


RESEARCH ARTICLE

Probability-turbulence divergence: A tunable allotaxonomic instrument for comparing heavy-tailed type-probability distributions

Peter Sheridan Dodds^{1,2,3,4*}, Joshua R. Minot⁵, Michael V. Arnold^{1,2}, Thayer Alshaabi^{6,7}, Jane Lydia Adams⁸, Andrew J. Reagan⁵, Christopher M. Danforth^{1,2,9}

1 Computational Story Lab, Vermont Advanced Computing Center, University of Vermont, Burlington, Vermont, United States of America, **2** Vermont Complex Systems Institute, MassMutual Center of Excellence for Complex Systems and Data Science, University of Vermont, Burlington, Vermont, United States of America, **3** Department of Computer Science, University of Vermont, Burlington, Vermont, United States of America, **4** Santa Fe Institute, Santa Fe, New Mexico, United States of America, **5** MassMutual Data Science, Amherst, Massachusetts, United States of America, **6** Howard Hughes Medical Institute, Janelia Research Campus, Ashburn, Vermont, United States of America, **7** Advanced Bioimaging Center, University of California Berkeley, Berkeley, California, United States of America, **8** Data Visualization Lab, Khoury College of Computer Sciences, Northeastern University, Boston, Massachusetts, United States of America, **9** Department of Mathematics & Statistics, University of Vermont, Burlington, Vermont, United States of America

* peter.dodds@uvm.edu



OPEN ACCESS

Citation: Dodds PS, Minot JR, Arnold MV, Alshaabi T, Adams JL, Reagan AJ, et al. (2026) Probability-turbulence divergence: A tunable allotaxonomic instrument for comparing heavy-tailed type-probability distributions. PLOS Complex Syst 3(7): e0000077. <https://doi.org/10.1371/journal.pcsy.0000077>

Editor: Réka Albert, Pennsylvania State University, UNITED STATES OF AMERICA

Received: March 24, 2025

Accepted: September 24, 2025

Published: July 2, 2026

Peer Review History: PLOS recognizes the benefits of transparency in the peer review process; therefore, we enable the publication of all of the content of peer review and author responses alongside final, published articles. The editorial history of this article is available here: <https://doi.org/10.1371/journal.pcsy.0000077>

Abstract

Real-world complex systems often comprise many distinct types of elements as well as many more types of networked interactions between elements. When the relative abundances of types can be measured well, we often observe heavy-tailed distributions for type probabilities (or relative rates). For the comparison of type-probability distributions of two systems or a system with itself at different points in time—a facet of allotaxonomy—a great range of probability divergences are available. Here, we introduce and explore ‘probability-turbulence divergence’, a tunable, straightforward, and interpretable instrument for comparing normalizable type-probability distributions. We model probability-turbulence divergence (PTD) after rank-turbulence divergence (RTD). While probability-turbulence divergence is more limited in application than rank-turbulence divergence, it is more sensitive to changes in type probability. We show how probability-turbulence divergence either explicitly or functionally generalizes many existing kinds of distances and measures, including, as special cases, L_q norms, the Sørensen-Dice coefficient (the F_1 statistic), and the Hellinger distance. We discuss similarities with the generalized entropies of Rényi and Tsallis, and the diversity indices (or Hill numbers) from ecology. We then build allotaxonomographs to display probability turbulence, incorporating a way to visually accommodate zero probabilities for ‘exclusive types’ which are types that appear in only one system. Using flipbooks, we show how tuning PTD’s single parameter informs the user how two systems diverge for types that are rare, common, and at

Copyright: © 2026 Peter Sheridan Dodds et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data availability

statement: Allotaxonomy home:

<https://compstorylab.org/allotaxonomy/>.

GitLab repository: <https://gitlab.com/compstorylab/allotaxonomer/>.

Zenodo record:

<https://doi.org/10.5281/zenodo.15006457>.

Funding: The authors were in part supported by National Science Foundation Award #2242829, <https://www.nsf.gov>, and a gift from the Massachusetts Mutual Life Insurance Company, <https://www.massmutual.com>. These funders played no role in study design, data collection and analysis, decision to publish, or preparation of any materials.

Competing interests: The authors have declared that no competing interests exist.

all scales in between. We demonstrate that PTD can be tuned to a ‘scale-equalizing’ view that is non-universal and dependent on the systems being compared. We explore comparisons of example distributions taken from literature, social media, and ecology. We close with thoughts on open problems concerning the optimization of the tuning of rank- and probability-turbulence divergence.

Author summary

Probability-turbulence divergence (PTD) is an allotaxonomic instrument designed to compare complex systems comprised of many element types which follow heavy-tailed abundance distributions. Allotaxonomographs for probability-turbulence divergence provide map-and-list visualizations that: (1) Properly accommodate zero probabilities, and (2) Surface which elements most differentiate the composition of any two systems. Constructed with a single parameter α , probability-turbulence divergence can, in the manner of a physical instrument, be ‘tuned’, foregrounding rare ($\alpha = 0$) or common types ($\alpha = \infty$) at its limits. For comparable complex systems that exhibit type turbulence, tuning to a ‘scale-equalizing’ value of α gives equal weighting to types of all abundances which in turn leads to a meaningful, summarizing ranking of types. Probability-turbulence divergence generalizes and unifies a wide range of existing distance measures.

1 Introduction

Driven by an interest in developing allotaxonomy [1]—the detailed comparison of any two complex systems comprising many types of elements—we often find we need to compare ‘normalized-size-rank’ distributions: Heavy-tailed categorical distributions of normalized type-sizes [2–5], often referred to as Zipf distributions. The normalized sizes we are most interested in are probabilities or relative rates. We take a relaxed definition of what a heavy tail means for a size-rank distribution: A slow decay over orders of magnitude in type rank. Though not required, power-law decay tails are emblematic signatures of heavy-tailed distributions commonly presented by complex systems [6–10], both observed and theoretical, and provide important examples to consider in our efforts to build a comparison tool.

Across fields, efforts to measure and explain how two probability distributions differ have led to the generation of a great many probability divergences [11–14]. Divergences have been developed for a host of motivations quite apart from our focus here on allotaxonomy, with example families scaffolded around L^p -norms, inner products, and information-theoretic constructions. As we will discuss, for heavy-tailed distribution comparisons which exhibit variable ‘probability turbulence’ [1, 15], we find these divergences lack appropriate adaptability.

Here, we introduce a tunable, interpretable instrument that we call probability-turbulence divergence (PTD) along with related allotaxonomographs—visualizations which show in detail how two categorical distributions differ according to a given

Table 1. Glossary.

