The False Promise of Sexiness:

Perceived Importance, Interest, and Citation Patterns for Counterintuitive Research Findings

Samuel G. B. Johnson ^{1,2,3}

Amanda Royka ⁴

Peter McNally ⁵

Frank C. Keil ⁶

¹ University of Warwick, Department of Psychology

² University of Bath, School of Management

³ University College London, Centre for the Study of Decision-Making Uncertainty

⁴ Queen Mary, University of London, School of Chemical and Biological Sciences

⁵ University of Pennsylvania, Center for Social Norms and Behavioral Dynamics

⁶ Yale University, Department of Psychology

Corresponding Author: Sam Johnson

Email for Correspondence: sgbjohnson@gmail.com

Word Count: 7,326 (main text)

Last Revised: 15 July 2020 (submitted version)

Acknowledgements

We thank Hannah Cohoon, Mallory Kidwell, and Brian Nosek for details about the OSC15 dataset. These results were presented at the Cognitive Science Conference (Johnson, Royka, McNally, & Keil, 2019), and acknowledge the helpful comments from the audience and reviewers. We also thank Richard Lucas, Simine Vazire and audiences at Yale University and the University of Bath for feedback.

Abstract

Decisions about what research is pursued, funded, published, and reported on by the media are often based on what scientific results are perceived as important and interesting. Here, we study one potential cue that scientists and the public may use for evaluating scientific findings—their counterintuitiveness. Two studies establish that counterintuitive findings are perceived as more important and interesting, both by scientists and the public. However, even though judgments of importance predict citations in both samples, counterintuitiveness does not. We argue that these results have implications for scientific institutions and practice.

Keywords: Folk science; lay theories; expertise; interest; public understanding of science; philosophy of science; science methodology

Introduction

Both scientists and the public often assess whether a piece of scientific research is interesting or important. As scientists, these judgments guide which research directions we think are worth pursuing; what work we decide to read and cite; and our peer reviews of others' work. (Indeed, out of a sample of 10 sets of recent journal reviews received by the first author, all but one included explicit assessments of “interest,” “importance,” or both.) Not least, this is because it is often very difficult to tell what research will truly stand the test of time—we have to guess. The public's views on these matters are also important because the public is the ultimate consumer of scientific research: They fund it through their tax contributions (hence, grant panels are responsible to the public) and journalists cover research they believe their readers will want to read about.

Despite their powerful roles in shaping scientific practice, interest and importance are slippery notions. Rather than applying strict definitions, it seems far likelier that we rely on heuristics to guide these judgments. One possibility may be that interest and importance are proportional to how *surprising* or *counterintuitive* a finding is. This heuristic is plausible because it loosely corresponds to defensible criteria derived from philosophy of science and because surprise is psychologically highly accessible. But it is also problematic because surprising findings are more likely to be non-replicable false-positives and can contribute to a discontinuous scientific literature.

In this article, we study this “counterintuitiveness heuristic” empirically. First, to what extent does this heuristic guide scientists' and the public's judgments of research findings? And second, is this heuristic valid in the sense that it tracks an objective metric of academic impact (citations)? Along the way, we collect data that allow us to test subsidiary questions about what other factors lead the public to view scientific findings as interesting and important.

Interest and Importance

Research in science education, science communication, and public understanding of science has tended to look at whether people *understand* (e.g., Bruine de Bruin & Bostrom, 2013) and *trust* (e.g., Jamieson et al., 2019; Kahan, 2010) findings, rather than how important or interesting they are. There are a few exceptions. For instance, non-experts (and particularly conservatives) tended to value scientific *fields* more highly to the extent that they produce more certain knowledge, although this does not seem to influence evaluations of particular findings from those fields (Broomell & Kane, 2017). As another example, Drummond and Fischhoff (2020) examined how individual differences in emotional reactions related to judgments of the quality—particularly the practical value—of a cancer biology finding. While anger was negatively related to judgments of quality, sadness and (most relevant to us) surprise were positively related to quality. Finally, people are known to be “seduced” by reductive information in scientific explanations, such as neuroscience information, finding such explanations more satisfying even if the reductive information is irrelevant (Hopkins, Weisberg, & Taylor, 2016; Weisberg et al., 2008). However, these results may have more to do with the believability of the findings rather than how interesting they are; for example, fMRI lie detection evidence is more persuasive in juror decision-making compared to other kinds of lie detection (McCabe et al., 2011).

Despite the relative lack of research on appraisals of scientific findings in particular, we have excellent *general* models of interest which can motivate predictions. According to appraisal theories of interest (e.g., Silvia, 2005, 2008), the function of interest is to facilitate learning by directing our attention. Thus information that evokes interest tends to be information that is comprehensible (possible to learn) and novel (worth learning). This framework can be used to understand the psychological appeal of counterintuitiveness as well as other potential correlates of scientific interest.

That is, counterintuitive research might be valued because it signals a gap in our knowledge (which can trigger curiosity; Loewenstein, 1994). Likewise, research that is easier to understand may be more interesting because the gap can be more readily traversed (which can trigger fluency; Graf & Landwehr, 2015). Finally, although not explicitly a part of Silvia's (2005) appraisal theory, this theory also seems consistent with people experiencing greater interest in research that is practically or personally relevant to them because the gap is worth traversing.

The Counterintuitiveness Heuristic

Philosophical perspective. If indeed people rely on counterintuitiveness as a proxy for interest or importance, to what extent is this heuristic likely to lead science toward the most fundamental truths? Let's consider how three traditions in philosophy of science might answer this question.

In Popperian theory, theories are scientific to the extent they can be falsified by some potential set of observations (Grünbaum, 1976; Mayo, 1991; Popper, 1959/1934). Popper thus favors "bold" theories (those that are improbable a priori and are thus more informative) combined with "severe" attempts to refute those theories' predictions (targeting predictions that are made uniquely by the theory but not by others). The intuition behind the latter criterion is that severe tests expose the theory to greater risk of falsification, and thus receives greater credit if it survives the test. Thus, Popperian philosophy of science seems to favor counterintuitive results both because they may correspond to stronger tests of a theory or because the theory itself may be more informative.

In Kuhnian theory, science proceeds within paradigms—sets of theories, assumptions, and methods that are taken as background (Kuhn, 1962). At most times, "normal science" prevails, with scientists making incremental contributions within a particular paradigm, with anomalous findings accommodated within the existing paradigm. But when a sufficiently anomalous finding arises, this can create a crisis that leads to a scientific revolution or paradigm shift. A refinement of this theory is provided by Lakatos (1978), who argues that scientific theories are combinations of fundamental and auxiliary assumptions; scientists give up auxiliary assumptions before the fundamental ones, but a sufficiently anomalous finding can cast doubt even on the fundamental assumptions. Thus, particularly counterintuitive findings can be valuable because they may be more likely to cast doubt on more fundamental assumptions and lead to broader theoretical progress.

Finally, in Bayesian theory, theories are evaluated according to the probability calculus (e.g., Howson & Urbach, 2006). A theory's posterior probability $P(H|E)$; plausibility after observing the evidence] is the product of its prior probability $P(H)$; plausibility before observing the evidence] and its relative fit with the data $P(E|H)/P(E)$; the ratio of the evidence's probability given that theory versus its probability ex ante]. Thus, a theory is *more* probable when it makes confirmed predictions that are surprising given other theories [$P(E)$ is low] but *less* probable when it is implausible a priori [$P(H)$]. These correspond to two very different versions of counterintuitiveness. As in Popper, a finding is very valuable if it is cleanly predicted by a theory but unexpected given its competitors. But it is no virtue of a theory to be counterintuitive a priori: Extraordinary claims require extraordinary evidence.

