

Human language reveals a universal positivity bias

Peter Sheridan Dodds^{a,b,1}, Eric M. Clark^{a,b}, Suma Desu^c, Morgan R. Frank^c, Andrew J. Reagan^{a,b}, Jake Ryland Williams^{a,b}, Lewis Mitchell^d, Kameron Decker Harris^e, Isabel M. Kloumann^f, James P. Bagrow^{a,b}, Karine Megerdooian^g, Matthew T. McMahon^g, Brian F. Tivnan^{b,g,1}, and Christopher M. Danforth^{a,b,1}

^aComputational Story Lab, Vermont Advanced Computing Core, and the Department of Mathematics and Statistics, University of Vermont, Burlington, VT 05401; ^bVermont Complex Systems Center, University of Vermont, Burlington, VT 05401; ^cCenter for Computational Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139; ^dSchool of Mathematical Sciences, The University of Adelaide, SA 5005, Australia; ^eApplied Mathematics, University of Washington, Seattle, WA 98195; ^fCenter for Applied Mathematics, Cornell University, Ithaca, NY 14853; and ^gThe MITRE Corporation, McLean, VA 22102

Edited by Kenneth W. Wachter, University of California, Berkeley, CA, and approved January 9, 2015 (received for review June 23, 2014)

Using human evaluation of 100,000 words spread across 24 corpora in 10 languages diverse in origin and culture, we present evidence of a deep imprint of human sociality in language, observing that (i) the words of natural human language possess a universal positivity bias, (ii) the estimated emotional content of words is consistent between languages under translation, and (iii) this positivity bias is strongly independent of frequency of word use. Alongside these general regularities, we describe interlanguage variations in the emotional spectrum of languages that allow us to rank corpora. We also show how our word evaluations can be used to construct physical-like instruments for both real-time and offline measurement of the emotional content of large-scale texts.

language | social psychology | happiness | positivity

Human language, our great social technology, reflects that which it describes through the stories it allows to be told and us, the tellers of those stories. Although language's shaping effect on thinking has long been controversial (1–3), we know that a rich array of metaphor encodes our conceptualizations (4), word choice reflects our internal motives and immediate social roles (5–7), and the way a language represents the present and future may condition economic choices (8).

In 1969, Boucher and Osgood (9) framed the Pollyanna hypothesis: a hypothetical, universal positivity bias in human communication. From a selection of small-scale, cross-cultural studies, they marshaled evidence that positive words are likely more prevalent, more meaningful, more diversely used, and more readily learned. However, in being far from an exhaustive, data-driven analysis of language, which is the approach we take here, their findings could only be regarded as suggestive. Indeed, studies of the positivity of isolated words and word stems have produced conflicting results, some pointing toward a positivity bias (10) and others toward the opposite (11, 12), although attempts to adjust for use frequency tend to recover a positivity signal (13).

Materials and Methods

To explore the positivity of human language deeply, we constructed 24 corpora spread across 10 languages. Our global coverage of linguistically and culturally diverse languages includes English, Spanish, French, German, Brazilian Portuguese, Korean, Chinese (Simplified), Russian, Indonesian, and Arabic. The sources of our corpora are similarly broad, spanning books (14), news outlets, social media, the web (15), television and movie subtitles, and music lyrics (16). Our work here greatly expands upon our earlier study of English alone, where we found strong evidence for a use-invariant positivity bias (17). In *SI Appendix*, we provide full details of our corpora (*SI Appendix, Table S1*), survey, and participants (*SI Appendix, Table S2*).

We address the social nature of language in two important ways: (i) We focus on the words people most commonly use, and (ii) we measure how those same words are received by individuals. We take word use frequency as the primary organizing measure of a word's importance. Such a data-driven approach is crucial for both understanding the structure of language and creating linguistic instruments for principled measurements (18, 19). By contrast, earlier studies focusing on meaning and emotion have used "expert" generated word lists, and these word lists fail statistically to match frequency distributions of natural language (10–12, 20), confounding attempts to make

claims about language in general. For each of our corpora, we selected between 5,000 and 10,000 of the most frequently used words, choosing the exact numbers so that we obtained ~10,000 words for each language.

Of our 24 corpora, we received 17 already parsed into words by the source: the Google Books Project (six corpora), the Google Web Crawl (eight corpora), and movie and television subtitles (three corpora). For the other seven corpora (five Twitter corpora, the *New York Times*, and music lyrics), we extracted words by standard white space separation. Twitter was easily the most variable and complex of our text sources, and required additional treatment. In parsing Twitter, we required strings to contain at least one Unicode character and no invisible control characters, and we excluded strings representing web links, bearing a leading @, ampersand (&), or other punctuation (e.g., Twitter IDs) but kept hashtags. Finally, for all corpora, we converted words to lowercase. We observed that common English words appeared in the Twitter corpora of other languages, and we have chosen simply to acknowledge this reality of language and allow these commonly used borrowed words to be evaluated.

Although there are many complications with inflections and variable orthography, we have found merit for our broad analysis in not collapsing related words. For example, we have observed that allowing different conjugations of verbs to stand in our corpora is valuable because human evaluations of such have proved to be distinguishable [e.g., present vs. past tense (18)]. As should be expected, a more nuanced treatment going beyond the present paper's bounds by involving stemming and word type, for example, may lead to minor corrections (21), although our central observations will remain robust and will in no way change the behavior of the instruments we generate.

There is no single, principled way to merge corpora to create an ordered list of words for a given language. For example, it is impossible to weight the most commonly used words in the *New York Times* against the most commonly used words in Twitter. Nevertheless, we are obliged to choose some method for doing so to facilitate comparisons across languages and for the purposes of building adaptable linguistic instruments. For each language where we had more than one corpus, we created a single quasi-ranked word list by finding the smallest integer r such that the union of all words with a rank $\leq r$ in at least one corpus formed a set of at least 10,000 words.

Significance

The most commonly used words of 24 corpora across 10 diverse human languages exhibit a clear positive bias, a big data confirmation of the Pollyanna hypothesis. The study's findings are based on 5 million individual human scores and pave the way for the development of powerful language-based tools for measuring emotion.

Author contributions: P.S.D., B.F.T., and C.M.D. designed research; P.S.D., E.M.C., S.D., M.R.F., A.J.R., J.R.W., L.M., K.D.H., I.M.K., J.P.B., K.M., M.T.M., and C.M.D. performed research; P.S.D., E.M.C., A.J.R., K.M., and M.T.M. contributed new reagents/analytic tools; P.S.D., E.M.C., S.D., M.R.F., A.J.R., J.R.W., J.P.B., and C.M.D. analyzed data; and P.S.D. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

Data deposition: Data are available in *Dataset S1* and at www.uvm.edu/storylab/share/papers/dodds2014a/index.html.

¹To whom correspondence may be addressed. Email: peter.dodds@uvm.edu, btivnan@mitre.org, or chris.danforth@uvm.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1411678112/-DCSupplemental.

We then paid native speakers to rate how they felt in response to individual words on a nine-point scale, with 1 corresponding to most negative or saddest, 5 to neutral, and 9 to most positive or happiest (10, 18) (*SI Appendix*). This happy-sad semantic differential (20) functions as a coupling of two standard five-point Likert scales. Participants were restricted to certain regions or countries (e.g., Portuguese was rated by residents of Brazil). Overall, we collected 50 ratings per word for a total of around 5 million individual human assessments. We provide all datasets as part of *SI Appendix*.

Results and Discussion

In Fig. 1, we show distributions of the average happiness scores for all 24 corpora, leading to our most general observation of a clear positivity bias in natural language. We indicate the above-neutral part of each distribution with yellow and the below-neutral part with blue, and we order the distributions moving upward by increasing median (vertical red line). For all corpora, the median clearly exceeds the neutral score of 5. The background gray lines connect deciles for each distribution. In *SI Appendix*, Fig. S1, we provide the same distributions ordered instead by increasing variance.

As is evident from the ordering in Fig. 1 and *SI Appendix*, Fig. S1, although a positivity bias is the universal rule, there are minor differences between the happiness distributions of languages. For example, Latin American-evaluated corpora (Mexican Spanish and Brazilian Portuguese) exhibit relatively high medians and, to a lesser degree, higher variances. For other languages, we see that those languages with multiple corpora have more variable medians, and specific corpora are not ordered by median in the same way across languages (e.g., Google Books has a lower median than Twitter for Russian, but the reverse is true for German and English). In terms of emotional variance, all four English corpora are among the highest, whereas Chinese and Russian Google Books seem especially constrained.