Glossary of Key Terms
Allotaxonomy: The comparison of any two complex systems comprising many different types of elements, likely of highly varying 'size' (e.g., number, probability) and compositional structure (e.g., hierarchies, networks). Coined from Greek roots: allo- ('other'), taxis ('order/arrangement'), and metron ('measure'). Introduced in Ref [1].
Allotaxonograph: A map-and-list visualization of how types (e.g., words, species) differ in size (e.g., rank, abundance) between two systems.
Probability-Turbulence Divergence (PTD): A tunable divergence measure introduced in this paper for comparing two probability distributions over categorical types. PTD generalizes similarity measures (e.g., Sørensen-Dice) by incorporating a parameter α that adjusts the weight given to rare vs. common types.
Rank-Turbulence Divergence (RTD): An instrument for comparing two ranked lists that come from heavy-tailed distributions of type sizes, which need not be known (i.e., only ranks matter).
Tuning parameter α: Real-valued parameter that is ≥ 0 which controls how PTD (or RTD) emphasizes rare vs. common types. Lower values highlight rare types; higher values prioritize common ones. The parameter can also be viewed as a diagnostic quantity.
Scale-equalizing α: Where probability turbulence scales well, a value of α for which rare and common types and types in between all contribute in a balanced way to probability-turbulence divergence. This value of α is not exact but rather represents a small range outside of which rare or common types begin to dominate the divergence score.
Flipbook sweeps over α: PDF booklets of allotaxonographs for α systematically varying from 0 to ∞ (see Eq 23). Allows users to observe how the map-and-list visualizations change with α , and to determine, if desired and if possible, a scale-equalizing value of α .

<https://doi.org/10.1371/journal.pcsy.0000077.t001>

measure. We refer the reader to Ref [1] for our motivation for creating allotaxonomy and allotaxonographs, the notion of rank turbulence, and a detailed justification for the form of rank-turbulence divergence (RTD) we developed there. We establish probability-turbulence divergence using largely the same arguments, though PTD is distinct from RTD in its application and properties. We will therefore be concise in our presentation and expand as needed when the probability version's behavior departs from that of its rank counterpart.

In Sect 2, we formally define probability-turbulence divergence. We describe the divergence's general analytic behavior as a function of its single parameter, α , and we determine its form for the two limits of the parameter, $\alpha = 0$ and $\alpha = \infty$. When $\alpha = 0$, in particular, we find an interesting departure from the equivalent tuning for rank-turbulence divergence, and which we will later connect to the Sørensen-Dice coefficient [16,17] and the F_1 score [18].

In Sect 3, we show that probability-turbulence divergence is either a generalization of or may be connected to a number of other kinds of divergences and similarities (e.g., the Sørensen-Dice coefficient), and then discuss limited functional similarities with the Rényi entropy and diversity indices [19–23].

In Sect 4, we then provide realizations of probability-turbulence divergence as an instrument through example allotaxonographs. For three disparate examples, we consider 1. Normalized frequency of n -gram use in Jane Austen's *Pride and Prejudice*, 2. Normalized frequency of n -gram use on Twitter, and 3. Tree species abundance [24]. We show how, for the kinds of heavy-tailed distributions we are interested in, a probability-turbulence divergence histogram can be constructed to accommodate both a logarithmic scale and the presence of zero probabilities. Similar to rank-turbulence divergence histograms, these graphs clearly show whether or not a probability-based divergence is a suitable choice for any given comparison. We explore allotaxonographs for the full range of parameter α , highlighting the existence of special, non-universal 'scale-equalizing' values of α . To demonstrate the tunability of probability-turbulence divergence, we provide allotaxonographs 'flipbooks' as part of the supplementary information [25]. We consider the construction and examination of these flipbooks essential to our method of comparison for any two systems.

We outline data, allotaxonomy code, and supplementary material in Sect 5, and we offer some concluding thoughts in Sect 6.

2 Probability-turbulence divergence

We aim to compare two systems Ω_1 and Ω_2 for both of which we have a list of component types and their probabilities. For simplicity, we will use probability in our general derivations. In describing any particular distribution, we will use an

appropriate terminology, such as normalized frequency, relative frequency, or rate of usage. All normalizations must be such that the sum of ‘sizes’ is 1.

We denote a type by τ and its probability in the two systems as $p_{\tau,1}$ and $p_{\tau,2}$. We represent the probability distributions for the two systems as P_1 and P_2 . We call types that are present in one system only ‘exclusive types’. We will use expressions of the form $\Omega^{(1)}$ -exclusive and $\Omega^{(2)}$ -exclusive to indicate to which system a type solely belongs.

We are interested in divergences that are some function of a sum of contributions by type. Here, we will consider a single-parameter family of divergences that are of the simplest form, i.e., a direct sum of contributions:

$$D_{\alpha}^P(P_1 \| P_2) = \sum_{\tau \in R_{1,2;\alpha}} \delta D_{\alpha,\tau}^P(P_1 \| P_2). \quad (1)$$

By the rank-ordered set $R_{1,2;\alpha}$, we indicate the union of all types from both systems, sequenced such that the contributions $\delta D_{\alpha,\tau}^P(P_1 \| P_2)$ are monotonically decreasing (hence the necessity of an α subscript). We impose this order for general good housekeeping, secondarily allowing us to handle possibilities such as truncated summations due to sampling, or convergence issues for theoretical examples.

Our motivation starts with the observed phenomenon of type turbulence in the comparison of complex systems with heavy-tailed distributions of types according to some kind of size [1,15]. Our objective is to create a tunable probability divergence measure that can contend with variable type turbulence. In doing so, we will incidentally arrive at a divergence that can be tuned from depending on only the most rare words to only the most common words.

As a start, and in a manner similar to how we approached rank-turbulence divergence [1], we consider the base difference quantity for type τ :

$$| [p_{\tau,1}]^{\alpha} - [p_{\tau,2}]^{\alpha} |, \quad (2)$$

where $\alpha \geq 0$ is a tuning parameter. Evidently, low α will dampen differences between probabilities and high α will accentuate them. However, increasing α will also reduce the difference quantity to 0, removing its ability to contribute. We partly solve this by raising the quantity to $1/\alpha$:

$$| [p_{\tau,1}]^{\alpha} - [p_{\tau,2}]^{\alpha} |^{1/\alpha}. \quad (3)$$

Now the large α limit functions well with $\lim_{\alpha \rightarrow \infty} | [p_{\tau,1}]^{\alpha} - [p_{\tau,2}]^{\alpha} |^{1/\alpha} = \max_{\tau}(p_{\tau,1}, p_{\tau,2})$. But now we have a problem for the $\alpha \rightarrow 0$ limit as the difference quantity will tend toward ∞ except in the rare case that the probabilities are the same. We adjust the difference quantity once more in a parsimonious way:

$$\frac{\alpha + 1}{\alpha} \sum_{\tau \in R_{1,2;\alpha}} | [p_{\tau,1}]^{\alpha} - [p_{\tau,2}]^{\alpha} |^{1/(\alpha+1)}. \quad (4)$$

The behavior for large α remains the same, and now for $\alpha \rightarrow 0$, the difference quantity converges and does so in a meaningful way. We explicate the $\alpha \rightarrow 0$ limit fully below in Sect 2.2.

We now define probability-turbulence divergence as:

$$\begin{aligned} D_{\alpha}^P(P_1 \| P_2) \\ = \frac{1}{\mathcal{N}_{1,2;\alpha}^P} \frac{\alpha + 1}{\alpha} \sum_{\tau \in R_{1,2;\alpha}} | [p_{\tau,1}]^{\alpha} - [p_{\tau,2}]^{\alpha} |^{1/(\alpha+1)}. \end{aligned} \quad (5)$$

where the parameter α may be tuned from 0 to ∞ and $\mathcal{N}_{1,2;\alpha}^P$ is a normalization factor. Per Ref [1] and below in Sect 2.2, the roles of the prefactor $(\alpha + 1)/\alpha$ and the power $1/(\alpha + 1)$ are to govern the behavior of PTD in the limit $\alpha \rightarrow 0$.

