



Using big data to track major shifts in human cognition

Simon DeDeo^{a,b,1}

Scheffer et al.'s (1) exciting new work reports an historic rearrangement, occurring in the late 20th century, of the balance between reason and emotion. Its approach is part of a new trend in the psychological sciences that uses extremely large volumes of text to study basic patterns of human cognition. Recent work in this vein has included studies of the universal properties of gender representations (2), the rise of causal thinking (3), and a cognitive bias towards positivity in language itself (4). The goal of going "from text to thought" (5) is an attractive one, and the promise of the machine learning era is that we will only get better at extracting the imprints left, in text, by the mechanisms of the mind.

To establish their claims, Scheffer et al. (1) use principal component analysis to identify two major polarities of correlated vocabulary words in the Google Books corpus (6). The first polarity (PC1) tracks a shift from archaic to modern, in both material life ("iron" is archaic, "computer" is modern) and culture ("liberty" is archaic, "privacy" is modern). The second polarity (PC2) that emerges is the intriguing one, and forms the basis of their paper: Its two poles, the authors argue, correspond to the distinction between "rational" and "intuitive" language.

Their main finding then has two pieces: a shift from the intuitive pole to the rational pole (the "rise" of rationality) and then back (the "fall") (1). The rise has begun by the start of their data in 1850, and unfolds over the course of a century or more. They attribute it to a society increasingly concerned with quantifying, and justifying, the world through scientific and impersonal language—a gradual tightening of Max Weber's famous "iron cage" of collectivized, rationalized bureaucracy in service of the capitalist profit motive (7). The fall, meaning a shift from the

rational back to the intuitive, begins in 1980, and is more rapid than the rise: By 2020, the balance is similar to that seen in the early 1900s. The fall appears to accelerate in the early 2000s, which leads the authors to associate it with social media use and a "post-truth era" where "feelings trump facts." Both these interpretations are supported by accompanying shifts toward "collective" pronouns (we, our, and they) in the Weberian period, and then toward the "individualistic" ones (I, my, he, and she) after.

The raw effect sizes the authors report are extraordinarily large (1). At the peak in 1980, rationality words outnumbered intuition words, on average, three to one. Forty years later (and 100 y earlier), however, the balance was roughly one to one. If these represent changes in actual language use, let alone the time devoted to the underlying cognitive processes, they are enormous shifts in the nature of human experience. Anyone interested in the history, or future, of human cognition will ask: Are they real?

The question comes in two parts. First, did the hypostatized shifts in word usage actually occur, or are they an artifact of the corpus or how it was analyzed? Second, what is the relationship between shifts in word usage and underlying psychological or social reality?

The Corpus Challenge

The primary corpus for Scheffer et al. (1) is Google Books. The release of that corpus was a major event, but researchers in cultural analytics later uncovered a range of deficits (8). While Scheffer et al. are aware of these, it is natural to worry about further unknown unknowns that accompany a proprietary database whose detailed composition has never been made public.

^aDepartment of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA 15213; and ^bSanta Fe Institute, Santa Fe, NM 87501
Author contributions: S.D. designed research, performed research, analyzed data, and wrote the paper.

The author declares no competing interest.

This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

See companion article, "The rise and fall of rationality in language," [10.1073/pnas.2107848118](https://doi.org/10.1073/pnas.2107848118).

¹Email: sdedeo@andrew.cmu.edu.

Published January 18, 2022.

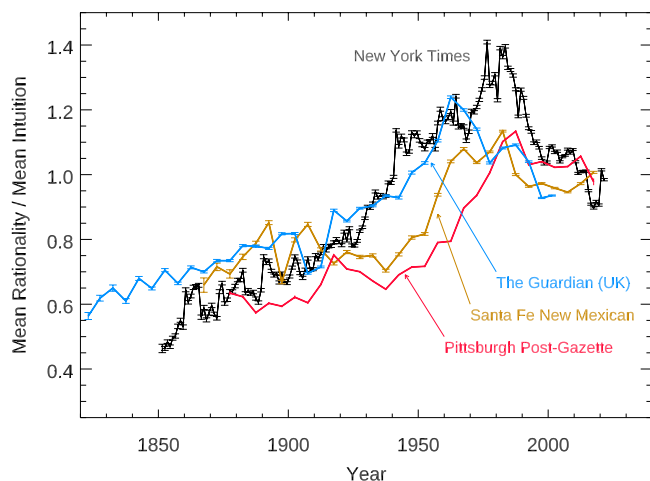


Fig. 1. The trajectory of Scheffer et al.'s (1) intuitive/system I and rational/system II word lists in four newspapers.

Happily, corroborating evidence comes from a separate analysis the authors conduct on *The New York Times* archive (1). While the effect sizes are smaller (a shift of 50%, rather than a factor of 3), the qualitative rise and fall structure persists. In Fig. 1 of this commentary, I replicate their analysis, plotting the relative average frequency of their system I intuition words (44 words in the belief and sense categories) and their system II rationality words (22 words in the science and quantification categories). The solid black line in Fig. 1 matches, as it should, their figure 3A with only minor differences (see below).