We now examine how individual words themselves vary in their average happiness score between languages. Owing to the scale of our corpora, we were compelled to use an online service, choosing Google Translate. For each of the 45 language pairs, we translated isolated words from one language to the other and then back. We then found all word pairs that (*i*) were translationally stable, meaning the forward and back translation returns the original word, and (*ii*) appeared in our corpora for each language.

We provide the resulting comparison between languages at the level of individual words in Fig. 2. We use the mean of each language's word happiness distribution derived from its merged corpora to generate a rough overall ordering, acknowledging that frequency of use is no longer meaningful, and moreover is not relevant, because we are now investigating the properties of individual words. Each cell shows a heat map comparison with word density increasing as shading moves from gray to white. The background colors reflect the ordering of each pair of languages: yellow if the row language had a higher average happiness than the column language and blue for the reverse. Also, in each cell, we display the number of translation-stable words between language pairs, N , along with the difference in average word happiness, Δ , where each word is equally weighted.

A linear relationship is clear for each language-language comparison, and is supported by Pearson's correlation coefficient r being in the range 0.73–0.89 ($P < 10^{-118}$ across all pairs; Fig. 2 and *SI Appendix*, Tables S3–S5). Overall, this strong agreement between languages suggests that approximate estimates of word happiness for unscored languages could be generated with no expense from our existing dataset. Some words will, of course, translate unsatisfactorily, with the dominant meaning changing between languages. For example “lying” in English, most readily interpreted as speaking falsehoods by our participants, translates to “acostado” in Spanish, meaning recumbent. Nevertheless, happiness scores obtained by translation will be serviceable for purposes where the effects of many different words are incorporated.




Fig. 1. Distributions of perceived average word happiness h_{avg} for 24 corpora in 10 languages. The histograms represent the 5,000 most commonly used words in each corpora (details are provided in *SI Appendix*), and native speakers scored words on a 1–9 double-Likert scale, with 1 being extremely negative, 5 being neutral, and 9 being extremely positive. Yellow indicates positivity ($h_{\text{avg}} > 5$), and blue indicates negativity ($h_{\text{avg}} < 5$). Distributions are ordered by increasing median (red vertical line). The background gray lines connect deciles of adjacent distributions. The same distributions arranged according to increasing variance are shown in *SI Appendix*, Fig. S1.




Fig. 2. Scatter plots of average happiness for words measured in different languages. We order languages from relatively most positive (Spanish) to relatively least positive (Chinese); a yellow background indicates the row language is more positive than the column language, and a blue background indicates the converse. The overall plot matrix is symmetrical about the leading diagonal, with the redundancy allowing for easier comparison between languages. In each scatter plot, the key gives the number of translation-stable words for each language pair, N ; the average difference in translation-stable word happiness between the row language and the column language, Δ ; and the Pearson correlation coefficient for the regression, r . All P values are less than 10^{-118} for the Pearson correlation coefficient and less than 10^{-82} for the Spearman correlation coefficient.

Stepping back from examining interlanguage robustness, we return to a more detailed exploration of the rich structure of each corpus's happiness distribution. In Fig. 3, we show how average word happiness h_{avg} is largely independent of word use frequency for four example corpora. We first plot use frequency rank r of the 5,000 most frequently used words as a function of their average happiness score, h_{avg} (background dots), along with some example evenly spaced words. We note that words at the extremes of the happiness scale are ones evaluators agreed upon

strongly, whereas words near neutral range from being clearly neutral [e.g., h_{avg} ("the") = 4.98] to contentious with high SD (17). We then compute deciles for contiguous sets of 500 words, sliding this window through rank r . These deciles form the vertical strands. We overlay randomly chosen, equally spaced example words to give a sense of each corpus's emotional texture.

We chose the four example corpora shown in Fig. 3 to be disparate in nature, covering diverse languages (French, Egyptian Arabic, Brazilian Portuguese, and Chinese), regions of the




Fig. 3. Examples of how word happiness varies little with use frequency. (A–D) Above each plot is a histogram of average happiness h_{avg} for the 5,000 most frequently used words in the given corpus, matching Fig. 1. Each point locates a word by its usage rank r and average happiness h_{avg} , and we show some regularly spaced example words. The descending gray curves of these jellyfish plots indicate deciles for windows of 500 words of contiguous use rank, showing that the overall histogram's form is roughly maintained at all scales. The “kkkkkk...” words represent laughter in Brazilian Portuguese, in the manner of “hahaha...,” and an English translation is provided in *SI Appendix, Fig. S2*.

world (Europe, the Middle East, South America, and Asia), and texts [Twitter, movies and television, the web (15), and books (14)]. The remaining 20 corpora yield similar plots (*SI Appendix*), and all corpora also exhibit an approximate self-similarity in SD for word happiness.

Across all corpora, we observe visually that the deciles tend to stay fixed or move slightly toward the negative, with some expected fragility at the 10% and 90% levels (due to the distributions' tails), indicating that the overall happiness distribution of each corpus approximately holds independent of word use. In Fig. 3, for example, we see that both the Brazilian Portuguese and French examples show a small shift to the negative for increasingly rare words, whereas there is no visually clear trend for the Arabic and Chinese cases. Fitting $h_{\text{avg}} = \alpha r + \beta$ typically returns α on the order of -1×10^{-5}

suggesting h_{avg} decreases 0.1 per 10,000 words. For SDs of happiness scores, we find a similarly weak drift toward higher values for increasingly rare words (correlations and linear fits for h_{avg} and h_{std} as a function of word rank r for all corpora are provided in *SI Appendix, Tables S6 and S7*). We thus find that, to first order, not just the positivity bias but the happiness distribution itself applies for common words and rare words alike, revealing an unexpected addition to the many well-known scalings found in natural language, as famously exemplified by Zipf's law (22).

In constructing language-based instruments for measuring expressed happiness, such as our hedonometer (18), this frequency independence allows for a way to “increase the gain” in a fashion resembling standard physical instruments. Moreover, we have earlier demonstrated the robustness of our hedonometer for the

English language, showing, for example, that measurements derived from Twitter correlate strongly with Gallup well-being polls and related indices at the state and city level for the United States (19).

Here, we provide an illustrative use of our hedonometer in the realm of literature, inspired by Vonnegut's shapes of stories (23, 24). In Fig. 4, we show "happiness time series" for three famous works of literature, evaluated in their original languages of English, Russian, and French, respectively: Melville's *Moby Dick* (www.gutenberg.org), Dostoyevsky's *Crime and Punishment* (25), and Dumas' *The Count of Monte Cristo* (www.gutenberg.org). We

slide a 10,000-word window through each work, computing the average happiness using a "lens" for the hedonometer in the following manner. We capitalize on our instrument's tunability to obtain a strong signal by excluding all words for which $3 < h_{\text{avg}} < 7$ (i.e., we keep words residing in the tails of each distribution) (18). Denoting a given lens by its corresponding set of allowed words L , we estimate the happiness score of any text T as $h_{\text{avg}}(T) = \sum_{w \in L} f_w h_{\text{avg}}(w) / \sum_{w \in L} f_w$, where f_w is the frequency of word w in T (16).

The three resulting happiness time series provide interesting detailed views of each work's narrative trajectory revealing