Psychological perspective. There are thus potentially good philosophical reasons to value counterintuitive results, whether conceptualized in terms of severe tests, paradigm shifts, or Bayesian probability. Psychological accounts suggest that the emotion of surprise may roughly track this logic. One theory holds that events are surprising when they contradict expectations (Teigen & Keren, 2003)—falsifying a hypothesis; another holds that events are surprising when they are difficult to integrate with existing theories (Maguire, Maguire, & Keane, 2011)—requiring a paradigm shift; and a third suggests that surprise is a rational statistical inference (Griffiths & Tenenbaum, 2007)—raising

the probabilistic evidence favoring a theory. Surprise is also highly accessible and influences many other psychological processes: It animates our memory for stories (Norenzayan et al., 2006), tendency to share information (Berger & Milkman, 2012), and propensity to learn (Rescorla & Wagner, 1972). As heuristics go, surprise may not be a bad proxy for scientific importance.

But this discussion so far is assuming that people are evaluating how improbable, counterintuitive, or surprising a result is relative to *scientific theories*. Yet, psychologists have demonstrated that people not only hold formal scientific theories in their minds, but that their intuitions are powerfully shaped by *intuitive* or *folk theories* (e.g., Shtulman, 2017), which are highly resistant to change. We have intuitive theories about physics, biology, economics, and psychology, among other domains, which often contradict formal scientific knowledge. Moreover, even scientists' folk intuitions are not supplanted by, but rather co-exist with, their formal scientific knowledge (Goldberg & Thompson-Schill, 2009; Shtulman & Varcarel, 2012). Thus, although the disagreement between a result and existing *scientific* theory is a plausible proxy for importance, disagreement with one's *folk* theory is not, if the folk theory is superseded by scientific understanding.

For example, suppose that a psychologist hypothesizes that behavior is guided by the unconscious activation of stereotypes. The original demonstrations of this phenomenon would have been highly counterintuitive—relative to both scientific and folk theories—and arguably of great scientific significance. However, even though conceptual replications of these priming effects (e.g., using different stereotypes) would no longer contradict *scientific theory*, they would remain counterintuitive relative to our *folk theory*. This creates a misalignment between the scientific and intuitive surprisingness of research. If scientists rely on their folk rather than scientific understanding for evaluating whether research is surprising, they may share this misalignment and their judgments of surprise may no longer be suitable proxies for scientific importance.

Practical perspective. Whether an apt heuristic or not, many scientists claim to prize counterintuitiveness highly. Scientists-in-training are often advised to seek out counterintuitive results. For example, one how-to guide for conducting “interesting” research advises:

One concrete test for evaluating ideas is to imagine the most surprising outcome possible (i.e., the best case scenario). If results were exactly as predicted, would they be interesting? If not, you should dream bigger when hypothesizing or perhaps consider the opposite of your hypothesis—if one way is intuitive, the opposite may be surprising. (Gray & Wegner, 2013, pg. 550.)

Whose intuitions are we contradicting? “Grandmothers, not scientists.... Ideally, research should counter both scientists' and laypeople's intuitions, but we emphasize the latter.” Though laudable as career advice, this raises concerns: Results with low prior probability are likelier to be false-positives, as emphasized by Bayesians. Indeed, the extent to which scientists find a result surprising empirically predicts the likelihood it will fail to replicate (Open Science Collaboration, 2015). This is presumably one reason that submission guidelines for the journal *Psychological Science* now distinguish between “theoretical significance” (which is an acceptance criterion) and “surprising novelty” (which is not).

The Current Studies

Two studies examine three sets of questions. First, do scientists and the public believe that counterintuitive findings are more important or interesting? We use two datasets to answer this question: (1) Data from the reproducibility project (Open Science Collaboration, 2015), which includes surprise and importance ratings from PhD-level scientists for a set of 112 psychology articles, along with several other relevant characteristics of the articles (such as subfield and authors' citations

and institutional prestige) (Study 1); and (2) New data we collected from non-experts on Amazon Mechanical Turk, including ratings of 40 findings from *Psychological Science* and 40 from *Science* (Study 2). Based on the prior literature on surprise and interest, we anticipated that both scientists and non-experts would indeed use counterintuitiveness as a heuristic for importance and interest.

Second, is this counterintuitiveness heuristic a valid predictor of scientific importance? Given that importance is inherently subjective, we instead rely on an objective (albeit imperfect) measure of academic impact, article citation counts. We use the Study 1 and 2 surprise and importance ratings to predict (or fail to predict) articles' actual citation counts between 4 and 12 years post-publication. Since there are good reasons to believe that surprise may or may not be a good proxy for importance (see above), we did not have a strong a priori prediction about this question.

Finally, Study 2 examines judgments aside from surprise that impact public perceptions of interest and importance, namely perceived scientific impact, practical value, and comprehensibility. Based on appraisal theories of interest—which hold that interest is an emotion functioning to guide attention toward information that is both feasible and desirable to learn—we anticipated that practical value and comprehensibility would be strongly associated with interest. We also examined whether these variables are associated with citations.

Study 1

To test the relationship between counterintuitiveness and importance among experts, we re-analyzed data from the Open Science Collaboration (2015) [hereafter OSC15]. The purpose of OSC15 was to attempt to replicate 100 psychology studies, with the headline result that more than half of these studies did not replicate. Less famously, OSC15 also measured the perceived surprise and scientific importance of each original finding among practicing scientists, and found that more surprising findings were less likely to replicate. Here, we re-analyze these data to test whether these expert judgments of surprise were linked with scientific importance, measured both subjectively (rating scales) and objectively (citations).

Method

The OSC15 project recruited groups of researchers to attempt 167 replication projects for studies drawn from 157 unique articles published in *Psychological Science*, *Journal of Personality and Social Psychology* (JPSP), and *Journal of Experimental Psychology: Learning, Memory, and Cognition* (JEP:LMC) in 2008, resulting in 100 completed replication studies. Procedures were followed to ensure that replication teams did not simply choose to replicate articles they found particularly interesting or surprising, with the OSC15 coordinators attempting to match researchers to articles on the basis of the articles the coordinators had chosen, rather than the converse. The study chosen for replication was the last in the article by default, unless it was deemed too infeasible or the original authors suggested a different study. For several studies, surprise and importance ratings were taken even though the replication was not ultimately completed, and we include these in our sample. For articles with multiple replication entries, we selected one entry to avoid double-counting—either the one for which a replication attempt was completed or, if this criterion did not break the tie, the first one listed in the OSC15 dataset. This results in a total of 112 unique articles.

For each finding, the researchers asked a group of 129 coders, separate from the replication teams, to rate the finding on surprise and importance. Coders were scientists who had either completed or were in the process of psychology PhD training, and who were associated with the OSC15 project but uninvolved with and blind to the outcome of the specific replication being rated.

To make these ratings, coders read reports written by the replication teams prior to data collection, summarizing the background, description of the effect, procedure, and analysis plan for each study. For surprise, coders rated “To what extent is the key effect a surprising or counterintuitive outcome?” on a 6-point scale, anchored at 1 (“Not at all surprising”), 2 (“Slightly surprising”), 3 (“Somewhat surprising”), 4 (“Moderately surprising”), 5 (“Very surprising”), and 6 (“Extremely surprising”). For importance, coders rated “To what extent is the key effect an exciting and important outcome?” on a 6-point scale, anchored at 1 (“Not at all exciting and important”), 2 (“Slightly exciting and important”), 3 (“Somewhat exciting and important”), 4 (“Moderately exciting and important”), 5 (“Very exciting and important”), and 6 (“Extremely exciting and important”). Coders could rate as many studies as they wished, and on average each effect was rated 4.52 times, with ratings averaged across coders.

Citation counts were taken at two time points. First, the OSC15 researchers determined the Google Scholar citation count for each article in March 2015. Second, the current authors repeated this procedure in June 2020. Citations were square-root transformed for analysis, as is standard for right-skewed count data. (On the pros and cons of Google Scholar versus other bibliometric databases, see Harzing & Alakangas, 2016.)

The OSC15 data also includes several other variables which are plausibly related to citations and which we include as covariates. First, two measures of the original study’s replicability—a subjective assessment of whether the original finding replicated and the sample size of the original study. Second, five characteristics of the authorship team—the number of citations on Google Scholar (in 2015) for the first and senior authors, the prestige of the institutions of the first and senior authors (on a 7-point scale, based on averaged ratings from both project team members and Mechanical Turk workers), and number of authors. Finally, two characteristics of the underlying research—the number of studies reported and the subfield of psychology (cognitive or social). Data for all studies are available at <https://osf.io/udtxe/>.