On top of this line, I have overlaid new curves for three other newspapers, built from a large-scale dataset available at othernewspapers.com: *The Guardian*, a general circulation newspaper in the United Kingdom that plays a role similar to *The New York Times*, and two newspapers (chosen semirandomly) serving a more local audience, the *Pittsburgh Post-Gazette* and the *Santa Fe New Mexican*. I used the same word lists as Scheffer et al. (1) but dropped the terms “unit,” “model,” and “size” because they appeared to be strongly associated with advertising copy, and I dropped the word “hunch,” as it appeared to generate a large number of false hits due to errors in optical character recognition.

All four datasets show evidence for the Weberian rise, but evidence for a coherent “post-truth” fall is weaker. *The New York Times* peaks around 1980, in rough synchrony with Google Books, but in *The Guardian* peak rationality is in the 1950s, 30 y earlier. Meanwhile, rationality peaks in both the *Pittsburgh Post-Gazette* and the *Santa Fe New Mexican* around the 1980s, but the decline is weaker, and none show the post-2000s acceleration required for the hypothesis of a social media–driven post-truth era.

Taken on their own, these results complicate the rise and fall hypothesis that the authors propose (1). When it comes to the rise portion, the fact that the local newspapers take longer to rise provides an intriguing elaboration of Weber’s tightening iron cage, suggesting that the process took time to propagate beyond the elite, managerial class readership of *The New York Times*. On March 1, 1950, a year when the gap is largest, the front page of *The New York Times* focused on exchange rates, international political strategy, and a corruption investigation, while the *Pittsburgh Post-Gazette* included a religious essay on the beginning of Lent, a human interest story (“Death ends career of dopey leopard”), and a local murder. In Pittsburgh, the cage had yet to tighten.

The lack of a fall is harder to square. One might say that the fall never happened in Pittsburgh or Santa Fe because rationality never reached the exalted heights of *The New York Times* or *The Guardian* that it needed to fall from. Or that it did, because readers in those smaller cities turned to *The New York Times* as they joined the managerial elite. Or, perhaps, that the further rise in London and New York was illusory, and driven by the emergence of mass market tabloids that took up the “intuitive” slack. More broadly, we are drawn to questions about the shifting print ecosystem that a reader inhabited. None of these archives were designed to sample the human experience in a uniform fashion.

This is where Google Books comes in: More comprehensive by far, it might be a better representation of what a reader encountered, particularly in a postwar era when a greater variety of books were available to read at lower relative cost. Even here, however, we ought to be cautious. Google Books is weighted by book, not readership, and, because readership is power-law distributed over books, the average book is not representative of the average reader’s experience. Worse, we expect the books that garner the most attention to have distinctive cognitive properties. This gap between the kinds of books a reader is likely to encounter and the uniform sample provided the Google Books corpus is part of Schmidt et al.’s (9) recent criticism.

In 1993, for example, sentences in Google Books beginning with the word “Derrida” were roughly half as common as those beginning with “Depression,” but this likely radically overestimates the relative consumption of (on the one hand) sentences about the poststructuralist philosopher and (on the other) the kinds of sentences that appeared in self-help books on depression, such as that year’s bestselling *Listening to Prozac*. If, as seems likely, the former are more likely to contain rationality words, the emergence of a new academic industry could lead to significant shifts in the rationality-to-intuition ratio, even if its books are only read by a tiny fraction of the public. I emphasize the word “could”: The fact is that nobody has systematically studied the effect of the book–reader distinction in the Google Books corpus, and it remains one of its many unknowns.

Such questions are familiar to social scientists and historians, who often work at the level of the institution rather than the individual. Consider, for example, recent work by Soni et al. (10) on the 19th century abolitionist movement. Like Charlesworth et al. (2), they use word embeddings to study collective representations, but they do so in order to contrast the representations that different newspapers produced. This simplifies the problem: While Soni et al. need only make sure their corpus is complete (i.e., contains as many of the articles from each newspaper as it can), Charlesworth et al. need to make the additional case that their corpora, have the right properties to calibrate an (implicit) “consensus” representation.

From Word to Thought

The second half of the question remains: To what extent do these word shifts reflect shifts in psychological or social reality? Did Westerners really become more, then less, rational?

Rationality is a cognitive process, something we can study in the laboratory with the cognitive reflection task, but it is also a social process, a method of coming to agreement—or at least justifying an agreement made by other means—through the use of language. To be rational is, in part, to accept a mind-to-world fit for beliefs (11) by deliberately committing oneself to conceptual constraints (12). The two go together, and public

use of reason in argument making appears to be a key part of triggering system II thinking in things like the Wason selection task (13).

It is the public aspect that Scheffer et al. (1) can hope to track most directly. It is natural, for example, to think that the use of abstract definitions and logical connectives is a sign of rational constraints, that it is difficult to talk about “chemicals” in an intuitive fashion, or that a preponderance of “We” over “I” might correspond to a shift from beliefs of the individually situated “I” to those of a collectively justified “We.”