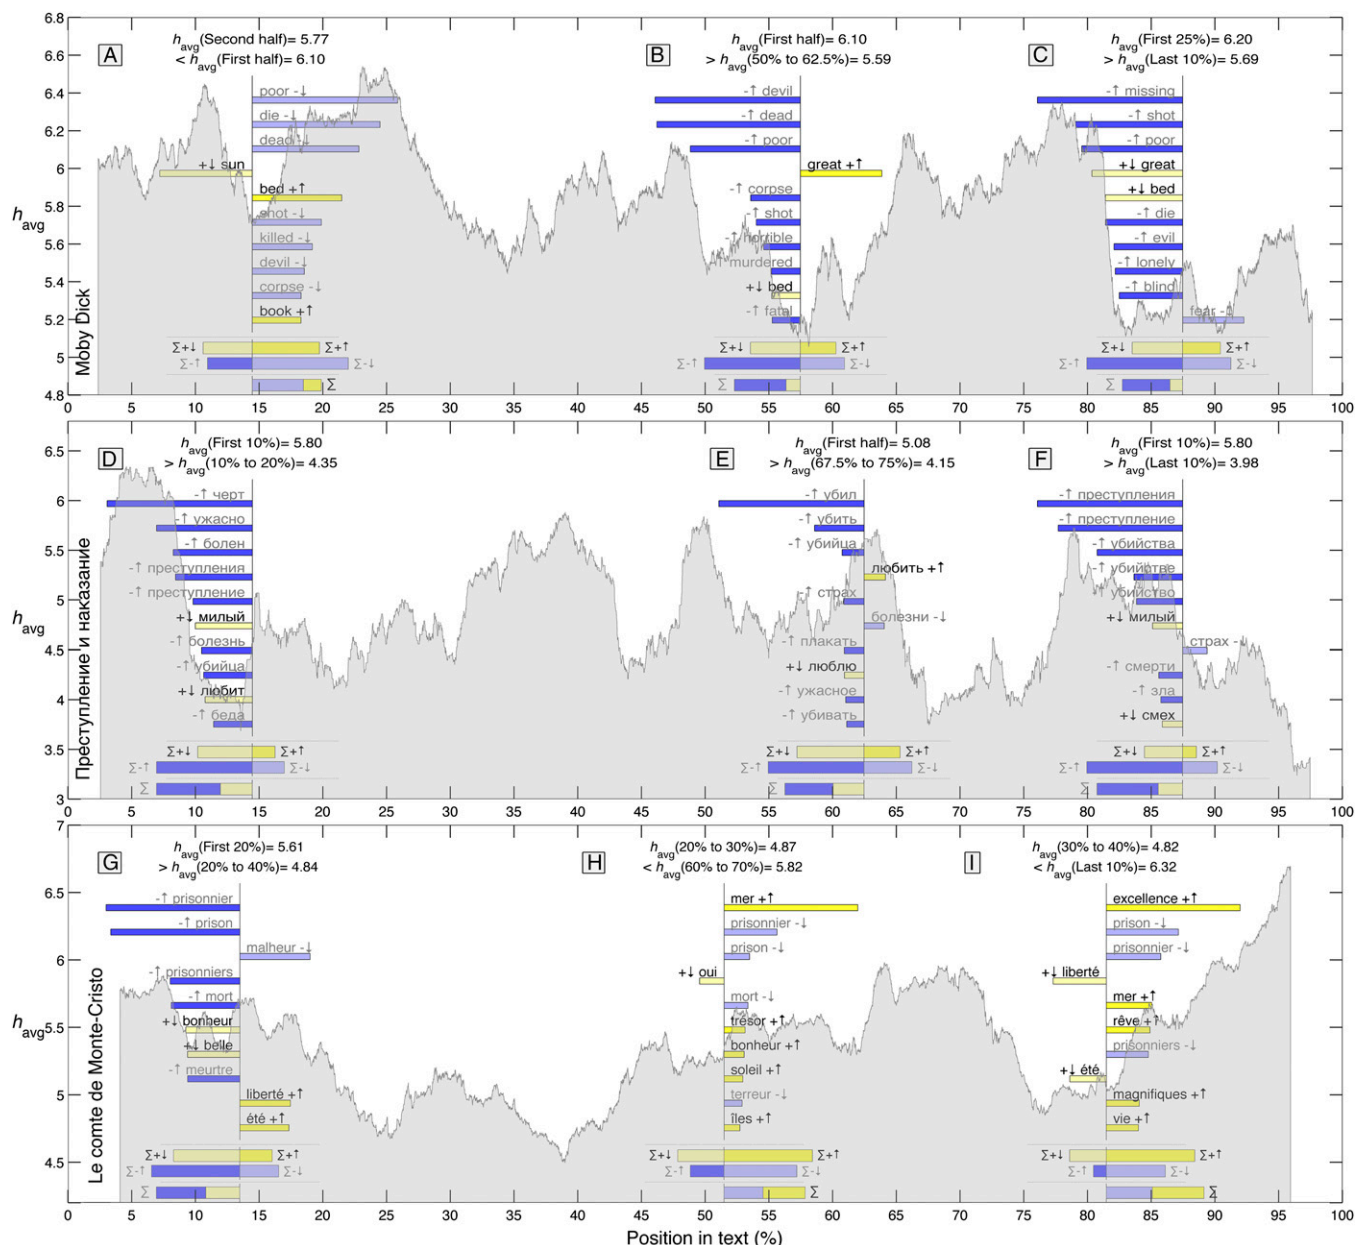


Fig. 4. Emotional time series for three famous 19th century works of literature: Melville's *Moby Dick* (Top), Dostoyevsky's *Crime and Punishment* (Middle), and Dumas' *The Count of Monte Cristo* (Bottom). Each point represents the language-specific happiness score for a window of 10,000 words (converted to lowercase), with the window translated throughout the work. The overlaid word shifts (A–I) show example comparisons between different sections of each work. Word shifts indicate which words contribute the most toward and against the change in average happiness between two texts (*SI Appendix*, pp. S9–S10). Although a robust instrument in general, we acknowledge the hedonometer's potential failure for individual words due to both language evolution and words possessing more than one meaning. For *Moby Dick*, we excluded "cried" and "cry" (to speak loudly rather than weep) and "Coffin" (surname, still common in Nantucket). Such alterations, which can be done on a case-by-case basis, do not noticeably change the overall happiness curves, although leaving the word shifts more informative. We provide online, interactive visualizations of the emotional trajectories of over 10,000 books at hedonometer.org/books.html.

numerous peaks and troughs throughout, at times clearly dropping below neutral. Both *Moby Dick* and *Crime and Punishment* end on low notes, whereas *The Count of Monte Cristo* culminates with a rise in positivity, accurately reflecting the finishing arcs of all three. The “word shifts” overlaying the time series compare two distinct regions of each work, showing how changes in word abundances lead to overall shifts in average happiness. Such word shifts are essential tests of any sentiment measurement, and are made possible by the linear form of our instrument (16, 18) [a full explanation is provided in *SI Appendix*, pp. S9–S10]. As one example, the third word shift for *Moby Dick* shows why the average happiness of the last 10% of the book is well below the average happiness of the first 25%. The major contribution is an increase in relatively negative words, including “missing,” “shot,” “poor,” “die,” and “evil.”

By adjusting the lens, many other related time series can be formed, such as those produced by focusing on only positive or negative words. Emotional variance as a function of text position can also be readily extracted. At hedonometer.org/books.html, we provide online, interactive emotional trajectories for over 10,000 works of literature where different lenses and regions of comparisons may be easily explored. Beyond this example tool we have created here for the digital humanities and our

hedonometer for measuring population well-being, the datasets we have generated for the present study may be useful in creating a great variety of language-based instruments for assessing emotional expression.

Overall, our major scientific finding is that when experienced in isolation and weighted properly according to use, words, which are the atoms of human language, present an emotional spectrum with a universal, self-similar positive bias. We emphasize that this apparent linguistic encoding of our social nature is a system-level property, and in no way asserts all natural texts will skew positive (as exemplified by certain passages of the three works in Fig. 4) or diminishes the salience of negative states (26). Going forward, our word happiness assessments should be periodically repeated and carried out for new languages, tested on different demographics, and expanded to phrases both for the improvement of hedonometric instruments and to chart the dynamics of our collective social self.

ACKNOWLEDGMENTS. We thank M. Shields, K. Comer, N. Berry, I. Ramiscal, C. Burke, P. Carrigan, M. Koehler, and Z. Henscheid, in part, for their roles in developing hedonometer.org. We also thank F. Henegan, A. Powers, and N. Atkins for conversations. P.S.D. was supported by National Science Foundation CAREER Award 0846668.