Results

First, we used stepwise regression to examine the factors predicting subjective ratings of importance. Across the four model specifications in Table 1, only subjective ratings of surprisingness predicted importance ratings, explaining about 23% of the variance. This held up even after controlling for characteristics of the study’s replicability, research team, and underlying research. This suggests that counterintuitiveness is indeed an aspect of what makes research intuitively important to scientists.

But to what extent do these intuitions about importance track the “objective” scientific importance of research, as proxied by citations? As shown in Table 2, surprise has no association with citations, either taken 7 or 12 years post-publication, explaining less than 1% of the variance in both models. Judgments of importance do have some predictive power, explaining around 5–6% of the variance in citations, suggesting that scientists do indeed have some ability to predict which articles would be more impactful to the scientific community, even though surprise (itself a component of importance judgments) does not.

This is not to say that citations are particularly difficult to predict. Model 3 in Table 2 adds several additional covariates, including whether the finding replicated or not. (This model excludes 15 observations for which the replication project was not completed; however, the results for other covariates are similar when the “replicated” variable is omitted). Several covariates were strong predictors of citations, especially citations to the first author’s work. (The first and senior author citations variable from OSC15 includes the target article’s Year 7 citations, but subtracting these citations and re-fitting the model yields an almost identical fit.) The first author’s institutional prestige also predicted citations, though perhaps to a smaller degree than one might expect. Finally, social psychology articles tended to be cited more than cognitive psychology articles. This may be because

many of the social psychology articles were published in a higher impact journal compared to the cognitive articles (*JPSP* vs *JEP:LMC*), since we do not find this tendency in Study 2, where all psychology articles were published in the same journal. Together, these covariates explained about 42% of the variance in citation counts.

DV: Importance Ratings				
	Model 1	Model 2	Model 3	Model 4
Surprisingness	.394 (.069) ***	.417 (.077) ***	.423 (.081) ***	.431 (.082) ***
Replicated (<i>Dummy</i>)		.042 (.141)	.038 (.147)	.095 (.151)
Original Sample Size (<i>Sqrt</i>)		-.001 (.001)	-.001 (.001)	-.001 (.001)
First Auth. Citations (<i>Sqrt</i>)			.000 (.002)	-.000 (.002)
Senior Auth. Citations (<i>Sqrt</i>)			-.000 (.001)	-.000 (.001)
First Auth. Inst. Prestige			.008 (.081)	.013 (.081)
Senior Auth. Inst. Prestige			.031 (.082)	.025 (.081)
Number of Authors			-.008 (.058)	-.024 (.058)
Number of Studies				.041 (.041)
Social Psychology (<i>Dummy</i>)				.194 (.148)
R ²	.227	.252	.258	.282
N	112	112	97	97
° < .10 * < .05 ** < .01 *** < .001				

Table 1. Regression models predicting importance ratings (Study 1).

Note. Entries are unstandardized *bs* and *SEs*. Sample size and first and senior author citations were square-root transformed; citations to the article under consideration were subtracted prior to this transformation. The Replicated variable is a dummy-coded indicator of whether the original finding replicated in the OSF15 project (15 missing observations), while Social Psychology is a dummy-coded indicator of the original article's subfield (1 = social, 0 = cognitive).

Of course, absence of evidence is not the same as evidence of absence, and *p*-values do not quantify evidence *favoring* the null hypothesis. Thus, in the Supplementary Materials (Appendix A), we report Bayesian versions of these analyses to quantify the evidence favoring the null hypothesis. Later in the main text, we also report an internal meta-analysis combining the data from Studies 1 and 2 to increase statistical power. These analyses all reach the same conclusion—there is no association between how surprising a finding is and how much it is cited.

Discussion

We can draw three conclusions from these results. First, scientists do indeed believe that more counterintuitive findings are more important, explaining about one-fourth of the variance in importance judgments. Second, citation counts *do not* track counterintuitiveness, with this variable explaining less than 1% of the variance in citations. Still, subjective judgments of importance are modestly related to citation counts. Third, citation counts are nonetheless reasonably predictable, with a range of covariates predicting nearly half of the variance. These covariates reflect typical findings in the scientometrics literature, which often finds that authors' reputation is associated with citations (Bornmann et al., 2012).

These data do have limitations. First, they reflect citation patterns only within a single field; it is possible that psychology is uniquely obsessed with counterintuitiveness. Yet, if this were the case, it would be surprising that counterintuitiveness does not drive citations at all, despite predicting

subjective judgments of importance. Perhaps in other fields, intuitions about importance are less reflective of counterintuitiveness. Second, judgments of counterintuitiveness may differ between scientists and the public. It is unclear which of these is more important for driving choices of what research questions to pursue or what articles to publish. In Study 2, we address these limitations by studying non-experts' judgments of research in both psychology and the natural sciences.

	DV: Year 7 Citations (<i>Sqrt</i>)			DV: Year 12 Citations (<i>Sqrt</i>)		
	Model 1	Model 2	Model 3	Model 1	Model 2	Model 3
Surprise	.361 (.393)	-.139 (.438)	.012 (.399)	.320 (.669)	-.624 (.741)	-.400 (.667)
Importance		1.267 (.529) *	.884 (.460) °		2.393 (.896) **	1.696 (.760) *
Replicated (<i>Dummy</i>)			.800 (.643)			1.135 (1.064)
Original Sample Size (<i>Sqrt</i>)			-.002 (.006)			.003 (.011)
First Auth. Citations (<i>Sqrt</i>)			.031 (.010) **			.048 (.016) **
Senior Auth. Citations (<i>Sqrt</i>)			.016 (.006) **			.027 (.010) **
First Auth. Inst. Prestige			.741 (.345) *			1.049 (.571) °
Senior Auth. Inst. Prestige			-.356 (.346)			-.381 (.573)
Number of Authors			.246 (.249)			.349 (.412)
Number of Studies			-.271 (.177)			-.380 (.292)
Social Psychology (<i>Dummy</i>)			2.301 (.635) ***			4.330 (1.050) ***
R ²	.008	.057	.465	.002	.063	.483
N	112	112	97	112	112	97
	° < .10	* < .05	** < .01	*** < .001		

Table 2. Regression models predicting citations (Study 1).

Note. Entries are unstandardized *bs* and *SEs*. Entries are unstandardized *bs* and *SEs*. Sample size and first and senior author citations were square-root transformed; citations to the article under consideration were subtracted prior to this transformation. The Replicated variable is a dummy-coded indicator of whether the original finding replicated in the OSF15 project (15 missing observations), while Social Psychology is a dummy-coded indicator of the original article's subfield (1 = social, 0 = cognitive).

Study 2

Whereas Study 1 examined scientists' judgments about psychology findings, Study 2 turned to non-experts' judgments about findings in psychology and the natural sciences. To do this, we asked groups of laypeople to rate non-technical summaries of articles from either *Psychological Science* (Study 2A) or *Science* (Study 2B) on several dimensions—most critically, interest and importance, as well as potential predictors including surprise, scientific impact, practical value, and comprehensibility. We predicted we would find a similar pattern compared to Study 1—that surprise would be linked to importance (as well as interest) but not to the ultimate citations of the underlying research.

Method

Participants. We recruited 120 participants from Amazon Mechanical Turk ($N = 60$ for each of Studies 2A and 2B). Of these, 57% were female, 42% had completed at least a 4-year college degree, and the average age was 35. Only 8% of participants had doctoral-level training in any field, so the vast majority of participants were laypeople in the specific fields featured in our studies.

Participants were excluded if they incorrectly answered more than 30% of a set of 20 check questions ($N = 8$ and 14 for Studies 2A and 2B).

Materials. For Study 2A, the materials were derived from abstracts of 40 articles appearing in the journal *Psychological Science* in the January, February, and March 2012 issues.

