kruger effect, suggesting it's very likely real.

But not so fast. The question is, *why* do the lowest performers at a given task consistently overestimate their performance? And do charts like the one above really demonstrate what they seem to?

What do these results mean?

Kruger and Dunning's explanation of the phenomenon was that, since the worst performers lacked the skill being tested, they also lacked the ability to accurately self-assess. This, according to them, is domain specific

so a person can be a victim of the Dunning-Kruger effect when self-assessing themselves on one skill but not another.

In recent years, several studies (such as this one³ and this one⁴) have been published that either deny or downplay the Dunning-Kruger effect, suggesting that the observed patterns in the graphs are due to subtle mathematical phenomena rather than psychological ones. Could it be that these results are simply artifacts of mathematics, rather than telling us something about the human mind?

To explore this question, we conducted different simulations of hypothetical people performing a skill. Each simulation varies these factors:

- How accurate the measure of skill is - is it a reliable test of skill being simulated, or a very noisy test of skill?
- How noisy people's estimates of their own skill are - do people know their own skill level accurately, or is there a lot of random error in these estimates?
- How much bias people have in their estimates of their own skill - do they systematically over (or under) estimate their abilities, and do they anchor their estimates toward the average (e.g., due to uncertainty about their actual skill level)?


We were then able to use these simulations to discover when an apparent Dunning-Kruger effect will appear - and uncover the conditions under which the effect is a real fact about psychology (i.e., a "real" Dunning-Kruger effect) rather than a mathematical artifact (i.e., a "fake" Dunning-Kruger effect).

Simulation 1: Unbiased self-predictors taking the perfect test

First, we examined what happens if study participants are very good at estimating their own skill (i.e., their estimates are accurate, and not biased towards over- or underestimation), and the test they take measuring their skill is perfectly accurate. Does it produce a Dunning-Kruger effect? As we can see in the chart, it does not. People's estimates follow a straight line, closely following the dashed gray line that reflects perfect self-estimation for a test that is a perfect measure of skill:

³Krueger, Joachim, and Ross A. Mueller. "Unskilled, Unaware, or Both? The Better-Than-Average Heuristic and Statistical Regression Predict Errors in Estimates of Own Performance." *Journal of Personality and Social Psychology* 82, no. 2 (2002): 180–188. <https://doi.org/10.1037/0022-3514.82.2.180>.


⁴Nuhfer, Edward, Steven Fleisher, Christopher Cogan, Karl Wirth, and Eric Gaze. "How Random Noise and a Graphical Convention Subverted Behavioral Scientists' Explanations of Self-Assessment Data: Numeracy Underlies Better Alternatives." *Numeracy* 10, no. 1, art. 4 (2017). <http://dx.doi.org/10.5038/1936-4660.10.1.4>.



If you're curious about the details of this scenario, we simulated 1000 people taking a test of skill where:

- We assigned the real skill score for each participant by sampling from a bell curve with a mean of 100 and standard deviation of 15, producing 1000 "real" skill scores
- We set each participant's "measured" score (i.e., the result of them taking a test of skill) to be the same as their real skill score, since the test of skill in this case is assumed to be a perfect measure of the skill.
- Finally, we calculated the self-assessed score for each person to be their real skill score plus a small amount of random noise (with a standard deviation of 5).

Interestingly, even if we introduce noise into people's self-estimates, we get basically the same result. When participants aren't systematically biased in their self-assessments, but they have a moderately high level of inaccuracy at estimating their own skill, the chart barely changes at all:



For this scenario we increased the standard deviation of the estimation error to be twice the standard deviation of the scores (a standard deviation of 30), and still the slope of the resulting trendline is about the same as that of the perfect self-assessments. This is because if people are not biased, as individuals they may overestimate or underestimate their skill, but as a group they are accurate (on average).