- Whorf BL (1956) *Language, Thought, and Reality: Selected Writings of Benjamin Lee Whorf*, ed Carroll JB (MIT Press, Cambridge, MA).
- Chomsky N (1957) *Syntactic Structures* (Mouton, The Hague).
- Steven Pinker (1994) *The Language Instinct: How the Mind Creates Language* (William Morrow and Company, New York).
- Lakoff G, Johnson M (1980) *Metaphors We Live By* (Univ of Chicago Press, Chicago).
- Campbell RS, Pennebaker JW (2003) The secret life of pronouns: Flexibility in writing style and physical health. *Psychol Sci* 14(1):60–65.
- Newman MEJ (2003) The structure and function of complex networks. *SIAM Review* 45(2):167–256.
- Pennebaker JW (2011) *The Secret Life of Pronouns: What Our Words Say About Us* (Bloomsbury Press, New York).
- Chen MK (2013) The effect of language on economic behavior: Evidence from savings rates, health behaviors, and retirement assets. *Am Econ Rev* 103(2):690–731.
- Boucher J, Osgood CE (1969) The Pollyanna hypothesis. *J Verbal Learning Verbal Behav* 8(1):1–8.
- Bradley MM, Lang PJ (1999) *Affective Norms for English Words (Anew): Stimuli, Instruction Manual and Affective Ratings. Technical Report C-1* (University of Florida, Gainesville, FL).
- Stone PJ, Dunphy DC, Smith MS, Ogilvie DM (1966) *The General Inquirer: A Computer Approach to Content Analysis* (MIT Press, Cambridge, MA).
- Pennebaker JW, Booth RJ, Francis ME (2007) *Linguistic Inquiry and Word Count: Liwc 2007*. Available at bit.ly/S1Dk2L. Accessed May 15, 2014.
- Jurafsky D, Chahuneau V, Routledge BR, Smith NA (2014) Narrative framing of consumer sentiment in online restaurant reviews. *First Monday*, 10.5210/fm.v19i4.4944.
- Michel J-B, et al. (2011) Quantitative analysis of culture using millions of digitized books. *Science* 331:176–182.
- Brants T, Franz A (2006) Google Web 1T 5-gram (Linguistic Data Consortium, Philadelphia). Version 1.
- Dodds PS, Danforth CM (2009) Measuring the happiness of large-scale written expression: Songs, blogs, and presidents. *J Happiness Stud* 11(4):441–456.
- Kloumann IM, Danforth CM, Harris KD, Bliss CA, Dodds PS (2012) Positivity of the English language. *PLoS ONE* 7(1):e29484.
- Dodds PS, Harris KD, Kloumann IM, Bliss CA, Danforth CM (2011) Temporal patterns of happiness and information in a global social network: Hedonometrics and Twitter. *PLoS ONE* 6(12):e26752.
- Mitchell L, Frank MR, Harris KD, Dodds PS, Danforth CM (2013) The geography of happiness: Connecting twitter sentiment and expression, demographics, and objective characteristics of place. *PLoS ONE* 8(5):e64417.
- Osgood C, Suci G, Tannenbaum P (1957) *The Measurement of Meaning* (Univ of Illinois, Urbana, IL).
- Warriner AB, Kuperman V (2014) Affective biases in English are bi-dimensional. *Cogn Emotion*, 10.1080/02699931.2014.968098.
- Zipf GK (1949) *Human Behaviour and the Principle of Least-Effort* (Addison-Wesley, Cambridge, MA).
- Vonnegut K, Jr (2005) *A Man Without a Country* (Seven Stories Press, New York).
- Vonnegut K (2010) Kurt Vonnegut on the shapes of stories. Available at www.youtube.com/watch?v=oP3c1h8v2ZQ. Accessed May 15, 2014.
- Dostoyevsky F (1866) *Crime and Punishment*. Original Russian text. Available at ilibrary.ru/text/69/p.1/index.html. Accessed December 15, 2013.
- Forgas JP (2013) Don't worry, be sad! On the cognitive, motivational, and interpersonal benefits of negative mood. *Curr Dir Psychol Sci* 22(3):225–232.

SUPPLEMENTARY INFORMATION

Online, interactive visualizations:

We provide the following resources online at <http://compstorylab.org/share/papers/dodds2014a/> and at <http://hedonometer.org>.

- Links to example scripts for parsing and measuring average happiness scores for texts, and for generating D3 word shifts: <http://compstorylab.org/share/papers/dodds2014a/code.html>
- Visualizations for exploring translation-stable word pairs across languages;
- Interactive, multi-language time series for over 10,000 works of literature including Moby Dick and Harry Potter;
- Jellyfish plots for all 24 corpora;
- An API for selected data streams and word lists at <http://hedonometer.org/api.html>
- Spatiotemporal hedonometric measurements of Twitter across all 10 languages, explorable at <http://hedonometer.org>.

Corpora

We used the services of Appen Butler Hill (<http://www.appen.com>) for all word evaluations excluding English, for which we had earlier employed Mechanical Turk (<https://www.mturk.com/> [17]).

English instructions (see the example below) were translated to all other languages and given to participants along with survey questions, and an example of the English instruction page is below. Non-english language experiments were conducted through a custom interactive website built by Appen Butler Hill, and all participants were required to pass a stringent aural proficiency test in their own language.


Measuring the Happiness of Words

Our overall aim is to assess how people feel about individual words. With this particular survey, we are focusing on the dual emotions of sadness and happiness. You are to rate 100 individual words on a 9 point unhappy-happy scale.

Please consider each word carefully. If we determine that your ratings are randomly or otherwise inappropriately selected, or that any questions are left unanswered, we may not approve your work. These words were chosen based on their common usage. As a result, a small portion of words may be offensive to some people, written in a different language, or nonsensical.

Before completing the word ratings, we ask that you answer a few short demographic questions. We expect the entire survey to require 10 minutes of your time. Thank you for participating!

Example:



Read the word and click on the face that best corresponds to your emotional response.

Demographic Questions

1. What is your gender? (Male/Female)
2. What is your age? (Free text)
3. Which of the following best describes your highest achieved education level?
Some High School, High School Graduate, Some college, no degree, Associates degree, Bachelors degree, Graduate degree (Masters, Doctorate, etc.)
4. What is the total income of your household?
5. Where are you from originally?
6. Where do you live currently?
7. Is _____ your first language? (Yes/No) If it is not, please specify what your first language is.
8. Do you have any comments or suggestions? (Free text)

Corpus:	# Words	Reference(s)
English: Twitter	5000	[18]
English: Google Books Project	5000	[14]
English: The New York Times	5000	[18]
English: Music lyrics	5000	[16]
Portuguese: Google Web Crawl	7133	[15]
Portuguese: Twitter	7119	Twitter API
Spanish: Google Web Crawl	7189	[15]
Spanish: Twitter	6415	Twitter API
Spanish: Google Books Project	6379	[14]
French: Google Web Crawl	7056	[15]
French: Twitter	6569	Twitter API
French: Google Books Project	6192	[14]
Arabic: Movie and TV subtitles	9999	The MITRE Corporation
Indonesian: Twitter	7044	Twitter API
Indonesian: Movie subtitles	6726	The MITRE Corporation
Russian: Twitter	6575	Twitter API
Russian: Google Books Project	5980	[14]
Russian: Movie and TV subtitles	6186	[15]
German: Google Web Crawl	6902	[15]
German: Twitter	6459	Twitter API
German: Google Books Project	6097	[14]
Korean: Twitter	6728	Twitter API
Korean: Movie subtitles	5389	The MITRE Corporation
Chinese: Google Books Project	10000	[14]

TABLE S1. Sources for all corpora.

Language	Participants' location(s)	Number of participants	Average words scored
English	United States of America, India	384	1302
German	Germany	196	2551
Indonesian	Indonesia	146	3425
Russian	Russia	125	4000
Arabic	Egypt	185	2703
French	France	179	2793
Spanish	Mexico	236	2119
Portuguese	Brazil	208	2404
Simplified Chinese	China	128	3906
Korean	Korea, United States of America	109	4587

TABLE S2. Number and main country/countries of location for participants evaluating the 10,000 common words for each of the 10 languages we studied. Also recorded is the average number of words evaluated by each participant (rounded to the nearest integer). We note that each word received 50 evaluations from distinct individuals. The English word list was evaluated via Mechanical Turk for our initial study [17]. The nine languages evaluated through Appen-Butler Hill yielded a higher participation rate likely due to better pay and the organization's quality of service.




FIG. S1. The same average happiness distributions shown in Fig. 1 re-ordered by increasing variance. Yellow indicates above neutral ($h_{avg} = 5$), blue below neutral, red vertical lines mark each distribution's median, and the gray background lines connect the deciles of adjacent distributions.

	Spanish	Portuguese	English	Indonesian	French	German	Arabic	Russian	Korean	Chinese
Spanish	1.00, 0.00	1.01, 0.03	1.06, -0.07	1.22, -0.88	1.11, -0.24	1.22, -0.84	1.13, -0.22	1.31, -1.16	1.60, -2.73	1.58, -2.30
Portuguese	0.99, -0.03	1.00, 0.00	1.04, -0.03	1.22, -0.97	1.11, -0.33	1.21, -0.86	1.09, -0.08	1.26, -0.95	1.62, -2.92	1.58, -2.39
English	0.94, 0.06	0.96, 0.03	1.00, 0.00	1.13, -0.66	1.06, -0.23	1.16, -0.75	1.05, -0.10	1.21, -0.91	1.51, -2.53	1.47, -2.10
Indonesian	0.82, 0.72	0.82, 0.80	0.88, 0.58	1.00, 0.00	0.92, 0.48	0.99, 0.06	0.89, 0.71	1.02, 0.04	1.31, -1.53	1.33, -1.42
French	0.90, 0.22	0.90, 0.30	0.94, 0.22	1.09, -0.52	1.00, 0.00	1.08, -0.44	0.99, 0.12	1.12, -0.50	1.37, -1.88	1.40, -1.77
German	0.82, 0.69	0.83, 0.71	0.86, 0.65	1.01, -0.06	0.92, 0.41	1.00, 0.00	0.91, 0.61	1.07, -0.25	1.29, -1.44	1.32, -1.36
Arabic	0.88, 0.19	0.92, 0.08	0.95, 0.10	1.12, -0.80	1.01, -0.12	1.10, -0.68	1.00, 0.00	1.12, -0.63	1.40, -2.14	1.43, -2.01
Russian	0.76, 0.88	0.80, 0.75	0.83, 0.75	0.98, -0.04	0.89, 0.45	0.93, 0.24	0.89, 0.56	1.00, 0.00	1.26, -1.39	1.25, -1.05
Korean	0.62, 1.70	0.62, 1.81	0.66, 1.67	0.77, 1.17	0.73, 1.37	0.78, 1.12	0.71, 1.53	0.79, 1.10	1.00, 0.00	0.98, 0.28
Chinese	0.63, 1.46	0.63, 1.51	0.68, 1.43	0.75, 1.07	0.71, 1.26	0.76, 1.03	0.70, 1.41	0.80, 0.84	1.02, -0.29	1.00, 0.00