We anticipated that many scientific abstracts would be too long, syntactically complex, and jargon-filled to be comprehensible to most laypeople. As such, the second author produced short summaries of each original abstract to improve comprehension. For example, the actual abstract of one article (Johnson & Fujita, 2012) read:

People are motivated to defend and rationalize the status quo, a phenomenon known as system justification. We propose the existence of a second, countervailing system-level motivation: system-change motivation, which is concerned with bettering the status quo over time. The opportunity to receive diagnostic information about the status quo pits the two system-level motives against each other. Whereas system justification promotes a preference for positive information about the status quo, system-change motivation promotes a preference for negative information about the status quo. In three experiments, we found that people preferred negative over positive feedback about the status quo when it was presented as being changeable. Our findings are the first to suggest the operation of a system-change motive.

Our summary version read:

Research suggests that people are motivated both to defend the status quo and to want to improve it over time. Experimental participants preferred to receive negative information about the status quo when the status quo can be changed. However, when the status quo is unlikely to change, people prefer to receive positive information about it.

Comparable summaries were constructed for all 40 abstracts. Summaries were written at a minimum Flesch–Kincaid grade level of 12 and were of similar length.

We conducted pretests to ensure as strong of a perceived correspondence between the real abstract and summary as possible. In an initial pretest, each participant was assigned to read 10 of the abstracts along with their summaries, and rated their correspondence on a 0 (“A very poor match”) to 10 (“An excellent match”) scale. Any abstract with a score below 7 was targeted for revision and re-normed in a second pretest. All correspondences were rated above the scale midpoint in the second pretest (except one item which was omitted due to a coding error).

For Study 2B, the materials were the “editor’s summaries” of 40 articles published in the January 6, January 13, and January 20, 2012 issues of *Science*. These summaries are written by the editorial staff of the journal, rather than by us, eliminating the risk of experimenter bias. We lightly edited the summaries to match the format of our Study 1 materials (replacing the authors’ names with “Researchers”). For example, the summary of one article (Fermi LAT Collaboration, 2012) read:

Binary star systems that contain a neutron star or a black hole are expected to emit gamma rays. These gamma-ray binaries are a rare class of objects, which are also expected to emit x-rays. Indeed, several such systems were initially detected through their x-ray emission. Researchers have reported the detection of a gamma-ray binary that was previously unknown as an x-ray source. Follow-up observations reveal that the system is also a source of x-rays and that the companion star is a class O star, a type that is very hot and very luminous.

Procedure. In both studies, participants each viewed 10 of the 40 summaries (balanced across participants). For each finding, participants first read the summary and then, on subsequent pages, made six ratings:

Interest. How interesting are these findings to you?

Importance. How important do you think these findings are?

Surprise. How surprising do you think these findings are?

Scientific impact. How much do you think these findings will change the way scientists think about this topic?

Practical value. How useful do you think this finding is on a practical level?

Comprehensibility. How well do you think that you understand the description of this finding?

These ratings were all made on scales from 0 (“Not at all...”) to 10 (“Very...”). Each rating was made on a separate page, with the summary repeated at the top of each page. Our central dependent measures—interest and importance—were always rated first (in a counterbalanced order), followed by the other four ratings in the order above.

After the main task, participants checked off, from a list of 20 concepts, those that had appeared in the summaries. Participants incorrectly answering more than 30% of these check questions were excluded to decrease noise due to inattentiveness.

As an objective measure of academic impact, we obtained the Google Scholar citation counts for each article approximately 4 years (on 26 March 2016 and 17 August 2016 for Studies 2A and 2B, respectively) and 7 years post-publication (on 29 and 30 January 2019). For analyses, each article was categorized as either cognitive or social psychology (in Study 2A, analogous to the OSF15 coding scheme) or as physical sciences/engineering or biomedical sciences/engineering (in Study 2B).

	Study 2A		Study 2B	
	Cognitive	Social	Physical	Biomedical
Importance	5.86 (0.86)	5.63 (1.01)	6.43 (0.92)	6.31 (1.11)
Interest	5.91 (0.97)	5.84 (0.73)	5.08 (1.39)	4.82 (1.57)
Surprise	4.78 (1.07)	3.79 (1.12)	4.55 (0.68)	5.08 (0.72)
Scientific Impact	5.89 (0.78)	4.87 (0.59)	6.19 (0.81)	6.05 (0.99)
Practical Value	5.69 (1.01)	5.33 (0.99)	5.75 (1.29)	5.61 (1.16)
Comprehensibility	7.42 (1.23)	8.14 (0.80)	4.92 (1.61)	4.37 (2.03)
Citations (Year 4)	31.4 (19.5)	59.8 (44.9)	175.2 (157.9)	116.2 (77.1)
Citations (Year 7)	59.5 (33.9)	134.2 (109.4)	298.9 (284.7)	194.2 (139.8)

Table 3. Descriptive statistics for Study 2.

Note. Entries are means (SDs) calculated by item.

Results

Overall, judgments of importance were predicted most robustly by perceived scientific impact and practical value, whereas judgments of interest were predicted most robustly by surprise and comprehensibility. However, these judgments were only poorly predictive of citations. In particular, surprising findings were not cited more (mirroring Study 1), and judgments of importance and interest also did not predict citations.

Descriptive statistics. For initial analyses, we averaged the ratings for each item across participants to calculate item means for each judgment. As shown in Table 3, the natural science findings were viewed as more important, but less interesting, compared to the psychology findings [$t(78) = 2.82, p = .009, d = 0.63$ and $t(78) = -3.51, p = .002, d = -0.78$; p -values corrected using FDR (Benjamini & Hochberg, 1995)]. The natural science findings were also viewed as more surprising and

scientifically impactful [$t(78) = 2.68, p = .011, d = 0.60$ and $t(78) = 3.84, p < .001, d = 0.78$], but of similar practical value [$t(78) = -0.67, p = .51, d = 0.15$]. Finally, the psychology findings were much easier to understand, [$t(78) = 9.39, p < .001, d = 2.10$]. Unsurprisingly, the natural science findings also tended to be cited much more.

<i>Study 2A (Psychology)</i>						
	Importance	Interest	Surprise	Scientific Impact	Practical Value	Comprehensibility
Interest	.60***	—				
Surprise	.49**	.59***	—			
Scientific Impact	.59***	.55***	.76***	—		
Practical Value	.79***	.61***	.36*	.56***	—	
Comprehensibility	.11	.66***	.16	-.05	.17	—
Citations (Year 4)	-.02	.17	-.15	-.21	-.03	.35*
Citations (Year 7)	-.08	.19	-.12	-.25	-.06	.38*
<i>Study 2B (Natural Science)</i>						
	Importance	Interest	Surprise	Scientific Impact	Practical Value	Comprehensibility
Interest	.39*	—				
Surprise	.18	.35*	—			
Scientific Impact	.87***	.40*	.27°	—		
Practical Value	.76***	.45**	-.02	.72***	—	
Comprehensibility	.22	.89***	.18	.18	.35*	—
Citations (Year 4)	.16	.16	-.23	.03	.43**	.27°
Citations (Year 7)	.18	.18	-.25	.04	.45**	.28°

° < .10 * < .05 ** < .01 *** < .001

Table 4. Correlation matrices (Studies 2A and 2B).

Note. Entries are first-order Pearson correlations among measures. Citations are square-root transformed.

Of greater interest for present purposes, Table 4 presents the correlations among variables calculated at the item level, separately for Studies 2A and 2B. Below we report regression models that aggregate these data, account for variability across participants, and tease apart the independent contributions of each predictor to the key variables (importance, interest, and citations). For now, however, we can note commonalities across both psychology and natural science findings.

In both cases, importance and interest were significantly correlated with one another. Yet, these two judgments were associated with somewhat different sets of predictors. Scientific impact and practical value are consistently the strongest predictors of importance, whereas surprise, comprehensibility, scientific impact, and practical value were all consistent predictors of interest.