Simulation 2: A test that some people perfectly ace or completely fail (i.e., a test with floor and ceiling effects)

So far we have assumed participants were taking a test that perfectly measures their real skill, but that is often not the case. For example, tests often are only capable of measuring skill within a specific range of abilities. If someone's real skill level is below the test's lower limit (the "floor"), they are likely to get the lowest possible score no matter how bad they are - in other words, the test can't distinguish between bad performers and extremely bad performers. In a sense, this inflates the score of extremely bad performers (since they get the


same score as merely bad performers). Similarly, if someone's real skill level is above the test's upper limit (the "ceiling"), they are likely to get the highest possible score no matter how exceptional they are - the test can't tell apart very good performers from incredible performers since they both tend to get the maximum score.

Ceiling and floor effects like these can generate graphs that suggest a small Dunning-Kruger effect when people are also making rough, noisy guesses about their own abilities. But these are "fake" Dunning-Kruger effects - they are mathematical artifacts that have nothing to do with human psychology.

Think of the people whose real skill level is far lower than what the test can measure (they're "below the floor"). As long as they know the test's minimum score, it is nearly impossible for this group of people to underestimate their own performance - the lowest prediction they can give is the minimum score (which is already above their true skill level). But it IS possible for them to overestimate their performance. Because it's nearly impossible for them to underestimate their performance, but possible to overestimate it, and they have uncertainty about their skill level, on average they will overestimate it - an apparent Dunning-Kruger effect! But this is just a mathematical necessity, even for participants making totally rational estimates.

A similar effect happens in reverse for participants who are so skilled that their true skill is above the ceiling of the test - it's nearly impossible for them to overestimate their own performance (since, at most, they can assign themselves the maximum score) but possible to underestimate it, so on average they underestimate.

We can see this phenomenon in action here, as we adjust the parameters of the simulation:



We start with a test without much floor or ceiling effect (left plot, above) and as we lower the ceiling and raise the floor on our test, a Dunning-Kruger effect appears (middle plot) - and then the Dunning-Kruger effect becomes stronger when we increase the noise of the test used to measure skill (right plot). But as we know, this is just due to limitations in the test plus facts about mathematics, so it's a "fake" effect.


What this means is that anyone attempting to study the Dunning-Kruger effect might accidentally end up with a fake Dunning-Kruger effect if they use a test that has floor or ceiling effects! Since some real world tests will be aced or completely failed by a sizable fraction of participants, this urges caution to watch out for this phenomenon.

Simulation 3: A very noisy test of skill

It turns out, there exists a totally different reason why studies may find a fake Dunning-Kruger effect, even when using a test with no floor or ceiling effects.

Real-life tests do not measure skill perfectly - they have uncertainty in their results. In other words, they measure each person's real skill plus some random noise. That is certainly the case, for instance, with commercial IQ tests: the correlation coefficient of the different tasks in the Wechsler Adult Intelligence Scale (WAIS) and general intelligence (g) ranges between 0.5 and 0.85, so each portion of the test is an imperfect measure of IQ. The hope is that by combining enough of these tasks together IQ is measured accurately, but uncertainty inevitably remains.

Our simulations show us that, the noisier the test is at measuring skill, the more it *seems* that lower performer participants overestimate their level of skill, making it look like an example of the Dunning-Kruger effect, when in fact it's just caused by the noisiness of the test itself:



We can intuitively understand why this happens: in the group of participants with the lowest measured scores, we find both those with genuinely low skill and those who were unlucky when they took the test (their measured skill was lower than their true ability). Since we're assuming in this simulation that all participants are perfect estimators of their own skill, the average self-assessed skill in this group (toward the left of the chart) should be higher than the measured score. A similar effect happens in reverse for people on the right of the chart: some of them are truly very skilled, and some were just "lucky" when taking the test of skill, meaning that on average they appear to overestimate their performance.

Like we saw with floor and ceiling effects, this has nothing to do with human psychology, it only has to do with the nature of testing and mathematics. This means that there is a second reason why a fake Dunning-Kruger effect may be found: it will appear to occur when a test of skill has too much measurement noise.

Simulation 4: Rational actors with uncertainty

Now let's explore what happens when we introduce bias in the way participants estimate their skill. "Bias" here simply means that people are not right, on average, about their own skill - they are systematically mispredicting.