TABLE S3. Reduced Major Axis (RMA) regression fits for row language as a linear function of the column language: $h_{\text{avg}}^{(\text{row})}(w) = mh_{\text{avg}}^{(\text{column})}(w) + c$ where w indicates a translation-stable word. Each entry in the table contains the coefficient pair m and c . See the scatter plot tableau of Fig. 2 for further details on all language-language comparisons. We use RMA regression, also known as Standardized Major Axis linear regression, because of its accommodation of errors in both variables.

	Spanish	Portuguese	English	Indonesian	French	German	Arabic	Russian	Korean	Chinese
Spanish	1.00	0.89	0.87	0.82	0.86	0.82	0.83	0.73	0.79	0.79
Portuguese	0.89	1.00	0.87	0.82	0.84	0.81	0.84	0.84	0.79	0.76
English	0.87	0.87	1.00	0.88	0.86	0.82	0.86	0.87	0.82	0.81
Indonesian	0.82	0.82	0.88	1.00	0.79	0.77	0.83	0.85	0.79	0.77
French	0.86	0.84	0.86	0.79	1.00	0.84	0.77	0.84	0.79	0.76
German	0.82	0.81	0.82	0.77	0.84	1.00	0.76	0.80	0.73	0.74
Arabic	0.83	0.84	0.86	0.83	0.77	0.76	1.00	0.83	0.79	0.80
Russian	0.73	0.84	0.87	0.85	0.84	0.80	0.83	1.00	0.80	0.82
Korean	0.79	0.79	0.82	0.79	0.79	0.73	0.79	0.80	1.00	0.81
Chinese	0.79	0.76	0.81	0.77	0.76	0.74	0.80	0.82	0.81	1.00

TABLE S4. Pearson correlation coefficients for translation-stable words for all language pairs. All p -values are $< 10^{-118}$. These values are included in Fig. 2 and reproduced here for to facilitate comparison.

	Spanish	Portuguese	English	Indonesian	French	German	Arabic	Russian	Korean	Chinese
Spanish	1.00	0.85	0.83	0.77	0.81	0.77	0.75	0.74	0.74	0.68
Portuguese	0.85	1.00	0.83	0.77	0.78	0.77	0.77	0.81	0.75	0.66
English	0.83	0.83	1.00	0.82	0.80	0.78	0.78	0.81	0.75	0.70
Indonesian	0.77	0.77	0.82	1.00	0.72	0.72	0.76	0.77	0.71	0.71
French	0.81	0.78	0.80	0.72	1.00	0.80	0.67	0.79	0.71	0.64
German	0.77	0.77	0.78	0.72	0.80	1.00	0.69	0.76	0.64	0.62
Arabic	0.75	0.77	0.78	0.76	0.67	0.69	1.00	0.74	0.69	0.68
Russian	0.74	0.81	0.81	0.77	0.79	0.76	0.74	1.00	0.70	0.66
Korean	0.74	0.75	0.75	0.71	0.71	0.64	0.69	0.70	1.00	0.71
Chinese	0.68	0.66	0.70	0.71	0.64	0.62	0.68	0.66	0.71	1.00

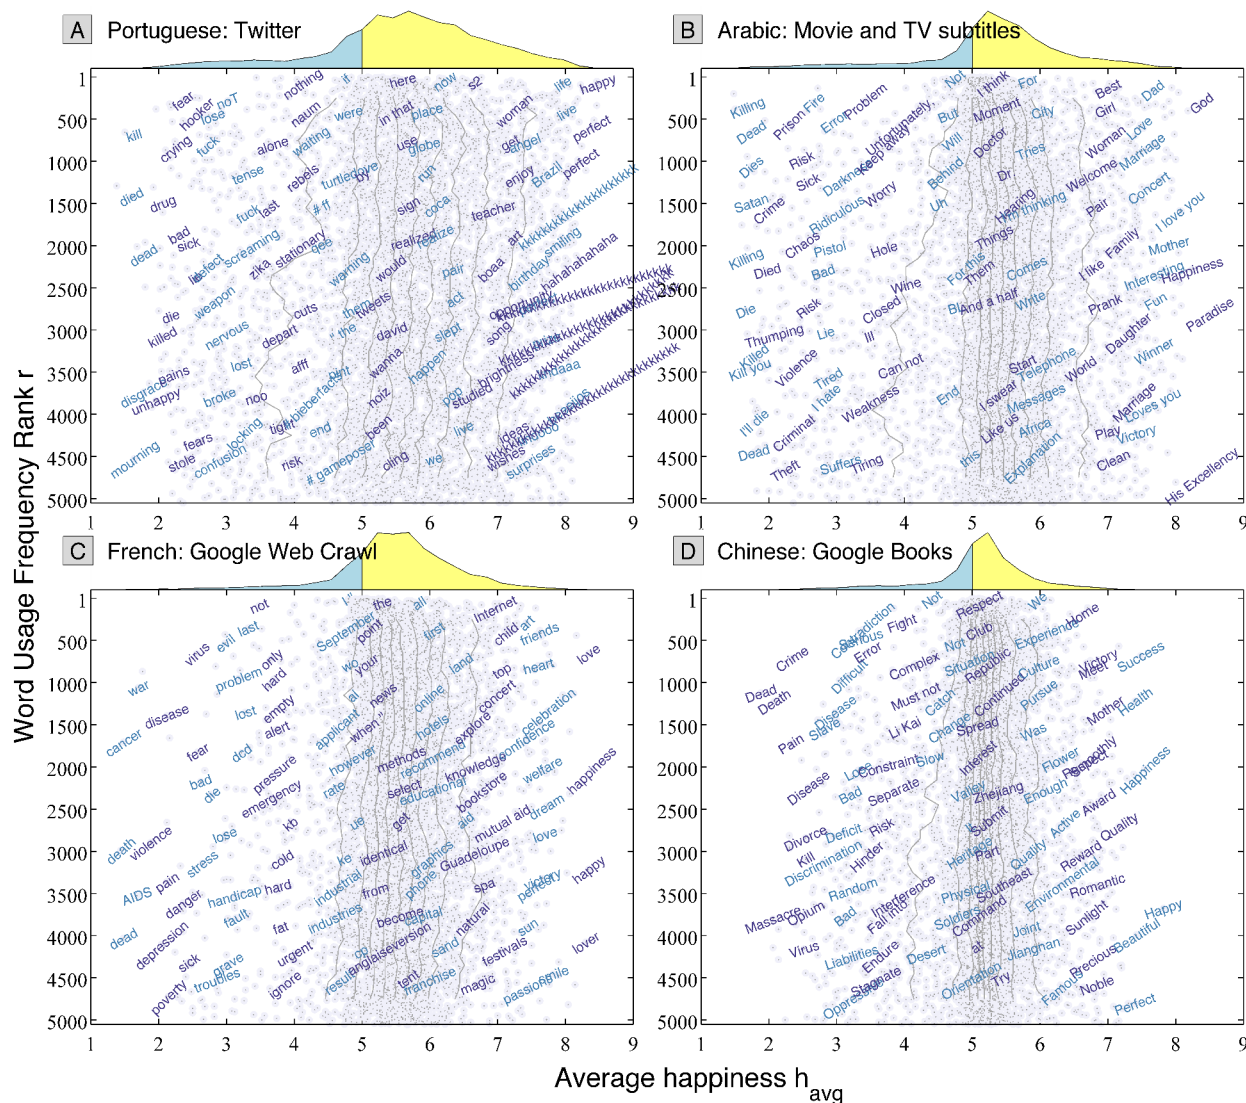
TABLE S5. Spearman correlation coefficients for translation-stable words. All p -values are $< 10^{-82}$.