As we will show below, sweeping across α allows the user to accentuate the importance of the most rare types (when $\alpha = 0$) through to the most common types ($\alpha = \infty$). Further, for certain special values of α , PTD corresponds in behavior to many extant divergences and differences. And finally, for system comparisons that exhibit scaling of probability-turbulence, users will be able to find a scale-equalizing α .

By construction and regardless of the choice of normalization factor, we can see from Eq 5 that probability-turbulence divergence will equal 0 when both distributions are the same. (We show below that for $\alpha = 0$, distinct distributions can also register $D_\alpha^P(P_1 \| P_2) = 0$.)

The core of Eq 5 is the absolute value of the difference of each type τ 's probability raised to the power of α :

$$| [p_{\tau,1}]^\alpha - [p_{\tau,2}]^\alpha |. \quad (6)$$

This α -tuned quantity controls the order of contributions by types to the overall value of PTD. As $\alpha \rightarrow 0$, lower probabilities—corresponding to the rare types—are relatively accentuated. For $\alpha \rightarrow \infty$, the higher of the two probabilities will dominate (unless they are equal), meaning the most common types will come to the fore.

2.1 Normalization for probability-turbulence divergence

As for rank-turbulence divergence, we choose $\mathcal{N}_{1,2;\alpha}^P$ so that when the two systems are entirely disjoint—that is, they share no types—then probability-turbulence divergence maximizes at 1. The normalization is thus specific to the two distributions being compared. We imagine that the types in each system have an extra descriptor specifying belonging to $\Omega^{(1)}$ or $\Omega^{(2)}$. With no matching types, the probability of a type present in one system is zero in the other, and the sum can be split between the two systems' types:

$$\mathcal{N}_{1,2;\alpha}^P = \frac{\alpha + 1}{\alpha} \sum_{\tau \in R_1} [p_{\tau,1}]^{\alpha/(\alpha+1)} + \frac{\alpha + 1}{\alpha} \sum_{\tau \in R_2} [p_{\tau,2}]^{\alpha/(\alpha+1)}, \quad (7)$$

where R_1 and R_2 are the rank-ordered sets of types for each system.

We can more compactly express the normalization as:

$$\mathcal{N}_{1,2;\alpha}^P = \frac{\alpha + 1}{\alpha} \sum_{\tau \in R_{1,2;\alpha}} \left([p_{\tau,1}]^{\alpha/(\alpha+1)} + [p_{\tau,2}]^{\alpha/(\alpha+1)} \right), \quad (8)$$

where, as for the definition of probability-turbulence divergence in Eq 5, the sum for the normalization is again over the ordered set $R_{1,2;\alpha}$. Types that appear in both systems will have their contribution $[p_{\tau,1}]^{\alpha/(\alpha+1)}$ and $[p_{\tau,2}]^{\alpha/(\alpha+1)}$ counted appropriately.

2.2 Normalization for $\alpha = 0$

The $\alpha \rightarrow 0$ limit requires some care and will vary from the equivalent limit for rank-turbulence divergence [1]. First, at the level of individual type contribution, if both $p_{\tau,1} > 0$ and $p_{\tau,2} > 0$ then

$$\lim_{\alpha \rightarrow 0} \frac{\alpha + 1}{\alpha} | [p_{\tau,1}]^\alpha - [p_{\tau,2}]^\alpha |^{1/(\alpha+1)} = \left| \ln \frac{p_{\tau,2}}{p_{\tau,1}} \right|. \quad (9)$$

If instead a type τ is exclusive to one system, meaning either $p_{\tau,1} = 0$ or $p_{\tau,2} = 0$, then the limit diverges as $1/\alpha$, which would seem problematic. We nevertheless will arrive at a well-behaved divergence through the normalization term $\mathcal{N}_{1,2;0}^P$.

Requiring as we have that the extreme of disjoint systems have a divergence of 1, we observe that each of the types in the case of disjoint systems would contribute $1/\alpha$. Therefore, in the $\alpha \rightarrow 0$ limit, we must have:

$$\mathcal{N}_{1,2;\alpha}^P \rightarrow \frac{1}{\alpha} (N_1 + N_2). \quad (10)$$

Because the normalization also diverges as $1/\alpha$, the divergence will be zero when there are no exclusive types and non-zero when there are exclusive types. We can combine these cases into a single expression:

$$D_0^P(P_1 \| P_2) = \frac{1}{(N_1 + N_2)} \sum_{\tau \in R_{1,2;0}} (\delta_{p_{\tau,1},0} + \delta_{0,p_{\tau,2}}). \quad (11)$$

The term $(\delta_{p_{\tau,1},0} + \delta_{0,p_{\tau,2}})$ returns 1 if either $p_{\tau,1} = 0$ or $p_{\tau,2} = 0$, and 0 otherwise when both $p_{\tau,1} > 0$ and $p_{\tau,2} > 0$. (By construction, we cannot have $p_{\tau,1} = p_{\tau,2} = 0$ as each type must be present in either one or both systems.)

We see then that $D_0^P(P_1 \| P_2)$ is the ratio of types that are exclusive to one system relative to the total possible such types, $N_1 + N_2$. If and only if all types appear in both systems with whatever variation in probabilities, then $D_0^P(P_1 \| P_2) = 0$. The limit of $\alpha = 0$ therefore exhibits special behavior. For $\alpha > 0$, probability-turbulence divergence only scores 0 for exactly matching distributions.

2.3 Type contribution ordering for the limit of $\alpha = 0$

In terms of contribution to the divergence score, all exclusive types supply a weight of $1/(N_1 + N_2)$. We can order them by preserving their ordering as $\alpha \rightarrow 0$, which amounts to ordering by descending probability in the system in which they appear.

And while types that appear in both systems make no contribution to $D_0^P(P_1 \| P_2)$, we can still order them according to the log ratio of their probabilities, Eq 9.

The overall ordering of types by divergence contribution for $\alpha = 0$ is then: (1) exclusive types by descending probability and then (2) types appearing in both systems by descending log ratio.

2.4 Normalization for $\alpha = \infty$

The $\alpha \rightarrow \infty$ limit is straightforward and in line with that of rank-turbulence divergence [1]:

$$D_\infty^P(P_1 \| P_2) = \frac{1}{2} \sum_{\tau \in R_{1,2;\infty}} (1 - \delta_{p_{\tau,1},p_{\tau,2}}) \max(p_{\tau,1}, p_{\tau,2}), \quad (12)$$

where the normalization from Eq 8 has become

$$\mathcal{N}_{1,2;\infty}^P = \sum_{\tau \in R_{1,2;\infty}} (p_{\tau,1} + p_{\tau,2}) = 1 + 1 = 2. \quad (13)$$

The dominant contributions to probability-turbulence divergence in the $\alpha \rightarrow \infty$ limit therefore come from the most common types in each system, providing they are not equally abundant.

3 Connections to other divergences and entropies

3.1 Links to existing probability-based divergences

Probability-turbulence divergence shares characteristics with and generalizes many other divergences (see Refs. [11] and [12] for two example compendia). In particular, we find known distances and similarities which correspond with or function similarly to probability-turbulence divergence for $\alpha = 0, 1/2, 1$, and ∞ .