The claim, however, sits somewhat uneasily with a diagnosis in terms of a Weberian bureaucracy. The architects of the Vietnam War, for example, were eager to demonstrate, in highly scientific and logical language, how rational and unsentimental they were. As Hannah Arendt pointed out in her discussion of the “Pentagon Papers” (14), however, the underlying thinking was a matter of fantasy—of making reality fit what they intuitively wanted to be true, and inventing contradictory justifications for actions they had already decided to take.

In part because of concerns like this, Scheffer et al. (1) are careful to talk about shifts of “public interest” in rationality, rather than rationality itself. This is a third thing yet: not a psychological phenomenon, nor its social manifestation, but a matter of style and the relative prestige of different dialects. In the extreme version, the decision to say “conclude” rather than “imagine” is just the difference between “beef” and “cow”: two words for the same thing.

Similar concerns apply to intuition. Scheffer et al. (1) suggest that the shift toward intuition words is associated with an individualistic turn away from a neoliberal rationalized bureaucracy. It is possible, however, that some of this shift is simply a new phase of the same system. One is reminded of the 1999 movie *Office*

Space, where the heroine is reprimanded for failing to wear the required “pieces of flair” (whimsical, “spiritual” buttons that are actually meaningless tokens of company culture), or the corporate manuals that suggest people use “I” talk to help defuse conflicts for better group cohesion, or the early-2000s “Army of One” recruitment campaign.

This raises the question, What fraction of the shift is driven by the merging and shifting of the meanings of words, rather than a true psychological rearrangement? More provocatively, What are the limits to how social processes can redefine basic lexical structures? [One useful feature of a word’s meaning is its ability to help us coordinate (15). If words are in danger of shifting their meaning, we might still be able to identify a common task stable enough to provide a standard measure for different words they’re associated with across time.]

Given the magnitude of Scheffer et al.’s (1) effect sizes, it is important not to overstate the concern. Only in truly pathological cases—the “Pentagon Papers” perhaps being one—should we expect the lexicon to completely decouple from psychological reality. In more ordinary circumstances, even when we ourselves use words in ways that clash with our inner experiences, we can never be sure that others are doing so as well. This means that, at the very least, a shift toward intuition words creates some level of initial uncertainty about whether or not we ought to be acting on intuition rather than reason.

Words matter, and one reason to trust the basics of Scheffer et al.’s (1) methods is that there are limits to society’s ability to redefine meanings. Only in special circumstances, corresponding to an advanced state of semantic decay, can “I imagine” truly mean, for both speaker and listener, “we conclude.” One hopes we are not there yet. That we are pushed to ask such questions, however, is part of this paper’s generative power.

- 1 M. Scheffer, I. van de Leemput, E. Weinans, J. Bollen, The rise and fall of rationality in language. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2107848118 (2021).
- 2 T. E. S. Charlesworth, V. Yang, T. C. Mann, B. Kurdi, M. R. Banaji, Gender stereotypes in natural language: Word embeddings show robust consistency across child and adult language corpora of more than 65 million words. *Psychol. Sci.* **32**, 218–240 (2021).
- 3 R. Iliev, R. Axelrod, Does causality matter more now? Increase in the proportion of causal language in English texts. *Psychol. Sci.* **27**, 635–643 (2016).
- 4 P. S. Dodds et al., Human language reveals a universal positivity bias. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 2389–2394 (2015).
- 5 J. C. Jackson et al., From text to thought: How analyzing language can advance psychological science. *Perspect. Psychol. Sci.*, 10.117/17456916211004899 (2021).
- 6 J. B. Michel et al.; Google Books Team, Quantitative analysis of culture using millions of digitized books. *Science* **331**, 176–182 (2011).
- 7 M. Weber, *The Protestant Ethic and the Spirit of Capitalism: And Other Writings* (Penguin, 2002).
- 8 E. A. Pechenick, C. M. Danforth, P. S. Dodds, Characterizing the Google Books corpus: Strong limits to inferences of socio-cultural and linguistic evolution. *PLoS One* **10**, e0137041 (2015).
- 9 B. Schmidt, S. T. Piantadosi, K. Mahowald, Uncontrolled corpus composition drives an apparent surge in cognitive distortions. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2115010118 (2021).
- 10 S. Soni, L. F. Klein, J. Eisenstein, Abolitionist networks: Modeling language change in nineteenth-century activist newspapers. *J. Cult. Anal.* **6**, 1–43 (2021).
- 11 J. Searle, *Rationality in Action* (Bradford, 2001).
- 12 R. Brandom, *A Spirit of Trust: A Reading of Hegel’s Phenomenology* (Harvard University Press, 2019).
- 13 H. Mercier, D. Sperber, Why do humans reason? Arguments for an argumentative theory. *Behav. Brain Sci.* **34**, 57–74. Discussion in: *Behav. Brain Sci.* **34**, 74–111 (2011).
- 14 H. Arendt, Lying in Politics: Reflections on the Pentagon Papers. *New York Rev. Books* **XVII**, 30–39 (1971).
- 15 D. Lewis, *Convention: A Philosophical Study* (Wiley, 2013).