Language: Corpus	ρ_p	p -value	ρ_s	p -value	α	β
Spanish: Google Web Crawl	-0.114	3.38×10^{-22}	-0.090	1.85×10^{-14}	-5.55×10^{-5}	6.10
Spanish: Google Books	-0.040	1.51×10^{-3}	-0.016	1.90×10^{-1}	-2.28×10^{-5}	5.90
Spanish: Twitter	-0.048	1.14×10^{-4}	-0.032	1.10×10^{-2}	-3.10×10^{-5}	5.94
Portuguese: Google Web Crawl	-0.085	6.33×10^{-13}	-0.060	3.23×10^{-7}	-3.98×10^{-5}	5.96
Portuguese: Twitter	-0.041	5.98×10^{-4}	-0.030	1.15×10^{-2}	-2.40×10^{-5}	5.73
English: Google Books	-0.042	3.03×10^{-3}	-0.013	3.50×10^{-1}	-3.04×10^{-5}	5.62
English: New York Times	-0.056	6.93×10^{-5}	-0.044	1.99×10^{-3}	-4.17×10^{-5}	5.61
German: Google Web Crawl	-0.096	1.11×10^{-15}	-0.082	6.75×10^{-12}	-3.67×10^{-5}	5.65
French: Google Web Crawl	-0.105	9.20×10^{-19}	-0.080	1.99×10^{-11}	-4.50×10^{-5}	5.68
English: Twitter	-0.097	6.56×10^{-12}	-0.103	2.37×10^{-13}	-7.78×10^{-5}	5.67
Indonesian: Movie subtitles	-0.039	1.48×10^{-3}	-0.063	2.45×10^{-7}	-2.04×10^{-5}	5.45
German: Twitter	-0.054	1.47×10^{-5}	-0.036	4.02×10^{-3}	-2.51×10^{-5}	5.58
Russian: Twitter	-0.052	2.38×10^{-5}	-0.028	2.42×10^{-2}	-2.55×10^{-5}	5.52
French: Google Books	-0.043	6.80×10^{-4}	-0.030	1.71×10^{-2}	-2.31×10^{-5}	5.49
German: Google Books	-0.003	8.12×10^{-1}	+0.014	2.74×10^{-1}	-1.38×10^{-6}	5.45
French: Twitter	-0.049	6.08×10^{-5}	-0.023	6.31×10^{-2}	-2.54×10^{-5}	5.54
Russian: Movie and TV subtitles	-0.029	2.36×10^{-2}	-0.033	9.17×10^{-3}	-1.57×10^{-5}	5.43
Arabic: Movie and TV subtitles	-0.045	7.10×10^{-6}	-0.029	4.19×10^{-3}	-1.66×10^{-5}	5.44
Indonesian: Twitter	-0.051	2.14×10^{-5}	-0.018	1.24×10^{-1}	-2.50×10^{-5}	5.46
Korean: Twitter	-0.032	8.29×10^{-3}	-0.016	1.91×10^{-1}	-1.24×10^{-5}	5.38
Russian: Google Books	+0.030	2.09×10^{-2}	+0.070	5.08×10^{-8}	$+1.20 \times 10^{-5}$	5.35
English: Music Lyrics	-0.073	2.53×10^{-7}	-0.081	1.05×10^{-8}	-6.12×10^{-5}	5.45
Korean: Movie subtitles	-0.187	8.22×10^{-44}	-0.180	2.01×10^{-40}	-9.66×10^{-5}	5.41
Chinese: Google Books	-0.067	1.48×10^{-11}	-0.050	5.01×10^{-7}	-1.72×10^{-5}	5.21

TABLE S6. Pearson correlation coefficients and p -values, Spearman correlation coefficients and p -values, and linear fit coefficients, for average word happiness h_{avg} as a function of word usage frequency rank r . We use the fit is $h_{\text{avg}} = \alpha r + \beta$ for the most common 5000 words in each corpora, determining α and β via ordinary least squares, and order languages by the median of their average word happiness scores (descending). We note that stemming of words may affect these estimates.

Language: Corpus	ρ_p	p -value	ρ_s	p -value	α	β
Portuguese: Twitter	+0.090	2.55×10^{-14}	+0.095	1.28×10^{-15}	1.19×10^{-5}	1.29
Spanish: Twitter	+0.097	8.45×10^{-15}	+0.104	5.92×10^{-17}	1.47×10^{-5}	1.26
English: Music Lyrics	+0.129	4.87×10^{-20}	+0.134	1.63×10^{-21}	2.76×10^{-5}	1.33
English: Twitter	+0.007	6.26×10^{-1}	+0.012	4.11×10^{-1}	1.47×10^{-6}	1.35
English: New York Times	+0.050	4.56×10^{-4}	+0.044	1.91×10^{-3}	9.34×10^{-6}	1.32
Arabic: Movie and TV subtitles	+0.101	7.13×10^{-24}	+0.101	3.41×10^{-24}	9.41×10^{-6}	1.01
English: Google Books	+0.180	1.68×10^{-37}	+0.176	4.96×10^{-36}	3.36×10^{-5}	1.27
Spanish: Google Books	+0.066	1.23×10^{-7}	+0.062	6.53×10^{-7}	9.17×10^{-6}	1.26
Indonesian: Movie subtitles	+0.026	3.43×10^{-2}	+0.027	2.81×10^{-2}	2.87×10^{-6}	1.12
Russian: Movie and TV subtitles	+0.083	7.60×10^{-11}	+0.075	3.28×10^{-9}	1.06×10^{-5}	0.89
French: Twitter	+0.072	4.77×10^{-9}	+0.076	8.94×10^{-10}	1.07×10^{-5}	1.05
Indonesian: Twitter	+0.072	1.17×10^{-9}	+0.072	1.73×10^{-9}	8.16×10^{-6}	1.12
French: Google Books	+0.090	1.02×10^{-12}	+0.085	1.67×10^{-11}	1.25×10^{-5}	1.02
Russian: Twitter	+0.055	6.83×10^{-6}	+0.053	1.67×10^{-5}	7.39×10^{-6}	0.91
Spanish: Google Web Crawl	+0.119	4.45×10^{-24}	+0.106	2.60×10^{-19}	1.45×10^{-5}	1.23
Portuguese: Google Web Crawl	+0.093	4.06×10^{-15}	+0.083	2.91×10^{-12}	1.07×10^{-5}	1.26
German: Twitter	+0.051	4.45×10^{-5}	+0.050	5.15×10^{-5}	7.39×10^{-6}	1.15
French: Google Web Crawl	+0.104	2.12×10^{-18}	+0.088	9.64×10^{-14}	1.27×10^{-5}	1.01
Korean: Movie subtitles	+0.171	1.39×10^{-36}	+0.185	8.85×10^{-43}	2.58×10^{-5}	0.88
German: Google Books	+0.157	6.06×10^{-35}	+0.162	4.96×10^{-37}	2.17×10^{-5}	1.03
Korean: Twitter	+0.056	4.07×10^{-6}	+0.062	4.25×10^{-7}	6.98×10^{-6}	0.93
German: Google Web Crawl	+0.099	2.05×10^{-16}	+0.085	1.18×10^{-12}	1.20×10^{-5}	1.07
Chinese: Google Books	+0.099	3.07×10^{-23}	+0.097	3.81×10^{-22}	8.70×10^{-6}	1.16
Russian: Google Books	+0.187	5.15×10^{-48}	+0.177	2.24×10^{-43}	2.28×10^{-5}	0.81

TABLE S7. Pearson correlation coefficients and p -values, Spearman correlation coefficients and p -values, and linear fit coefficients for standard deviation of word happiness h_{std} as a function of word usage frequency rank r . We consider the fit is $h_{\text{std}} = \alpha r + \beta$ for the most common 5000 words in each corpora, determining α and β via ordinary least squares, and order corpora according to their emotional variance (descending).