Although the term "bias" seems to imply irrationality, this isn't necessarily the case. For instance, imagine that you have no idea how good you are at a given skill. What is it most rational to believe about your skill? Well,

predicting that your skill is about average (or, that your percentile is around 50%) is actually the best prediction you can make, if you truly lack any information about your skill.

Suppose then that you get a small amount of evidence that you are not that great - but the evidence is far from definitive. The rational solution is to then believe that your skill is a bit below average (or that your percentile is slightly below 50%) - that is, to nudge yourself away from thinking you're average based on the evidence. The stronger the evidence is that you have low skill, the more your prediction should diverge from the average.

To make this concrete, suppose you've never played darts (or any game that's similar) before. Initially, you predict that you have average dart skill (among people who also haven't played). If you miss with your first three darts, you might nudge this prediction down a little bit - it's a small amount of evidence that you may be below average - but it's not much evidence since even okay players will miss sometimes, and plus, you're just getting the hang of it. As you keep missing throw after throw though it's rational to keep adjusting your estimate of how bad you are at darts. After 100 terrible throws in a row, you conclude you really are indeed very unskilled at darts.


Such an estimation procedure is actually the rational thing to do! Yet it involves bias, in the sense that, using this procedure, on average, people will mispredict their own skill. People who are less skilled will tend to put themselves closer to the average than they really are (since they start with a prediction of being at the average and adjust down as they get evidence). People with above-average skills who use this method will do the opposite: they will start predicting they are about average, and as they get evidence about their ability they'll keep bumping up their prediction. In other words, both low-skilled and high-skilled people will tend to think they are closer to average than they really are - creating a Dunning-Kruger effect!

A more technical way of viewing the above procedure is through the lens of Bayesianism. A Bayesian rational agent doesn't know what their skill is, so they start with a very flat "prior" encompassing a wide range of possibilities, but which is centered around having average skill. As they acquire evidence about their skill level, they "update" on this evidence, causing the distribution of possibilities for their skill level to start to shift, transforming (as the evidence accumulates) from a very flat distribution to one that is more peaked around a specific skill level. In the interim period while they are accumulating evidence, their best estimate typically lies between their true skill level and an average skill level.

Note that if this occurs it would be a real (not "fake") Dunning-Kruger effect, because if this happens then it's a genuine fact about human behavior, not a mathematical artifact. On the other hand, this sort of Dunning-Kruger effect would occur even from rational actors! In other words, while the Dunning-Kruger effect is usually thought of as a form of irrationality, this effect may occur even if nobody is being irrational!

Let's give this phenomenon of people estimating themselves to be somewhere between the mean and their true level a name: the "Closer-To-The-Average Effect."

To model the Closer-To-The-Average Effect, we created a simulation where participants are relatively accurate at predicting their scores but their estimates of their own ability lie between the mean ability and their true ability - precisely what a rational actor might do if they only had limited evidence about their own ability. Lo and behold, we obtain a Dunning-Kruger effect in our charts from this simulation. Here the parameter "k" reflects how far towards the mean people's estimates of their own ability are skewed - a higher k reflects the situation where a rational actor has less evidence about their own ability and so skews their predictions closer to the mean. As k goes up we see that the Dunning-Kruger effect gets stronger:




Of course, we don't know for sure that what's really happening psychologically is that people are starting with estimating their ability near the mean and then adjusting based on evidence. But it's interesting to observe that the results we see are very similar to what we'd get if that's indeed what they were doing.

Simulation 5: Nearly-rational actors with uncertainty who think they're better than average

The results from the prior simulation come very close to matching our empirical study data - but they don't quite match!


Here we show the relationship again between IQ and self-assessed IQ from our study, and we have added (in green) the trendline of the data with the strongest Dunning-Kruger effect from Simulation 4. See how the trendline of the study data is parallel to that in the simulation, but lies above it on the chart:



Participants in our study, as well as those in Kruger and Dunning's studies, estimated their performance above the 60th percentile on average. This is what's known as the Better-Than-Average (BTA) Effect, a cognitive bias where people rate their abilities as better than average even though it is statistically impossible for most people to have better-than-median abilities. If you are interested, you can read our research on this phenomenon in the New York Times⁵. While people do not show a Better-Than-Average effect for every single skill, it is a common pattern found across a wide range of skills.