3.1.1 Case 1: $\alpha = 0$. For $\alpha = 0$, probability-turbulence divergence partners the similarity measure Sørensen-Dice coefficient, $S^{S-D}(P_1 \parallel P_2)$ [16,17,26], which was independently developed in the context of ecology by Dice (1945) and Sørensen (1948) (see also Ref [27]). For two systems, the Sørensen-Dice coefficient is the number of shared types relative to the mean of the number of types in each system. Using our notation, and referring back to Eq 11, we have:

$$\begin{aligned} S^{S-D}(P_1 \parallel P_2) &= \frac{2}{(N_1 + N_2)} \sum_{\tau \in R_{1,2,0}} (1 - \delta_{p_{\tau,1},0} - \delta_{0,p_{\tau,2}}) \\ &= 1 - D_0^P(P_1 \parallel P_2), \end{aligned} \quad (14)$$

where we are again summing over the union of types $R_{1,2,0}$. The quantity $(1 - \delta_{p_{\tau,1},0} - \delta_{0,p_{\tau,2}})$ is 1 when a type appears in both systems and 0 otherwise.

The Sørensen-Dice coefficient has arisen in many settings, with different names. For example, in statistics, the Sørensen-Dice coefficient is the F_1 score of a test's accuracy [18,28].

3.1.2 Case 2: $\alpha = 1/2$. Examples of divergences matching the internal structure of $D_{1/2}^P$ include the Hellinger distance [29] (1909), $\mathcal{D}^{\text{Hel}}$, the proportional Matusita distance [30] (1955), $\mathcal{D}^{\text{Mat}}$, and the squared-chord distance [31] (2003), $\mathcal{D}^{\text{sqd-chd}}$.

$$\begin{aligned} &[D_{1/2}^P(P_1 \parallel P_2)]^{3/4} \\ &\propto 2^{-1/2} \mathcal{D}^{\text{Hel}}(P_1 \parallel P_2) \\ &= \mathcal{D}^{\text{Mat}}(P_1 \parallel P_2) \\ &= [\mathcal{D}^{\text{sqd-chd}}(P_1 \parallel P_2)]^{1/2} \\ &= \left[\sum_{\tau \in R_{1,2,1/2}} ([p_{\tau,1}]^{1/2} - [p_{\tau,2}]^{1/2})^2 \right]^{1/2}. \end{aligned} \quad (15)$$

$$= \left[\sum_{\tau \in R_{1,2,1/2}} ([p_{\tau,1}]^{1/2} - [p_{\tau,2}]^{1/2})^2 \right]^{1/2}. \quad (16)$$

3.1.3 Case 3: $\alpha = 1$. In terms of the probability-turbulence divergence's internal structure of $|[p_{\tau,1}]^\alpha - [p_{\tau,2}]^\alpha|$, a large selection of divergences match up with the $\alpha = 1$ instance. These include all constructions built around the L_q -space Minkowski distance [12,32–34]:

$$\mathcal{D}_q(P_1 \parallel P_2) = \left[\sum_{\tau \in R_{1,2,1}} |[p_{\tau,1}]^1 - [p_{\tau,2}]^1|^q \right]^{1/q}. \quad (17)$$

We have explicitly included the powers of 1 for the probabilities to emphasize the match with $\alpha = 1$ for PTD.

For example, for $q = 1$, we have the city block (or Manhattan) distance:

$$\mathcal{D}_1(P_1 \parallel P_2) = \sum_{\tau \in R_{1,2,1}} |[p_{\tau,1}]^1 - [p_{\tau,2}]^1|, \quad (18)$$

and for $q = 2$, the standard spatial (or Euclidean) distance:

$$\mathcal{D}_2(P_1 \parallel P_2) = \left[\sum_{\tau \in R_{1,2,1}} \left| [p_{\tau,1}]^1 - [p_{\tau,2}]^1 \right|^2 \right]^{1/2}. \quad (19)$$

Note that in the $q \rightarrow 0$ and $q \rightarrow \infty$ limits, the connection between the L_q -distance and D_1^P breaks down. In general, as $q \rightarrow 0$, for two vectors $\vec{z}_1$ and $\vec{z}_2$, $L_q(\vec{z}_1, \vec{z}_2)$ will be: 0 if $\vec{z}_1 = \vec{z}_2$; 1 if the vectors differ in exactly one coordinate; and ∞ otherwise. Because two probability distributions cannot differ in just one coordinate, $\mathcal{D}_0(P_1 \parallel P_2)$ is either 0 or ∞ . And because of how we have constructed D_α^P to behave well and meaningfully in the limit $\alpha \rightarrow 0$, the link breaks down.

For the limit $\alpha \rightarrow \infty$, $\mathcal{D}_\infty(P_1 \parallel P_2) = \max_\tau |p_{\tau,1} - p_{\tau,2}|$ which is the Chebyshev distance [35]. For D_∞^P , the sum remains and maxima are determined for each element rather than overall (Eq 12).

Now, while the overall values for these L_q distances will vary, the rank orderings of their contributing types will all be identical to that of D_1^P providing $0 < q < \infty$. To see this, we first leave aside the overall power of $1/q$ for $\mathcal{D}_q(P_1 \parallel P_2)$, as well as prefactors for D_1^P (see Eq 5). For type τ , we then have the summands $|p_{\tau,1} - p_{\tau,2}|^q$ for $\mathcal{D}_q(P_1 \parallel P_2)$, and $|p_{\tau,1} - p_{\tau,2}|^{1/2}$ for $D_1^P(P_1 \parallel P_2)$. Because if $\beta > 0$, $|z_1|^\beta < |z_2|^\beta$ for all $|z_1| < |z_2|$, the orderings of types by contribution will be the same (i.e., the powers q and $1/2$ do not matter, just that they are both between 0 and ∞).

3.1.4 Case 4: $\alpha = \infty$. Finally, in the $\alpha = \infty$ limit, D_∞^P is akin to, and in practice the same as, the Motyka distance [11], $\mathcal{D}^{\text{Mot}}$:

$$\mathcal{D}^{\text{Mot}} = \frac{\sum_{\tau \in R_{1,2,\infty}} \max_\tau (p_{\tau,1}, p_{\tau,2})}{\sum_{\tau \in R_{1,2,\infty}} (p_{\tau,1} + p_{\tau,2})}. \quad (20)$$

Because of the factor $\delta_{p_{\tau,1}, p_{\tau,2}}$ in Eq 12, if there are any types which have the same probability in both systems, then there will be disagreement. However, for real, large-scale systems, the total number of types will likely differ, greatly reducing the chance of any equal probabilities. And where the type set is fixed for both systems, large sample numbers should prevent equivalences.

We summarize the connections between probability-turbulence divergence and other measures in Table 2.

3.2 Generalized entropies and Hill numbers

While none of these other divergences provide direct tunability of the type probability—a severe limitation, as we intend our examples below in Sect 4 will convey—there are well-established quantities which do.

As we observed for rank-turbulence divergence in [1], the parameter α 's effect is similar to that of its counterparts in various kinds of generalized entropy including those of Rényi and Tsallis [19–21] and, more directly, the diversity indices (or Hill numbers) from ecology [22,23]. We make some links with Rényi entropy, leaving other connections for possible future work.