Whereas in Study 1, judgments of importance among scientists were somewhat predictive of citations, this was not the case for laypeople's judgments either for psychology or natural science. Like Study 1, however, surprise did not predict citations, and if anything was in the opposite direction. That said, lay judgments did have some predictive power for citations, with comprehensibility somewhat associated with citations in both samples and practical value robustly associated with citations among the natural science findings.

DV: Importance				
	Combined	Study 2A (Psychology)	Study 2B (Natural Science)	
Interest	0.27 (0.02) ***	0.36 (0.03) ***	0.18 (0.04) ***	
Surprise	0.03 (0.02)	−0.03 (0.03)	0.14 (0.04) ***	
Scientific Impact	0.22 (0.03) ***	0.18 (0.04) ***	0.23 (0.04) ***	
Practical Value	0.38 (0.02) ***	0.38 (0.03) ***	0.35 (0.03) **	
Comprehensibility	−0.11 (0.02) ***	−0.09 (0.03) **	−0.07 (0.03) *	
N	80	40	40	
DV: Interest				
	Combined	Study 2A (Psychology)	Study 2B (Natural Science)	
Importance	0.40 (0.04) ***	0.49 (0.05) ***	0.28 (0.06) ***	
Surprise	0.16 (0.03) ***	0.12 (0.03) ***	0.19 (0.05) ***	
Scientific Impact	0.03 (0.04)	0.07 (0.04) °	−0.00 (0.05)	
Practical Value	0.11 (0.03) ***	0.15 (0.04) ***	0.08 (0.05) °	
Comprehensibility	0.42 (0.02) ***	0.25 (0.03) ***	0.54 (0.03) ***	
N	80	40	40	
° < .10 * < .05 ** < .01 *** < .001				

Table 5. Multilevel regression models predicting perceived importance and interest (Study 2).

Note. Entries are unstandardized *bs* and *SEs*. *P*-values were calculated using Satterthwaite's method.

Modeling importance and interest. We use multilevel modeling to analyze judgments of importance and interest, because this approach maximizes statistical power by accounting for all of the data in one model, while treating participant and item as random factors. The first column of Table 5 summarizes models predicting importance (top panel) and interest (bottom panel), including random intercepts for participant and for item (nested within study). The VIFs for these models (< 2.25) reflect an acceptable degree of multicollinearity. The second and third columns report models on subsets of the data from Studies 2A and 2B only, with random intercepts for participant and item. We also fit a separate model on the combined data (not depicted in Table 5) that included interaction terms between study and each of the predictors, and we report results from this model below to test whether the parameters differ significantly between the Study 2A and 2B models. In no case does adding subfield improve the explanatory power of the model [$F(3,99.3) = 1.37, p = .26$ and $F(3,99.4) = 0.09, p = .97$ for importance and interest, respectively], so we omit this factor.

The results are mostly consistent with the raw correlations. First, interest and importance predict each other over-and-above the other variables. This effect was stronger for psychology than for the natural sciences [$t(953.4) = 3.50, p < .001$ in the model predicting importance].

Looking next at the other predictors of importance over-and-above interest, importance is positively and strongly predicted by scientific impact and practical value, with these effects similar in magnitude for psychology and the natural sciences ($ps > .39$). The effect for surprise, however, was inconsistent across fields: Surprise was a fairly strong positive predictor for natural science but not for psychology, with the difference in coefficients reaching significance [$t(967.7) = 3.68, p < .001$]. This is probably because surprise is a very strong predictor of scientific impact for psychology but not the natural sciences (Table 4), and the raw correlation coefficients may therefore be more meaningful.

Finally, comprehensibility was a significant *negative* predictor of importance: More comprehensible findings were consistently seen as *less* important. This last effect is consistent with some prior research: Technical language, such as methods descriptions, is associated with a greater perception of “scientificness” and of trust (Thomm & Bromme, 2011), even irrelevant technical details improve perceived explanatory quality (Weisberg et al., 2008), and people often are willing to embrace more complex explanations (Johnson, Valenti & Keil, 2019; Zemla et al., 2017). However, these results are seemingly in tension with research on fluency, which suggests people associate more complex texts with lower intelligence (Oppenheimer, 2006).

Judgments of interest showed a different pattern. Surprise and comprehensibility are consistently robust predictors of interest both for psychology and for natural science, with the effect of surprise similar across fields [$t(937.0) = 1.41, p = .16$], but comprehensibility larger for the natural sciences [$t(595.1) = 5.63, p < .001$]. That is, unlike for importance judgments where comprehensibility seems to be harmful, comprehensibility has an enhancing effect on interest. Scientific impact was not a robust predictor of interest while practical value had a more consistent effect, with these effects not significantly differing in size across fields (p s $> .22$).

	DV: Year 4 Citations (<i>Sqrt</i>)			DV: Year 7 Citations (<i>Sqrt</i>)		
	Model 1	Model 2	Model 3	Model 1	Model 2	Model 3
Importance	1.21 (0.48) *	0.07 (0.87)	−0.34 (0.75)	1.27 (0.65) °	−0.03 (1.18)	−0.75 (1.06)
Interest	−0.52 (0.38)	0.92 (0.89)	−0.83 (0.79)	−0.35 (0.52)	1.07 (1.20)	−1.10 (1.10)
Surprise		−0.27 (0.62)	−0.08 (0.54)		−0.37 (0.83)	−0.04 (0.75)
Scientific Impact		−0.95 (0.95)	−0.77 (0.86)		−1.42 (1.29)	−0.94 (1.20)
Practical Value		1.43 (0.73) °	1.66 (0.61) **		2.08 (0.99) *	2.42 (0.85) **
Comprehensibility		−1.09 (0.51) *	0.91 (0.55)		−1.15 (0.69) °	1.27 (0.77)
Social Psychology (<i>Dummy</i>)			0.85 (1.31)			2.12 (1.83)
Physical Sciences (<i>Dummy</i>)			8.57 (1.42) ***			10.96 (1.98) ***
Biomedical Sciences (<i>Dummy</i>)			7.25 (1.46) ***			9.20 (2.03) ***
R ²	.081	.174	.468	.048	.137	.401
N	80	80	80	80	80	80
	° < .10	* < .05	** < .01	*** < .001		

Table 6. Regression models predicting citations (Study 2).

Note. Entries are unstandardized *bs* and *SEs*. Dummy-coded subfield variables are relative to cognitive psychology (coded as ‘0’).

Modeling citations. Since citations are an item-level rather than participant-level dependent variable, we take item as the unit of observation. Table 6 summarizes OLS results including subfield (dummy-coded) as predictors, but a multilevel model treating subfield as random yields almost identical coefficients and significance levels compared to Model 3. Adding interaction terms between subfield and the other predictors does not improve model fit [$F(20,52) = 1.27, p = .24$], signaling a lack of evidence that predictors of citations differed across subfields.

This analysis reveals that the only judgment that consistently predicts practical value citations is practical value. Interest, importance, scientific impact had no predictive power once practical value is included in the model, and indeed the coefficients on these predictors were, if anything, negative. There is perhaps a slight, but non-significant, trend for more readily comprehensible findings to be cited more; although there is a negative coefficient when subfield is not adjusted for, this trend is reversed when subfield is included in the model, since the natural science findings were both less comprehensible but more heavily cited, relative to psychology.

Discussion

The key results are similar to Study 1: Counterintuitiveness is linked to interest (though less consistently to perceived importance), but has no relationship with citations. Thus, the findings we find among practicing scientists appear to generalize to non-experts.

This study also allowed us to assess other predictors of interest, importance, and citations. Interest was predicted by comprehensibility and practical value, perceived importance by practical value and perceived scientific impact (and *negatively* by comprehensibility), and citations only by practical value. This last result is particularly important because it establishes that our method has sufficient precision and power to explain at least some of the variance in citations—our model including all judgments accounted for around 13–18% of the variance in citations. But surprise is, if anything, negatively associated with citations (albeit not significantly).