To model the Better-Than-Average Effect we added 6 points to each person's self-assessed skill score in the previous simulation so that the mean self-assessed score percentile would be the 65th, just like in our study. If we now compare the results of the simulation to the real study data, here's what we find - the two trend lines (real data, and simulated) are now essentially on top of each other:

⁵Greenberg, Spencer, and Seth Stephens-Davidowitz. "You Are Not as Good at Kissing as You Think. But You Are Better at Dancing." *The New York Times*, April 6, 2019. <https://www.nytimes.com/2019/04/06/opinion/sunday/overconfidence-men-women.html>.




So is there a Dunning-Kruger effect?

The simulations above are remarkable because they show that when researchers are careful to avoid "fake" Dunning-Kruger effects, the real patterns that emerge in Dunning-Kruger studies, can typically be reproduced with just two assumptions:

1. Closer-To-The-Average Effect: people predict their skill levels to be closer to the mean skill level than they really are. This could be rational (when people simply have limited evidence about their true skill level), or irrational (if people still do this strongly when they have lots of evidence about their skill, then they are not adjusting their predictions enough based on that evidence).
2. Better-Than-Average Effect: on average, people tend to irrationally predict they are above average at skills. While this does not happen on every skill, it is known to happen for a wide range of skills. This bias is not the same thing as the Dunning-Kruger effect, but it shows up in Dunning-Kruger plots.

Designing a test that doesn't introduce the "fake" Dunning-Kruger effects can be quite challenging, even when one is aware of it. In our study, our estimate of the correlation coefficient between general intelligence (g) and measured IQ is 0.75 (corresponding to a noise standard deviation of 10 in the above examples). That measurement error, while fairly small, contributes to the flattening of the self-assessments line. We can't be confident about how much our own result can be attributed to a "fake" Dunning-Kruger effect vs. a real one. For many study designs it is difficult to rule out measurement error and floor/ceiling effects as the cause of any apparent Dunning-Kruger effect - very high quality measurements must be made to definitively prove there is a real effect.

But what of Kruger and Dunning's original explanation? Well, it might still be true - perhaps it really is the case that people with lower skill lack the ability to judge their own skill, and that, additionally, high skilled people are suffering from a false-consensus effect. Indeed, Kruger and Dunning provided some evidence of this in their original studies. But it seems to us that these two biases (one impacting low skilled people, the other impacting higher skilled people) don't necessarily provide the most parsimonious explanation. An alternative explanation is that people's judgements undergo a Closer-To-The-Average Effect - that would also explain the plots, and it does so both for low-skilled and high-skilled participants. Interestingly, while the Dunning-Kruger effect is related to people's self-assessments versus their real skill level, a similar phenomenon may emerge in some cases with other positive attributes that are not skills - but that are simply positive attributes we possess. For instance, a study on self-reported physical attractiveness⁶ postulates that "Unattractive people are unaware of their own (un)attractiveness" and produces graphs comparing how physically attractive people have been evaluated to be by third parties versus how attractive they rate themselves to be. And we find the exact same effect - the self-assessment trendline (marked as "subjective") here is flatter than it should be (compare it to the trendline marked as "objective" showing the trend line of ratings by 3rd parties):



⁶Greitemeyer, Tobias. "Unattractive People Are Unaware of Their (Un)Attractiveness." *Scandinavian Journal of Psychology* 61, no. 4 (2020): 471–483. <https://doi.org/10.1111/sjop.12631>.

The above plot shows the same phenomena we've discussed in this article: both a Better-Than-Average Effect and a Closer-To-The-Average Effect. So perhaps these effects don't apply just to skills, but to self-assessments of positive traits more generally.

Are people being irrational?