Table 2. Correspondences between probability-turbulence divergence and extant difference measures for distributions. Equivalence is at the level of the contribution by individual types to overall score.

α	Measures with matching type orderings
0	Similarity measure Sørensen-Dice coefficient [16,17,26], F_1 score of a test's accuracy [18,28]
1/2	Hellinger distance [29], Matusita distance [30], squared-chord distance [31]
1	Many including all L_q -distance constructions for $0 < q < \infty$ [12]
∞	Motyka distance [11] (providing no types have the exact same probabilities)

<https://doi.org/10.1371/journal.pcsy.0000077.t002>

Rényi entropy, ${}^\alpha H$, and the associated diversity index, ${}^\alpha N$, are defined as:

$${}^\alpha H = \ln [{}^\alpha N] = \frac{1}{1-\alpha} \ln \left(\sum_{\tau} [p_{\tau}]^{\alpha} \right), \quad (21)$$

and

$${}^\alpha N = \left(\sum_{\tau} [p_{\tau}]^{\alpha} \right)^{1/(1-\alpha)}, \quad (22)$$

where $\alpha \geq 0$. We acknowledge that, at the risk of a minor dislocation from relevant literature, we have had to confront some notation peril here as a standard notation for the diversity index is ${}^\alpha D$. We have also already used N in our present paper but this choice tracks sensibly: As $\alpha \rightarrow 0$, we retrieve the natural logarithm of the number of distinct types N (species richness in ecology) for Rényi entropy, and therefore the diversity index is ${}^0 N = N$. As $\alpha \rightarrow \infty$, the most abundant type will dominate, with min-entropy the limit: ${}^\infty N = \min_{\tau} 1/p_{\tau} = 1/\max_{\tau} p_{\tau}$. In the $\alpha \rightarrow 1$ limit, we recover Shannon's entropy, H , as well as ${}^1 N = e^{H_1} = e^H$.

There are similar aspects for probability-turbulence divergence and the diversity index in the limits of $\alpha = 0$ and ∞ . For $\alpha = 0$, for example, both reduce to quantities involving simple counts of distinct types. Nevertheless, we note that we cannot construct probability-turbulence divergence from manipulations of Rényi entropy or the diversity index. We can, roughly speaking, only create a difference of sums whereas we need a sum of absolute differences with suitable exponents.

4 Probability-turbulence divergence allotaxonographs

We assert that the successful use of our rank- and probability-turbulence divergences is best achieved through consideration of rich graphical representations which we have called allotaxonographs [1]. In this section, we present and describe allotaxonographs comparing probability distributions using probability-turbulence divergence for:

- Normalized usage frequencies of 2-grams in the first and second halves of Jane Austen's *Pride and Prejudice* [36] for $\alpha = 0, 1/2, 3/4$ (scale-equalizing), 1, and ∞ (Figs 1–5);
- Normalized usage frequencies of n -grams in all English-identified tweets on 2020/03/12 and 2020/05/30 for $n = 1, 2$, and 3 and for scale-equalizing values of $\alpha = 1/3, 5/6$, and ∞ (Figs 6, 7, and 8);
- Relative abundances of tree species on Barro Colorado Island for five-year censuses concluding in 1985 and 2015 for a scale-equalizing $\alpha = 1/3$ (Fig 9).

In particular, the *Pride and Prejudice* examples for 2-grams will show how PTD may be adjusted to: emphasize rare n -grams ($\alpha = 0$), emphasize the most common n -grams ($\alpha = \infty$), be scale-equalizing ($\alpha = 3/4$), or behave like many extant distances and divergences ($\alpha = 1/2$ and 1). The choices of α for the three Twitter examples and the one from Barro Colorado Island further showcase how scale-equalizing fits are achieved by a range of values of α . There is no universal α that scale-equalizes turbulence between probability-rank distributions.

The examples for 2-grams and 3-grams can also be seen as demonstrations of possible comparisons of features of complex networks and systems (e.g., 2-grams in text as directed edges).

Based on our experience using RTD and PTD, we advise that either comparative instrument always be used to produce a sequence of allotaxonographs with the following 24 values of α :

$$\vec{\alpha} = \left\{ 0, \frac{1}{12}, \frac{2}{12}, \dots, \frac{18}{12}, 2, 3, 5, 10, \infty \right\}. \quad (23)$$




Fig 1. Allotaxonograph comparing 2-gram usage in the first and second halves of Jane Austen's *Pride and Prejudice* using probability-turbulence divergence with $\alpha = 0$, D_0^P . To create our allotaxonographs, we bin all non-zero probability pairs ($\log_{10} p_{\tau,1}$, $\log_{10} p_{\tau,2}$) in logarithmic space. Colors indicate counts of 2-grams per cell, and we highlight example 2-grams along the edges of the histogram. For pairs where one of the probabilities is zero, we add a separate rectangular panel along the bottom of each axis (lighter gray and lighter blue). Contour lines indicate where probability-turbulence divergence is constant (the jump to the zero probability region necessitates a break in smoothness). The gray scale for 2-grams is indexed by their percentage contribution to probability-turbulence divergence, $\delta D_{0,\tau}^P$. **Ranked list on the right:** We order the most salient 2-grams according to their overall contribution $\delta D_{0,\tau}^P$ which we mark by bar length. We show the rank pair for each 2-gram in light gray opposite each 2-gram. **For $\alpha = 0$ at one extreme of the parameter's range, probability-turbulence divergence elevates exclusive types above all types that appear in both systems, and the ranked list on the right comprises only system-exclusive 2-grams.** Per Sect 2.2 and Eq 11, exclusive types each equally contribute $\frac{1}{(N_1+N_2)}$ to D_0^P while types appearing in both systems have zero weight. The ordering of 2-grams is determined by maintaining their contribution order as α approaches 0. We force the contour lines in the main body of the histogram to remain equally spaced, even as they all represent 0 in the $\alpha \rightarrow 0$ limit. Outside of the main body of the histogram, all contour lines travel to the (0,0) point (which is $(-\infty, -\infty)$ in log-space). See Sect 3 for the connection between D_0^P and the Sørensen-Dice coefficient [16,17,26] and the F_1 score [18,28]. See also S1, S2, and S3 Flipbooks in the supplementary material, per Sect 5.3.

<https://doi.org/10.1371/journal.pcsy.0000077.g001>

All of our comparisons here begin with this sequence of the α parameter.

In the same fashion as our rank-turbulence divergence allotaxonographs [1]—but with some necessary and key modifications—our allotaxonographs for probability-turbulence divergence pair two complementary visualizations: A map-like histogram and a ranked list.

In isolation, both the histogram and the ranked list have important but limited descriptive power. The histogram helps us see how well our choice of α performs, information that is entirely lost by the ranking process. And the ranked list would be difficult to intuit from the histogram alone.




Fig 2. Allotaxonograph comparing 2-gram usage in the first and second halves of Jane Austen's *Pride and Prejudice* using probability-turbulence divergence with $\alpha = 1/2$, $D_{1/2}^P$.

<https://doi.org/10.1371/journal.pcsy.0000077.g002>

Many aspects of our allotaxonographs are configurable. On [GitLab](#), we provide our universal code for generating allotaxonographs for rank-turbulence divergence, probability-turbulence divergence, and other probability divergences (see [Sect 5.2](#)).