Internal Meta-Analysis

To further increase power, we performed a meta-analysis at the item level, using the data from both Studies 1 and 2 (a total of 192 findings). We fit a multi-level model with a random intercepts for field (cognitive psychology, social psychology, physical science, or biomedical science) and study (1 or 2). We included surprise and importance as predictors, since these were measured in both Studies 1 and 2, and we centered and scaled these predictors separately for each study to account for the differing measurement scales used across the studies. We chose Year 7 citations as the dependent variable since this was measured for both datasets.

DV: Year 7 Citations (<i>Sqrt</i>)		
	Model 1	Model 2
Surprise	−0.12 (0.30)	−0.44 (0.34)
Importance		0.74 (0.34) *
N	192	192

Table 7. Meta-analysis of Studies 1 and 2

Note. Entries are unstandardized *bs* and *SEs*.

Table 7 shows, consistent with the results of each study individually, importance is a significant predictor of citations, whereas surprise is not. We also tested several other specifications, including an analysis in which study is treated as a fixed effect and interacted with surprise and importance; however, this model did not yield an improved fit over the reduced model without the interaction terms [$F(2,184.6) = 0.14, p = .87$], indicating a lack of evidence that the judgments of the experts in Study 1 and non-experts in Study 2 differ in their predictive power for citations.

Appendix A also reports a Bayesian version of this meta-analysis, which confirms that the evidence favors the null hypothesis: The data are about 9 times more probable under the null model that there is no effect of surprise on citations.

General Discussion

How do scientists and the public decide what research is interesting and important? Here, we find that both scientists and the public rely on a *counterintuitiveness heuristic*—believing that more surprising findings are more interesting and important. However, more counterintuitive findings were not associated with higher citation rates in either study or in an internal meta-analysis.

These results contribute to discussions about intuitions about science (Gottlieb & Lomrozo, 2018; Hopkins et al., 2016; Johnson, Kim, & Keil, 2016; Weisberg et al., 2008) by establishing what judgments are associated with perceived scientific importance; to scientometrics (e.g., Bornmann et al., 2012) by studying the (non)predictors of citation patterns; and to the psychology of interest generally (Silvia, 2005) by uncovering predictors of interest that broadly mirror appraisal theories.

The data do merit some caution in interpretation. None of these datasets are perfect. The OSC15 data set used in Study 1 is limited to psychology articles. The Study 2 data relied on brief summaries of findings, rather than the extensive descriptions given to the Study 1 participants. And both samples include a limited number of articles. Yet, these data sources are complementary; for example, they include both experts and non-experts, results published in several different fields, subfields, and journals (including highly elite journals such as *Science* and somewhat more specialized journals such as *JEP:LMC*), and a combined sample size of 192 articles. It is of course possible that an effect of counterintuitiveness on citations will be discovered with a much larger sample, but such an effect would likely be very small.

Interpretations and Implications

What should we make of these results? We consider two possible interpretations.

Counterintuitiveness as a rational cue to importance. As we reviewed in the Introduction, there are good reasons to believe that more surprising findings are more important. The key argument is that if a result is uniquely predicted by one theory, it is likely to be surprising if that theory's competitors currently dominate in our priors. That is, surprising results can be especially informative in distinguishing between theories. Introductions to articles often outline multiple different theoretical perspectives, explaining what hypotheses each theory would generate and typically concluding that only one theory is consistent with the result. Thus, scientists certainly seem to be adept at this kind of reasoning.

Why, then, are more surprising findings not cited more, despite a significant effect of perceived *importance* on citations in both studies? One possibility is that citations simply are a poor proxy for actual scientific importance. Some might even argue that particularly counterintuitive findings are disadvantaged because they do not fit into the dominant paradigm. In some ultimate sense, there is no way to measure scientific quality absent a *theory* or definition of scientific quality—which is of course controversial, as well as question-begging in the current context if surprise is a part of the definition. Still, we know that factors independent of scientific quality (such as authors' reputation, discipline and sub-discipline, etc.; e.g., Bornmann et al., 2012) influence citations, so it may well be the case that the signal to noise ratio is very small indeed. From this perspective, our participants' *intuitions* are rational in using the counterintuitiveness heuristic, while the broader scientific community is irrational in ignoring it in their citation patterns.

Counterintuitiveness as a cognitive bias. An alternative perspective—with which, we confess, we have greater sympathy—is that the counterintuitiveness heuristic is a cognitive bias. On this view,

when scientists and the public evaluate a finding as counterintuitive, this evaluation is often made relative to their *intuitive* theory rather than their scientific theories. Research on intuitive theories has demonstrated that they co-exist with formal scientific knowledge among both experts and non-experts (Goldberg & Thompson-Schill, 2009; Shtulman, 2015, 2017). Thus, we are unlikely to fully internalize counterintuitive scientific discoveries—we will continue to be intuitively surprised the next time we encounter that finding. Thus, judgments of surprise are, at best, a blend of our scientific and intuitive judgment, making them a potentially poor proxy for scientific importance.

Although making this strong claim would go beyond the data, several pieces of evidence are at least consistent with this claim. First, counterintuitiveness is equally (non-)predictive of citations among experts and non-experts. If scientists were tapping into sophisticated scientific intuitions when rating how surprising a finding is, one would expect their judgments to be better predictors compared to those of the public. Second, social psychology is classically considered a more “counterintuitive” discipline relative to lay intuitions compared to cognitive psychology, and, befitting this reputation the experts in Study 1 rated the social psychology findings somewhat more surprising [$t(110) = 1.89, p = .062, d = 0.36$]. It seems unlikely that this is because social psychology has stronger theories than cognitive psychology, with disconfirmations of those theories’ highly constrained predictions especially surprising. Rather, it seems likely that people (including the expert coders) simply have stronger folk intuitions about the phenomena studied in social psychology, which they use to evaluate those phenomena. Finally, anecdotally, scientists are often advised to “be counterintuitive” and specifically toward “grandmothers, not scientists” (Gray & Wegner, 2013). In our experience, scientists often seem to follow this advice, harnessing their inner grandmothers when evaluating whether a result is “cute,” “sexy,” or counterintuitive.

Implications. It is important to understand how people think about science because scientific progress tracks social priorities—scientists serve largely at the pleasure of society. To the extent that we have systematic misconceptions about science, we must understand how those misconceptions might thwart the dissemination of science to the public, or even scientific progress itself. To the extent that experts have irreducible preferences over the kind of science they like to do or non-experts over the kind of science they like to see, we must understand how those preferences might be reflected in the kind of research produced by scientific institutions.

Together with OSC15, we now know two stylized facts about counterintuitive findings: They are not particularly well-cited, and they are less likely to survive replication attempts (Open Science Collaboration, 2015). Yet, we find that experts find them more important—and of course, editors, reviews, and grant panel members are typically experts—and non-experts find them more interesting—and journalists and news consumers are typically non-experts.

These facts are a matter of the scientific record, but what to do about them is a challenging and ultimately value-laden question. Here we present a personal perspective.

We suggest that, at the very least, editors and reviewers should be aware of this pattern as they evaluate other scientists’ work and that journalists should be aware of this pattern when deciding what research to cover. Editors often attach higher priority to research that relies on field studies, large samples, or real behavior as opposed to surveys. Such preferential treatment can be justified scientifically, as these are strong cues to external, statistical, and construct validity, respectively. To the extent that additional weight is similarly given to counterintuitive findings, however, this practice may ultimately do more harm than good: The normative basis of such a preference is arguably undermined by the lower replicability of such findings coupled with a lack of evidence that such effects are actually more scientifically important.

However, personal caution may not do much to change the balance of surprising to non-surprising results published in journals or covered in magazines. Editors are accountable to publishers and society

boards, which certainly wish to maximize citations to their journals but also their readership among scientists and coverage among journalists. Journalists are responsible to their editors and ultimately to their readers—and readers may very well demand counterintuitive findings. Although one occasionally hears complaints among the public that too much science has become “click bait,” journalistic practice seems to have changed relatively little and our results suggest that they may have a good economic rationale for their practices as they attempt to market their work. With these pressures in place, scientists themselves have little incentive to change course.