We've now seen the same phenomena across three different studies: the original Dunning-Kruger work, our own study on IQ scores, and a study on physical attractiveness. In all cases, there is a "Better-Than-Average" effect where people overestimate themselves a bit, on average. Additionally, in all cases, those with the higher measured scores rate themselves as higher than those with the lowest measured scores - but not by much! In other words, people's self-assessed ability seems to go up only slowly with their true ability.

It seems hard to square this with rationality. Presumably, when it comes to logical ability (as in the original Dunning and Kruger study), IQ (as in our study), and physical attractiveness, people have quite a lot of evidence about their own level. This suggests that people may be under-reacting to the evidence they have - predicting themselves closer to the mean than they truly are. So even though it's true that a Closer-To-The-Average Effect can sometimes be rational, people may be underreacting to the evidence they have, which is not. We might call this an "Under Adjustment Bias" - when people irrationally judge themselves to be closer to the average than the evidence warrants. It would be interesting to study this in greater detail to carefully test when it does and does not occur.

Final takeaways

So what did we learn here?

There are some key takeaways from these simulations:

- Fake Dunning-Kruger effects that are just due to mathematical artifacts (rather than psychology) can easily occur - for instance, when a test of skill has a high degree of measurement error, or when the test has ceiling and floor effects. Therefore, great care must be taken in research to be confident one has found a real Dunning-Kruger effect, rather than a fake one. In fact, unless really high-quality skill measurements are made, we can't be sure any given study is showing real Dunning-Kruger effects rather than fake ones! This is true in the case of our IQ study. We hope that more such studies will be conducted.
- A real Dunning-Kruger effect probably still exists, however it may well be the result of rational deliberation. While we can't rule out Kruger and Dunning's original explanation for the effect (i.e., that the lowest performers lack the metacognitive skills needed to accurately assess their own performance), the data matches simulated data where people estimate their skill by starting by assuming they are average and then updating their views away from the average based on evidence they have (the Closer-To-The-Average Effect). The results we see could be rational, if people have a high degree of uncertainty about their skill level, but the observed data is pushed so far towards the average that it suggests people may not be updating their estimates about their skill sufficiently based on the evidence they have (Under Adjustment Bias).
- Generally people are also overconfident in their skills, on average - a Better-Than-Average Effect, that acts above and beyond the Dunning-Kruger effect. This effect must be added to the Dunning-Kruger effect to match the empirical data.

References

1. Kruger, Justin, and David Dunning. "Unskilled and Unaware of It: How Difficulties in Recognizing One's Own Incompetence Lead to Inflated Self-Assessments." *Journal of Personality and Social Psychology* 77, no. 6 (1999): 1121–1134. <https://doi.org/10.1037//0022-3514.77.6.1121>.
2. Positly. "Positly: Find Verified Survey and Study Participants." <https://www.positly.com/>.
3. Krueger, Joachim, and Ross A. Mueller. "Unskilled, Unaware, or Both? The Better-Than-Average Heuristic and Statistical Regression Predict Errors in Estimates of Own Performance." *Journal of Personality and Social Psychology* 82, no. 2 (2002): 180–188. <https://doi.org/10.1037/0022-3514.82.2.180>.
4. Nuhfer, Edward, Steven Fleisher, Christopher Cogan, Karl Wirth, and Eric Gaze. "How Random Noise and a Graphical Convention Subverted Behavioral Scientists' Explanations of Self-Assessment Data: Numeracy Underlies Better Alternatives." *Numeracy* 10, no. 1, art. 4 (2017). <http://dx.doi.org/10.5038/1936-4660.10.1.4>.
5. Greenberg, Spencer, and Seth Stephens-Davidowitz. "You Are Not as Good at Kissing as You Think. But You Are Better at Dancing." *The New York Times*, April 6, 2019. <https://www.nytimes.com/2019/04/06/opinion/sunday/overconfidence-men-women.html>.
6. Greitemeyer, Tobias. "Unattractive People Are Unaware of Their (Un)Attractiveness." *Scandinavian Journal of Psychology* 61, no. 4 (2020): 471–483. <https://doi.org/10.1111/sjop.12631>.