We complement all of our allotaxonographs with PDF flipbooks which move systematically through a range of α values. These flipbooks can be found in the paper's supplementary material on [Zenodo](#) [25], the paper's Online Appendices (<https://compstoriylab.org/allotaxonometry/>), and the associated GitLab repository (<https://gitlab.com/compstoriylab/allotaxonometer/>).

4.1 Allotaxonographs for *Pride and Prejudice*

For a primary, familiar example to help us explain our probability-turbulence divergence allotaxonographs, we compare the normalized usage frequency distributions of 2-grams between the first and second halves of *Pride and Prejudice* [36]. We emphasize that we are mostly intent here on showing how allotaxonographs function. We are not attempting to reveal astonishing insights into one of the most well-regarded and well-studied novels of all time.

In [S1](#), [S2](#), and [S3](#) Flipbooks in the paper's supplementary material [25], we give the reader a view into how our allotaxonometric comparisons of the usage frequencies of 1-grams, 2-grams, and 3-grams in the two halves of *Pride and Prejudice* behave as a function of α (see [Sect 5.3](#)). All of the flipbooks are especially informative in showing how contour




Fig 3. Allotaxonograph comparing 2-gram usage in the first and second halves of Jane Austen's *Pride and Prejudice* using probability-turbulence divergence with a scale-equalizing $\alpha = 3/4$, $D_{3/4}^P$. By scale-equalizing, we mean that the contour lines reasonably fit with the outer shape of the histogram, thereby giving the list on the right a mixture of contributions from rare to common 2-grams.

<https://doi.org/10.1371/journal.pcsy.0000077.g003>

lines and ranked lists change with α . Here in the main paper, we present and discuss five α values, all of which tell a story.

4.1.1 Focus on most rare types: $\alpha = 0$. To help demonstrate how tuning α affects the ranked list of dominant contributions, in Fig 1 we start with an allotaxonograph for $\alpha = 0$ for the first and second halves of *Pride and Prejudice*. Much of the present section describes allotaxonographs for PTD in general.

First, we describe the format of our allotaxonographs. The histogram on the left bins all pairs $(\log_{10} p_{\tau,1}, \log_{10} p_{\tau,2})$. The shape we see here is typical of comparisons between systems with heavy-tailed size-rank distributions. We rotate the axes so as not to privilege one system over the other, as this might lead to a false sense of an independent-dependent variable relationship [1,39].

We indicate counts per cell using the perceptually uniform colormap magma [40]. Because of the logarithmic scale, the cells start to separate for lower values of probability, corresponding to counts of 0, 1, 2, and so on. We emphasize that for each cell, the counts as indicated by color refer to the number of types (not tokens) that have a pair of probabilities in the two systems that fall within that cell. For example, the 2-gram "of the" is extremely common in both the first and second halves of *Pride and Prejudice*, and is slightly more common in the second half. There are no other 2-grams that have a pair of probabilities that are the same or close enough to "of the"'s pair, and hence it is represented with the lightest color indicating a count of 1.




Fig 5. Allotaxonograph comparing 2-gram usage in the first and second halves of Jane Austen's *Pride and Prejudice* using probability-turbulence divergence with $\alpha = \infty$, D_{∞}^P . The most common 2-grams now fully dominate the contributions as indicated in the ranked list and by the darker shading of the 2-grams at the top of the histogram. As for $\alpha = 0$ in Fig 1, the contours do not conform to the edges of the histogram. The dominant 2-grams for $\alpha = \infty$ largely comprise function words with the exception of 'Lady Catherine' and 'Miss Bingley' (characters most commonly referred to with a 2-gram).

<https://doi.org/10.1371/journal.pcsy.0000077.g005>

lines in Fig 1, which again are present because we anchor them at evenly spaced locations along the bottom of the main histogram, all correspond to a divergence of $\delta D_{0,\tau}^P = 0$. The dashed parts of the visible contour lines then collapse to the bottom zero point, showing how the exclusive types provide the only non-zero contributions to $\delta D_{0,\tau}^P$.

For all of our allotaxonographs, we place 10 contour lines on each half of the histogram's diamond. These contour lines are evenly spaced not by height but are rather anchored to the bottom axes of the main histogram where they are evenly spaced in logarithmic space.

While it may appear that we have omitted annotations internal to the histogram for convenient purposes of visualizing the histogram more cleanly, our annotations are intentional. Because individual 2-grams internal to the histogram will never dominate standard divergences, highlighting them would be badly misleading [14,41]. For our allotaxonographs, the annotations along the bottom of the histogram potentially fall into this trap: 'rejecting him' and 'escape all', each appearing once overall, are just two examples of tens of thousands of 2-gram *hapax legomena*. Such types may matter in aggregate but not individually.

Now, our allotaxonographs for probability-turbulence divergence must depart from those of rank-turbulence divergence because we have to accommodate instances of $\log_{10} p$ when $p = 0$. For ranks, types with 0 counts in one system—exclusive types—are assigned a tied rank for last place, necessarily a finite number. Here, on a logarithmic scale our




Fig 6. Allotaxonograph using probability-turbulence divergence to compare normalized 1-gram usage rates on two days of Twitter, 2020/03/12 and 2020/05/30—key dates in the US for the COVID-19 pandemic and the Black Lives Matter protests following George Floyd’s murder. We assess $\alpha = 1/3$ to be reasonably scale-equalizing for 1-grams. Details are the same as for Fig 1. The days are according to Coordinated Universal Time (UTC) and the 1-grams are those containing Latin characters found in English-language tweets [37,38]. See S4 Flipbook for the instrument’s variation as a function of α .

<https://doi.org/10.1371/journal.pcsy.0000077.g006>

exclusive types would have to be located on one axis at $\log_{10} p = -\infty$. We end the main histogram's domain for the lower value of $\log_{10} p_{\tau,i}$ such that $\log_{10} p_{\tau,i} > 0$. We then add lighter-colored regions to the bottom of both sides of the histogram, and locate $\log_{10} p_{\tau,i} = 0$ along their midlines. The transition is a discrete jump (we do not smoothly interpolate), and we connect the contour lines with a dotted line.

The last piece for the histogram is the list of balances at the bottom right. These summary quantities are intended to be both informative and diagnostic, and they have important if subtle differences. The first bars show the balance of total 2-gram counts which for our example of *Pride and Prejudice* is 50/50 by construction. The second and third balances refer to sizes of lexicons, and these are well balanced too. If we create a lexicon of all 2-grams for *Pride and Prejudice*, about 58.2% of them appear in the first half and 58.3% in the second. If we instead create separate lexicons of 2-grams for the two halves of *Pride and Prejudice*, the third line of the balances records the percentage of 2-grams that are exclusive to each half.

While it could be said that we are ultimately creating a simple 2-d histogram for a joint probability distribution with heavy tails, to our knowledge, there have been relatively few other attempts to do so [14,41]. As we have described them, we believe our histograms are crafted with a number of special details that make them well suited to their task.