Yet, the past decade has seen remarkable advances in scientific practice across many disciplines, perhaps most impressively in psychology. Many journals ask for sample size justifications, praise preregistered studies, and demand open data and extensive robustness checks. In 2012, many practices we now consider ethically suspect were widespread—more than half of psychologists admitted to omitting dependent measures from a study write-up or making decisions on collecting additional data based on whether a sample yielded significant results (John, Loewenstein, & Prelec, 2012). We believe these practices are far less common now, only 8 years later. Perhaps a similar reckoning is in order over the relative priority of counterintuitive science.

References

- Benjamini, Y., and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society Series B*, 57, 289–300.
- Berger, J., & Milkman, K. L. (2012). What makes online content go viral? *Journal of Marketing Research*, 49, 192–205.
- Bornmann, L., Schier, H., Marx, W., & Daniel, H. (2012). What factors determine citation counts of publications in chemistry besides their quality? *Journal of Infometrics*, 6, 11–18.
- Broomell, S. B., & Kane, P. B. (2017). Public perception and communication of scientific uncertainty. *Journal of Experimental Psychology: General*, 146, 286–304.
- Bruine de Bruin, W., & Bostrom, A. (2013). Assessing what to address in science communication. *Proceedings of the National Academy of Sciences*, 110, 14062–14068.
- Drummond, C., & Fischhoff, B. (2020). Emotion and judgments of scientific research. *Public Understanding of Science*, 29, 319–334.
- Fermi LAT Collaboration (2012). Periodic Emission from the Gamma-Ray Binary 1FGL J1018.6–5856. *Science*, 335, 189–193.
- Goldberg, R.F., & Thompson-Schill, S.L. (2009). Developmental “roots” in mature biological knowledge. *Psychological Science*, 20, 480–487.
- Gottlieb, S., & Lombrozo, T. (2018). Can science explain the human mind? Intuitive judgments about the limits of science. *Psychological Science*, 29, 121–130.
- Graf, L. K. M., & Landwehr, J. R. (2015). A dual-process perspective on fluency-based aesthetics: The pleasure–interest model of aesthetic liking. *Personality and Social Psychology Review*, 19, 395–410.
- Gray, K., & Wegner, D. M. (2013). Six guidelines for interesting research. *Perspectives on Psychological Science*, 8, 549–553.
- Grünbaum, A. (1976). Is the method of bold conjectures and attempted refutations justifiably the method of science? *British Journal for the Philosophy of Science*, 27, 105–136.
- Harzing, A., & Alakangas, S. (2016). Google Scholar, Scopus, and the Web of Science: A longitudinal and cross-disciplinary comparison. *Scientometrics*, 106, 787–804.
- Hopkins, E.J., Weisberg, D.S., & Taylor, J.C.V. (2016). The seductive allure is a reductive allure: People prefer scientific explanations that contain logically irrelevant reductive information. *Cognition*, 155, 67–76.
- Howson, C., & Urbach, P. (2006). *Scientific reasoning: The Bayesian approach*. La Salle, IL: Open Court.
- Jamieson, K. H., McNutt, M., Kiemer, V., & Sever, R. (2019). Signaling the trustworthiness of science. *Proceedings of the National Academy of Sciences*, 116, 19231–19236.
- John, L. K., Loewenstein, G., & Prelec, D. (2012). Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological Science*, 23, 524–532.
- Johnson, I. R., & Fujita, K. (2012). Change we can believe in: Using perceptions of changeability to promote system-change motives over system-justification motives in information search. *Psychological Science*, 23, 133–140.
- Johnson, S.G.B., Kim, K., & Keil, F.C. (2016). The determinants of knowability. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Johnson, S.G.B., Royka, A., McNally, P., & Keil, F.C. (2019). When is science considered interesting and important? In *Proceedings of the 41st Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Johnson, S. G. B., Valenti, J. J., & Keil, F. C. (2019). Simplicity and complexity preferences in causal explanation: An opponent heuristic account. *Cognitive Psychology*.
- Kahan, D. M. (2010). Fixing the communications failure. *Nature*, 463, 296–297.

- Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In *Heuristics and biases: The psychology of intuitive thought* (pp. 49–81). New York, NY: Cambridge University Press.
- Kuhn, T.S. (1962). *The structure of scientific revolutions*. Chicago, IL: University of Chicago Press.
- Lakatos, I. (1978). Falsification and the methodology of scientific research programmes. In J. Worrall and G. Currie (Eds.), *The Methodology of Scientific Research Programmes (Philosophical Papers: Vol. 1)*. Cambridge, UK: Cambridge University Press.
- Leshner, A. I. (2003). Don't let ideology trump science. *Science*, 302, 1479.
- Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin*, 116, 75–98.
- Lombrozo, T. (2007). Simplicity and probability in causal explanation. *Cognitive Psychology*, 55, 232–257.
- Maguire, R., Maguire, P., & Keane, M.T. (2011). Making sense of surprise: An investigation of the factors influencing surprise judgments. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37, 176–186.
- Mayo, D. (1991). Novel evidence and severe tests. *Philosophy of Science*, 58, 523–552.
- McCabe, D. P., Castel, A. D., & Rhodes, M. G. (2011). The influence of fMRI lie detection evidence on juror decision-making. *Behavioral Sciences & The Law*, 29, 566–577.
- Norenzayan, A., Atran, S., Faulkner, J., & Schaller, M. (2006). Memory and mystery: The cultural selection of minimally counterintuitive narratives. *Cognitive Science*, 30, 531–553.
- Open Science Collaboration (2015). Estimating the reproducibility of psychological science. *Science*, 349, aac4716.
- Oppenheimer, D.M. (2006). Consequences of erudite vernacular utilized irrespective of necessity: Problems with using long words needlessly. *Applied Cognitive Psychology*, 20, 139–156.
- Popper, K. (1959). *The logic of scientific discovery*. London, UK: Routledge. (Original work published 1934.)
- Rescorla, R.A., & Wagner, A.R. (1972). A theory of Pavlovian conditioning: Variations on the effectiveness of reinforcement and non-reinforcement. In *Classical conditioning II: Current research and theory* (pp. 64–99). New York, NY: Appleton-Century-Crofts.
- Shtulman, A. (2015). How lay cognition constrains scientific cognition. *Philosophy Compass*, 10/11, 785–798.
- Shtulman, A. (2017). *Scienceblind: Why our intuitive theories about the world are so often wrong*. New York, NY: Basic Books.
- Shtulman, A., & Valcarcel, J. (2012). Scientific knowledge suppresses but does not supplant earlier intuitions. *Cognition*, 124, 209–215.
- Silvia, P. J. (2005). What is interesting? Exploring the appraisal structure of interest. *Emotion*, 5, 89–102.
- Silvia, P. J. (2008). Interest—the curious emotion. *Current Directions in Psychological Science*, 17, 57–60.
- Teigen, K. H., & Keren, G. (2003). Surprises: Low probabilities or high contrasts? *Cognition*, 87, 55–71.
- Thomm, E., & Bromme, R. (2011). “It should at least seem scientific!” Textual features of “scientificness” and their effect on lay assessments of online information. *Science Education*, 96, 187–211.
- Weisberg, D.S., Keil, F.C., Goodstein, J., Rawson, E., & Gray, J.R. (2008). The seductive allure of neuroscience explanations. *Journal of Cognitive Neuroscience*, 20, 470–477.
- Zemla, J. C., Sloman, S., Bechlivanidis, C., & Lagnado, D. A. (2017). Evaluating everyday explanations. *Psychonomic Bulletin & Review*, 24, 1488–1500.

SUPPLEMENTARY MATERIALS

The False Promise of Sexiness:

Perceived Importance, Interest, and Citation Patterns for Counterintuitive Research Findings

Johnson, Royka, McNally, & Keil (2020)

Appendix A: Bayesian Analyses

Since one of the primary effects of interest—the effect of surprise on citations—has turned out to be null, we followed up on this null effect with a Bayesian analysis. Whereas traditional null hypothesis significance testing does not lend itself to quantifying the evidence for the *absence* of an effect, Bayesian methods do.