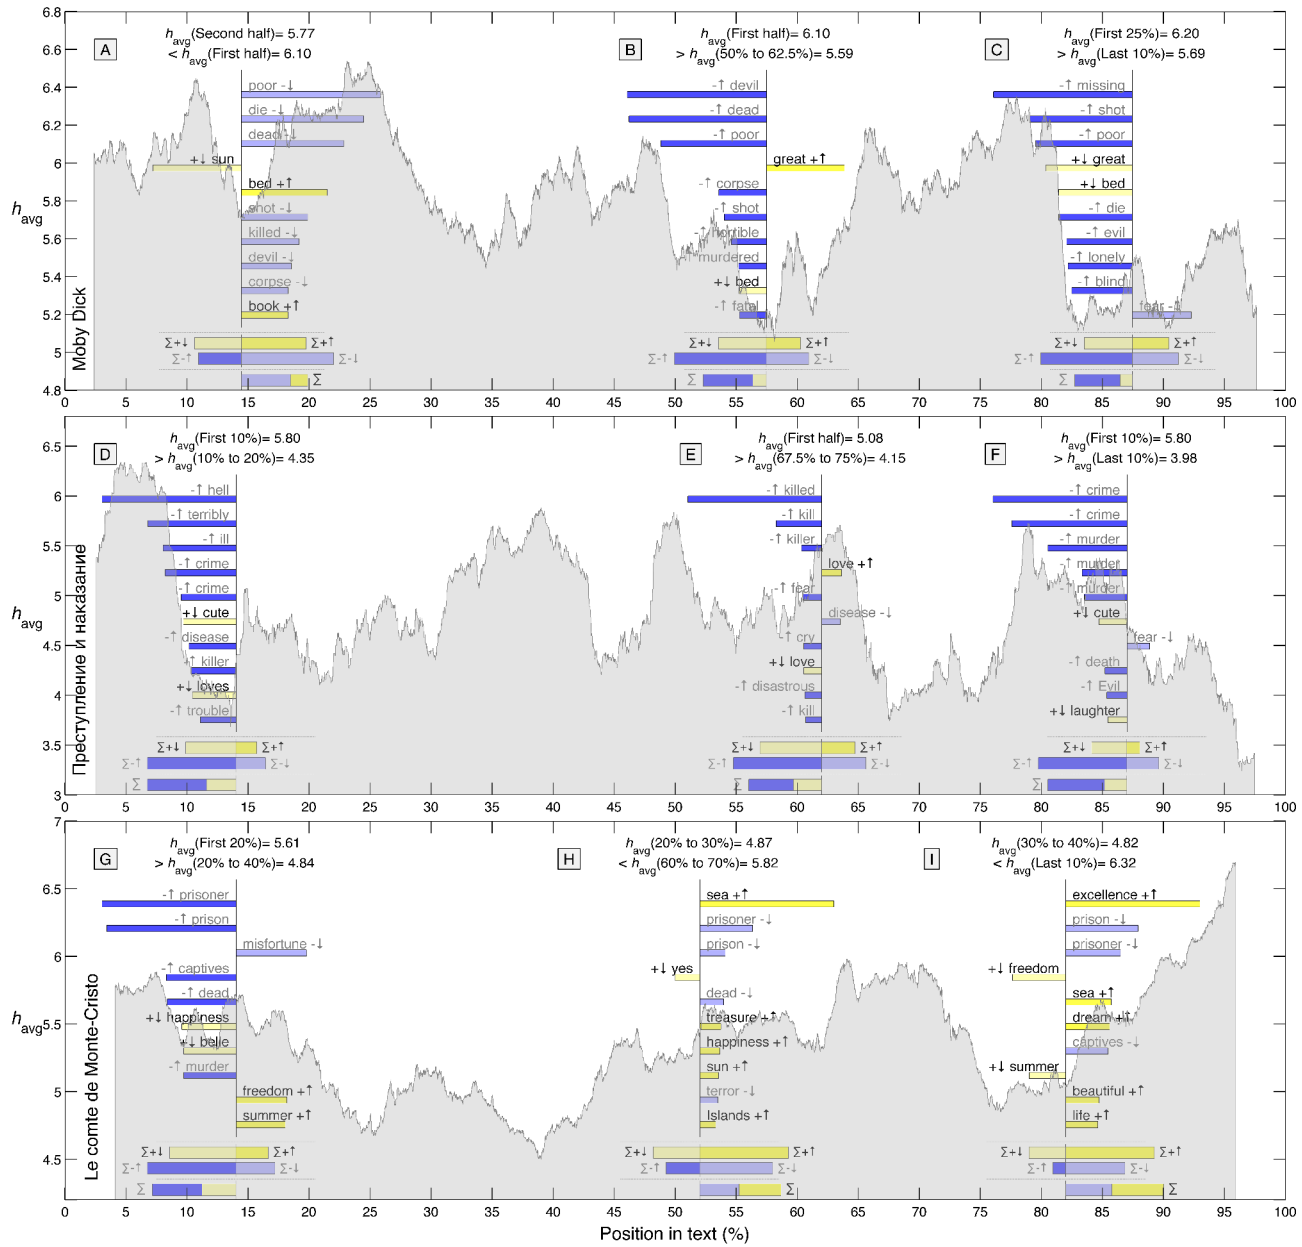


FIG. S3. Fig. 4 from the main text with Russian and French translated into English. Online, interactive visualizations of over 10,000 books can be found here: <http://hedonometer.org/books.html>.