Fig 7. Allotaxonograph using probability-turbulence divergence to compare normalized 2-gram usage ranks on two days of English-language Twitter, 2020/03/12 and 2020/05/30. Details are the same as for Fig 1 and 6. We see that the comparison of 2-gram distributions produces a different, broader histogram than that formed by 1-gram distributions (Fig 6). We choose $\alpha = 5/6$ to provide a balance of 2-grams across five orders of magnitude for non-zero probability. In contrast to the 1-gram version, the top 2-grams are more evenly distributed on both sides of the list. While some 2-grams are function words combined with the 1-grams we saw in Fig 6, meaningful 2-grams also appear ('Tom Hanks', 'toilet paper', 'George Floyd', and 'police brutality'). See S5 Flipbook for the instrument's variation as a function of α .

<https://doi.org/10.1371/journal.pcsy.0000077.g007>

The ranked list on the right maps the two dimensions of the histogram onto an ordered single dimension of divergence contributions, largest first. The left-right arrangement is solely done to be consistent with the histogram—all contributions are positive. The light gray numbers opposite each 2-gram (e.g., 44,185 $\Rightarrow$ 433.5 for 'the Parsonage') indicate the 2-gram's rank in the first and second halves of Pride and Prejudice (and in general, $\Omega^{(1)}$ and $\Omega^{(2)}$). We note that because we use tied ranks, types will have ranks that are either integers or half-integers [1].

In the ranked list, we add open triangles to types if they are exclusive to one system, corresponding to those appearing in the zero probability expansions of the histogram. For example, 'to Brighton' appears only in the first half of Pride and Prejudice, and 'the Parsonage' only in the second.

In reducing such a high-dimensional categorical space—where each unique type represents a dimension—we have first collapsed the data to a 2-d histogram, and then to a 1-d list. Being able to find the shape in the histogram to which we can apply an instance of probability-turbulence divergence gives us some suggestive proof in the pudding.

4.1.2 Squared-chord distances: $\alpha = 1/2$. Fig 2 shows how PTD changes as we jump from $\alpha = 0$ to $\alpha = 1/2$. The contour lines are now showing some alignment with the histogram. We now see 2-grams that appear in both halves in the




Fig 8. Allotaxonograph using probability-turbulence divergence to compare 3-gram usage ranks on two days of English-language Twitter, 2020/03/12 and 2020/05/30. Details are the same as for Fig 1 and 6. The histogram has broadened even further out from the 2-gram, and now is well suited to probability-turbulence divergence with the extreme of $\alpha = \infty$, D_{∞}^P . Fragmentary and meaningful 3-grams appear alongside each other, including 'World Health Organization' and 'black lives matter'. Social amplification is also apparent as 3-grams for highly retweeted tweets dominate the rank list. See S6 Flipbook for the instrument's variation as a function of α .

<https://doi.org/10.1371/journal.pcsy.0000077.g008>

top contributors list. For example, 'Miss Bingley' and 'Sir William'. That said, many of the top contributors remain exclusive types, from either half. The contour lines trim into the main histogram as they come down, and then still show some sharp inward turn for the exclusive types.

We point out that any distance measure of the squared-chord family (Sect 3.1.2) could be represented by a similar allotaxonograph. The contour lines and type ordering would remain the same. And the user would be given immediate visual feedback on the appropriateness of using such a measure for whatever probability distributions they are comparing.

4.1.3 Scale-equalizing fit with $\alpha = 3/4$. In Fig 3, we show an allotaxonograph with $\alpha = 3/4$ which we deem to be scale-equalizing. We see that the choice of $\alpha = 3/4$ generates a list with 2-grams from across the rare-to-common spectrum. The balanced darker shadings of annotations in the histogram add further support. Of course, $\alpha = 3/4$ is in itself rough—we are never going to be looking for many significant figures for the scale-equalizing value, but rather a small range.

4.1.4 L_q -distances: $\alpha = 1$. In Fig 4, we step up to the $\alpha = 1$ case. In comparison to $\alpha = 3/4$, we now start to see common 2-grams rise to the top of the contribution list. The 2-grams 'in the', 'had been', and 'to be' are the top 3, and 'Miss Bingley' has dropped to 5th. In moving from $\alpha = 1/2$ to $\alpha = 1$, we have crossed a threshold of scale-equalizing. Again,




Fig 9. Allotaxonograph using probability-turbulence divergence to compare tropical forest tree species abundance on Panama's Barro Colorado Island (BCI) for 5-year censuses completed in 1985 and 2015 [24]. The scale-equalizing choice of $\alpha = 1/3$ produces a set of dominant species reasonably well balanced across the abundance spectrum. See Ref [1] for the corresponding rank-turbulence divergence allotaxonographs. See S7 Flipbook for the instrument's variation as a function of α (Sect 5.3).

<https://doi.org/10.1371/journal.pcsy.0000077.g009>

$\alpha = 3/4$ is rough but for the user's interpretation of how all scales contribute, it provides a sufficiently balanced allotaxonograph.

4.1.5 Focus on most common types: $\alpha = \infty$. Finally, in Fig 5, we show the form of the allotaxonograph for the extreme of $\alpha = \infty$. The most common types now dominate (providing they do not have the same probability in each system, itself a very unlikely event for real systems). The contour lines are now parallel to the upper borders of the diamond.

4.2 Allotaxonographs for Twitter

For our second set of allotaxonographs, we compare two key dates of two major events through the lens of English-speaking Twitter: 2020/03/12, the date that COVID-19 became the major story in the United States, and 2020/05/30, five days after the murder of George Floyd in Minneapolis, Minnesota, by police officer Derek Chauvin. We compare day-scale normalized usage frequency distributions for 1-, 2-, and 3-grams for these two dates in Figs 6, 7, and 8. These datasets are far larger than Pride and Prejudice with approximately $10^{7.5}$ types and $10^{8.5}$ tokens. Nevertheless, producing an allotaxonograph on a standard present-day laptop takes on the order of minutes, which is almost entirely accounted for by loading and merging of the distributions.

We choose 2020/03/12 as a key date for the COVID-19 pandemic for several reasons. First and primarily, the World Health Organization (WHO) officially declared the COVID-19 outbreak to be a pandemic on 2020/03/11, a decision that was amplified immediately online but discussion of which most strongly appeared in the news and on Twitter on the following day.

The date of 2020/03/11 also saw a confluence of three major events that jolted the United States and dramatically elevated the story of the pandemic, all occurring tightly around a 15-minute period between 9 and 10 pm EDT (1 am to 2 am UTC). First, the National Basketball Association (NBA) abruptly suspended its season. The central event was the abandoning of a game just before tipoff between the Utah Jazz and Oklahoma City Thunder, upon the league learning that Rudy Gobert, a center for the Utah Jazz, had tested positive for COVID-19. Other players would test positive in the coming days and weeks, as would staff for teams and members of the media. Just a few days earlier, Gobert had joked with the media about his perception of institutional overreaction to the coronavirus by touching microphones at an interview.

Second, Tom Hanks announced that both he and his wife Rita Wilson had tested positive for COVID-19 while Hanks was working on a Baz Luhrmann film in Australia. Hanks was at the time the most high-profile figure known to have contracted COVID-19.

Third, President Donald Trump gave an Oval Office Address, the second of his presidency, “On the Coronavirus Pandemic.” The address marked a strong shift in Trump’s rhetoric regarding the danger of the COVID-19 outbreak. The main decision announced was the ban on travel from Europe to the US for 30 days, which was later clarified to not also mean a ban on trade. Futures on the US stock market dropped during the speech.