Table A1 presents a Bayesian regression model, fit using the ‘BayesFactor’ package in R (Morey et al., 2018), which implements the method of Liang et al. (2008; see also Rouder & Morey, 2013). In this model, we include all predictors used in the most complex model summarized in Table 2 of the main text, except we omit the dummy-coded indicator for whether the original effect replicated, as this variable has 15 missing observations and thus its inclusion reduces statistical power.

The ‘BF’ entry for each predictor represents the relative fit of the full model to a reduced model excluding that predictor. When the evidence favors the reduced model, it is labeled as ‘H₀’, indicating that the evidence favors a null effect of that predictor. When the evidence favors the full model, it is labeled ‘H₁’, indicating that the evidence is more consistent with a non-zero effect. The BF itself is a ratio (the “likelihood ratio”) of how well the data fit each model. Thus, a BF of 3 favoring the null hypothesis means that the data would be three times likelier if there is no effect of that parameter than if there is a non-zero effect.

The ‘Credible Interval’ entry presents the Bayesian equivalent of a 95% confidence interval. It summarizes the posterior distribution for each predictor’s parameter estimate, based on a 10,000 sample MCMC chain. Thus, ‘50%’ refers to the median of the posterior distribution, while ‘2.5%’ and ‘97.5%’ refer to the lower and higher bounds on the credible interval. Credible intervals (unlike confidence intervals) can be interpreted as a range in which there is a 95% chance that the true value of the parameter falls, given the prior distribution and the data.

The results are consistent with the frequentist analysis reported in Table 2 of the main text. BF’s favor the alternative hypothesis for importance, first and senior author citations, first author institutional prestige, and the social psychology dummy variable. Most importantly, the evidence favors a null effect for surprise. For Year 7 citations, the BF is 2.1, indicating that the data are about twice as likely under the null rather than alternative hypothesis, indicating weak evidence favoring the null hypothesis. For Year 12 citations, the BF is 1.1—but this is because the model is indifferent between a null effect and a *negative* effect of surprise on citations. Thus, these data corroborate the frequentist analysis in the main text.

	DV: Year 7 Citations (<i>Sqrt</i>)					DV: Year 12 Citations (<i>Sqrt</i>)			
	BF	Credible Interval			BF	Credible Interval			
		2.5%	50%	97.5%		2.5%	50%	97.5%	
Surprise	H ₀ : 2.1	−0.99	−0.29	0.40	H ₀ : 1.1	−2.01	−0.85	0.31	
Importance	H ₁ : 2.8	0.02	0.85	1.66	H ₁ : 4.9	0.23	1.62	3.05	
Original Sample Size (<i>Sqrt</i>)	H ₀ : 3.0	−0.01	0.00	0.01	H ₀ : 2.5	−0.01	0.01	0.03	
First Auth. Citations (<i>Sqrt</i>)	H ₁ : 21.5	0.01	0.02	0.04	H ₁ : 18.5	0.01	0.04	0.07	
Senior Auth. Citations (<i>Sqrt</i>)	H ₁ : 8.1	0.00	0.01	0.02	H ₁ : 7.5	0.00	0.02	0.04	
First Auth. Inst. Prestige	H ₁ : 2.3	−0.02	0.63	1.28	H ₁ : 1.1	−0.25	0.86	1.95	
Senior Auth. Inst. Prestige	H ₀ : 1.9	−0.94	−0.31	0.33	H ₀ : 2.7	−1.37	−0.29	0.82	
Number of Authors	H ₀ : 2.5	−0.31	0.12	0.55	H ₀ : 2.8	−0.55	0.15	0.86	
Number of Studies	H ₀ : 1.7	−0.47	−0.16	0.14	H ₀ : 2.3	−0.70	−0.19	0.31	
Social Psychology (<i>Dummy</i>)	H ₁ : 40.1	0.62	1.79	2.93	H ₁ : 223.0	1.64	3.49	5.41	

Table A1. Bayesian regression for Study 1.

Note. Entries are the Bayes Factors favoring either the null (H₀) or alternative hypothesis (H₁) and the 95% credible interval on each parameter estimate. Sample size and first and senior author citations were square-root transformed; citations to the target article were subtracted prior to this transformation. Social Psychology is a dummy-coded indicator of the original article's subfield (1 = social, 0 = cognitive).

We also performed a similar analysis on the citation data for Study 2, based on Model 3 in Table 6. The results again mirror the frequentist analysis. Practical value was the only judgment with clear evidence of a positive effect on citations, with a BF between 7 and 11. Similarly to the Study 1 model, the BFs for surprise were between 2.5 and 3.0, favoring the null hypothesis.

		DV: Year 4 Citations (<i>Sqrt</i>)				DV: Year 7 Citations (<i>Sqrt</i>)			
		BF	Credible Interval			BF	Credible Interval		
			2.5%	50%	97.5%		2.5%	50%	97.5%
Importance	H ₀ : 2.7		-1.70	-0.29	1.08	H ₀ : 2.1	-2.54	-0.62	1.30
Interest	H ₀ : 1.8		-2.21	-0.71	0.73	H ₀ : 1.7	-2.91	-0.91	1.05
Surprise	H ₀ : 2.9		-1.06	-0.07	0.93	H ₀ : 2.6	-1.37	-0.02	1.34
Scientific Impact	H ₀ : 2.0		-2.27	-0.65	0.93	H ₀ : 2.0	-2.93	-0.80	1.35
Practical Value	H ₁ : 7.9		0.30	1.42	2.55	H ₁ : 10.5	0.47	2.01	3.56
Comprehensibility	H ₁ : 1.1		-0.25	0.77	1.82	H ₁ : 1.2	-0.32	1.06	2.48
Social Psychology (<i>Dummy</i>)	H ₀ : 2.4		-1.70	0.72	3.17	H ₀ : 1.5	-1.60	1.75	5.09
Physical Sciences (<i>Dummy</i>)	H ₁ : >1000		4.52	7.36	10.17	H ₁ : 734.6	5.25	9.14	13.0
Biomedical Sciences (<i>Dummy</i>)	H ₁ : >1000		3.39	6.25	9.05	H ₁ : >1000	3.79	7.65	11.6

Table A2. Bayesian regression for Study 2.

Note. Entries are the Bayes Factors favoring either the null (H₀) or alternative hypothesis (H₁) and the 95% credible interval on each parameter estimate. Sample size and first and senior author citations were square-root transformed; citations to the target article were subtracted prior to this transformation. Social Psychology, Physical Sciences, and Biomedical Sciences are dummy-coded indicators of the original article's subfield (where 0 = Cognitive Psychology).

Finally, we also computed BFs for the meta-analysis described in the main text. We used Rouder and Morey's (2011) method for combining *t*-statistics into a Bayesian meta-analytic estimate (see also Rouder et al., 2009). We separately fit Model 2 of Table 7 for each of the two studies, yielding *t*-

statistics of -0.729 and -0.940 with sample sizes of 112 and 80, respectively. Using a standard Cauchy prior, this yields a BF of 8.9 favoring the null hypothesis.

References

- Liang, F., Paulo, R., Molina, G., Clyde, M. A., Berger, J. O. (2008). Mixtures of g-priors for Bayesian Variable Selection. *Journal of the American Statistical Association*, 103, 410–423.
- Morey, R. D., Rouder, J. N., Jamil, T., Urbanek, S., Forner, K., & Ly, A. (2018). Package 'BayesFactor'. <https://cran.r-project.org/web/packages/BayesFactor/BayesFactor.pdf>.
- Rouder, J. N. & Morey, R. D. (2011). A Bayes Factor meta-analysis of Bem's ESP claim. *Psychonomic Bulletin & Review*, 18, 682–689.
- Rouder, J. N., & Morey, R. D. (2013) Default Bayes Factors for model selection in regression. *Multivariate Behavioral Research*, 47, 877–903.
- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t -tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, 16, 225–237.