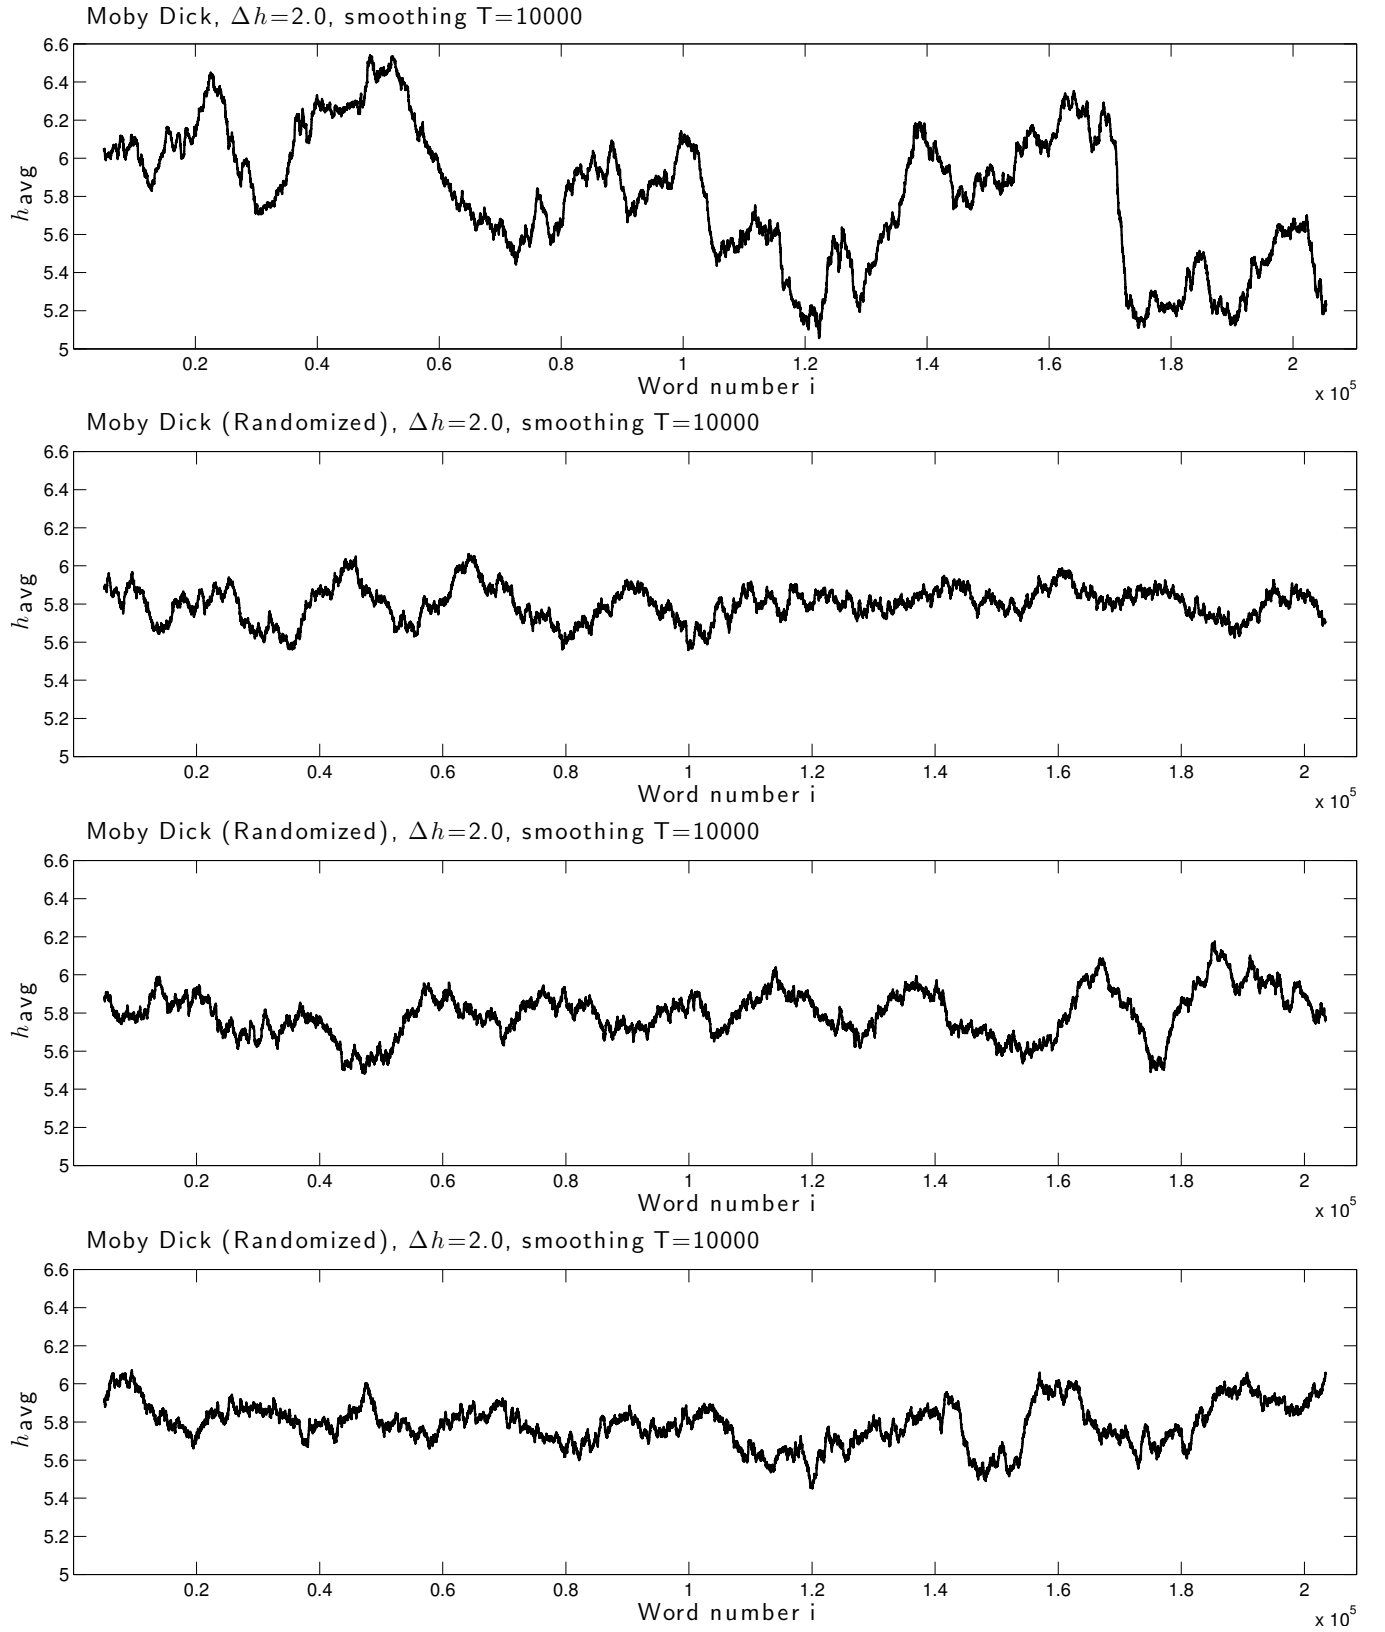



FIG. S4. Comparison of the emotional trajectory of Moby Dick with the results for three example randomized versions of the same text, showing the loss of structure and variability.

EXPLANATION OF WORD SHIFTS

In this section, we explain the word shifts included in Figs. 4 and S3. We expand upon the approach described in [16] and [18] to rank and visualize how words contribute to this overall upward shift in happiness.

Shown below is the third inset word shift used in Fig 4 for the Count of Monte Cristo, a comparison of words found in the last 10% of the book (T_{comp} , $h_{\text{avg}} = 6.32$) relative to those used between 30% and 40% (T_{ref} , $h_{\text{avg}} = 4.82$). For this particular measurement, we employed the ‘word lens’ which excluded words with $3 < h_{\text{avg}} < 7$.



We will use the following probability notation for the normalized frequency of a given word w in a text T :

$$\Pr(w|T; L) = \frac{f(w|T; L)}{\sum_{w' \in L} f(w'|T; L)}, \quad (1)$$

where $f(w|T; L)$ is the frequency of word w in T with word lens L applied [16]. (For the example word shift above, we have $L = \{[1, 3], [7, 9]\}$.) We then estimate the happiness score of any text T as

$$h_{\text{avg}}(T; L) = \sum_{w \in L} h_{\text{avg}}(w) \Pr(w|T; L), \quad (2)$$

where $h_{\text{avg}}(w)$ is the average happiness score of a word as determined by our survey.

We can now express the happiness difference between

two texts as follows:

$$\begin{aligned} & h_{\text{avg}}(T_{\text{comp}}; L) - h_{\text{avg}}(T_{\text{ref}}; L) \\ &= \sum_{w \in L} h_{\text{avg}}(w) \Pr(w|T_{\text{comp}}; L) - \sum_{w \in L} h_{\text{avg}}(w) \Pr(w|T_{\text{ref}}; L) \\ &= \sum_{w \in L} h_{\text{avg}}(w) [\Pr(w|T_{\text{comp}}; L) - \Pr(w|T_{\text{ref}}; L)] \\ &= \sum_{w \in L} [h_{\text{avg}}(w) - h_{\text{avg}}(T_{\text{ref}}; L)] \\ &\quad \times [\Pr(w|T_{\text{comp}}; L) - \Pr(w|T_{\text{ref}}; L)], \end{aligned} \quad (3)$$

where we have introduced $h_{\text{avg}}(T_{\text{ref}}; L)$ as base reference for the average happiness of a word by noting that

$$\begin{aligned} & \sum_{w \in L} h_{\text{avg}}(T_{\text{ref}}; L) [\Pr(w|T_{\text{comp}}; L) - \Pr(w|T_{\text{ref}}; L)] \\ &= h_{\text{avg}}(T_{\text{ref}}; L) \sum_{w \in L} [\Pr(w|T_{\text{comp}}; L) - \Pr(w|T_{\text{ref}}; L)] \\ &= h_{\text{avg}}(T_{\text{ref}}; L) [1 - 1] = 0. \end{aligned} \quad (4)$$

We can now see the change in average happiness between a reference and comparison text as depending on how these two quantities behave for each word:

$$\delta_h(w) = [h_{\text{avg}}(w) - h_{\text{avg}}(T_{\text{ref}}; L)] \quad (5)$$

and

$$\delta_p(w) = [\Pr(w|T_{\text{comp}}; L) - \Pr(w|T_{\text{ref}}; L)]. \quad (6)$$

Words can contribute to or work against a shift in average happiness in four possible ways which we encode with symbols and colors:

- $\delta_h(w) > 0$, $\delta_p(w) > 0$: Words that are more positive than the reference text’s overall average and are used more in the comparison text (+↑, strong yellow).
- $\delta_h(w) < 0$, $\delta_p(w) < 0$: Words that are less positive than the reference text’s overall average but are used less in the comparison text (-↓, pale blue).
- $\delta_h(w) > 0$, $\delta_p(w) < 0$: Words that are more positive than the reference text’s overall average but are used less in the comparison text (+↓, pale yellow).
- $\delta_h(w) < 0$, $\delta_p(w) > 0$: Words that are more positive than the reference text’s overall average and are used more in the comparison text (-↑, strong blue).

Regardless of usage changes, yellow indicates a relatively positive word, blue a negative one. The stronger colors indicate words with the most simple impact: relatively positive or negative words being used more overall.

We order words by the absolute value of their contribution to or against the overall shift, and normalize them as percentages.

Simple Word Shifts

For simple inset word shifts, we show the 10 top words in terms of their absolute contribution to the shift.

Returning to the inset word shift above, we see that an increase in the abundance of relatively positive words ‘excellence’ ‘mer’ and ‘rêve’ ($+\uparrow$, strong yellow) as well as a decrease in the relatively negative words ‘prison’ and ‘prisonnier’ ($-\downarrow$, pale blue) most strongly contribute to the increase in positivity. Some words go against this trend, and in the abbreviated word shift we see less usage of relatively positive words ‘liberté’ and ‘été’ ($+\downarrow$, pale

yellow).

The normalized sum total of each of the four categories of words is shown in the summary bars at the bottom of the word shift. For example, $\Sigma+\uparrow$ represents the total shift due to all relatively positive words that are more prevalent in the comparison text. The smallest contribution comes from relatively negative words being used more ($-\uparrow$, strong blue).

The bottom bar with Σ shows the overall shift with a breakdown of how relatively positive and negative words separately contribute. For the Count of Monte Cristo example, we observe an overall use of relatively positive words and a drop in the use of relatively negative ones (strong yellow and pale blue).