Combined, these disparate events were a major part of the COVID-19 pandemic becoming the dominant story for what would become weeks, months, and then years ahead.

The murder of George Floyd on 2020/05/25, Memorial Day in the US, precipitated Black Lives Matter protests and civilian-police confrontations in Minneapolis. The protests would grow over the following weeks, and begin to spread around the world. And, at least in the first week, George Floyd’s murder overtook coronavirus as the dominant story in the US [42].

With the above context in mind, we can sensibly examine the allotaxonographs of Figs 6, 7, and 8.

Our primary observation is that the three histograms vary considerably as we move through 1-, 2-, and 3-grams. The histograms broaden with increasing n , with the 3-gram histogram losing a scaling form and squaring up in the axes. The rapidly growing combinatorial possibilities of n -grams with increasing n mean that we see more and more exclusive n -grams as we look across the three allotaxonographs. For 1-grams, around 60% of each date’s lexicon are exclusive, for 2-grams, the percentage increases to around 70%, and for 3-grams we reach 80% (see the bottom of the three balance summaries in each allotaxonograph).

The maximum count per cell is 10^6 for 1-grams, 10^7 for 2-grams, and 10^8 for 3-grams. The cells with the most n -grams are of course the *hapax legomena*—the bottommost two cells in the histogram—those n -grams which appear once on one of the dates and not at all on the other.

To obtain good balance for the most dominant n -grams, we select $\alpha = 1/3$, $5/6$, and ∞ . Different kinds of terms dominate depending on n with ‘coronavirus’, ‘the coronavirus’, and ‘tested positive for’ leading on 2020/03/12, and ‘Minneapolis’, ‘George Floyd’, and ‘of George Floyd’ at the top on 2020/05/30.

Because social amplification is encoded in Twitter’s data stream through retweets, dominant 2-grams and especially 3-grams are liable to belong to the most retweeted messages of the day, and may lead to some variation in the dominant n -grams. (By contrast, we do not have a measure of popularity of individual phrases or sentences within Pride and Prejudice with just the bare text.) For example, ‘toilet paper’ and ‘World Health Organization’ appear as dominant 2-grams and 3-grams but none of their five distinct 1-grams are near the top of the ranked list in Fig 6. On the other hand, some dominant 1-grams may be used in diverse 2-grams and 3-grams and thus may not appear in the ranked lists for 2-grams and 3-grams. Examples from Fig 6 are ‘antifa’ and ‘Breonna’.

For all three n -gram comparisons of these two dates on Twitter, we provide [S4](#), [S5](#), and [S6](#) [25] Flipbooks (Sect 5.3). Readers may use these to easily explore how the choice of α affects the fit for the contour lines in the histogram and the ordering of which n -grams dominate probability-turbulence divergence.

4.3 Allotaxonographs for Barro Colorado Island

We include one final allotaxonograph from an entirely different field of research, ecology. In [Fig 9](#) we show a probability-turbulence divergence allotaxonograph for tree species abundance on Barro Colorado Island for censuses completed in 1985 and 2015. This example also shows how allotaxonographs can be used to inspect how well divergence measures perform for datasets that are much smaller than our examples from literature and Twitter. The species that dominates the overall divergence score is one that has diminished in abundance, *Piper cordulatum* [43–46]. In Ref [1], we compared these distributions with rank-turbulence divergence, and the overall orderings of dominant species are broadly consistent.

In [S7 Flipbook](#), we show how the dominant contributions of species vary as a function of α .

5 Data, code, and supplementary material

5.1 Datasets

Pride and Prejudice: We sourced a plain text version of Jane Austen's *Pride and Prejudice* from Project Gutenberg (<http://www.gutenberg.org/ebooks/1342>).

Normalized n -gram usage frequency on Twitter: We collected around 10% of all tweets sent on these dates based on Coordinated Universal Time (UTC), meaning they covered 8:00:00 pm to 7:59:59 pm Eastern Daylight Time (EDT) and 4:00:00 pm to 3:59:59 pm Pacific Daylight Time (PDT). We provide historical access to the top 10^6 1-grams, 2-grams, and 3-grams across more than 100 languages as part of our Storywrangler for Twitter project [38].

Species abundance on Barro Colorado Island: We accessed the dataset for BCI censuses performed roughly every 5 years over 35 years through the online repository described in Ref [24].

5.2 Code for divergence calculation and rendering allotaxonographs

All scripts and documentation reside on [GitLab](#) at <https://gitlab.com/compstorylab/allotaxonometer/>. For the present paper, we wrote the scripts to generate the allotaxonographs in MATLAB. We originally produced all figures and flipbooks using MATLAB Version R2020a, while endeavoring to keep the code functional with future versions. As is, the script also generates allotaxonographs for rank-turbulence divergence [1], the present paper's probability-turbulence divergence, and a parametrized symmetric entropy divergence that generalizes Jensen-Shannon divergence. We welcome ports to other languages.

5.3 Supplementary material

We list all relevant sites and links in [Table 3](#). The base site <https://compstorylab.org/allotaxonomy/> provides a home for our work on allotaxonomy. We use [Zenodo](#) to store flipbooks [25], and we give direct links to all flipbooks in [Table 3](#). Flipbooks are designed to be examined with a PDF viewer in single-page mode. Some PDF viewers within browsers do not accommodate single-page mode, and we recommend using an alternate browser or downloading flipbooks and viewing locally. The same flipbooks are available on the base site and in our [GitLab repository](#).

6 Concluding remarks

We have defined, analyzed, and demonstrated the use of probability-turbulence divergence as an instrument of allotaxonomy, both analytically and through the diagnostic visualizations afforded by our allotaxonographs. As the probability-based analog of our rank-turbulence divergence, the instrument is able to perform well when comparing heavy-tailed size-rank distributions of type frequencies. We emphasize the importance of creating allotaxonographs for a range of values of

α (see Eq 23). From there, the user can assess whether or not a scale-equalizing value of α exists. We have shown further that probability-turbulence divergence generalizes a range of existing probability-based divergences, either matching in exact form or equating in how types are ordered by type contribution.

While we view rank-turbulence divergence as our most general, interpretable instrument, for systems in which probabilities (or rates) of types occurring are well defined—and the resulting distributions involved are heavy-tailed—probability-turbulence divergence provides a more nuanced instrument.

We also favor divergences which compare distributions in as transparent a way as possible. To that end, we have made the core of probability-turbulence divergence a simple difference of powers of probabilities (Eq 6). By contrast, we view some divergences as being problematic in being overly constructed. We venture that Jensen-Shannon divergence (JSD), which we ourselves have used elsewhere, is one such instrument. The creation of an artificial mixed distribution is a contrivance we avoid here, and is perhaps indicative of taking information theory too far [47].

In our experience, we have also found that the visual information delivered by our allotaxonographs, especially in their coupling of histograms and ranked lists, has been essential to working effectively with divergences of all kinds.

One caution we make is that in the examples we have explored in the present paper, we have taken distributions as they are. That is, we have not contended with issues of sub-sampling and missing tail data [48]. We can say that types appearing with a high rate (e.g., common n -grams on Twitter) will not be affected by accessing more data, as they are well-estimated rates. In our paper on rank-turbulence divergence, we examined how truncation of distributions affects allotaxonographs, and such an approach is always available for any divergence.

Finally, in our present paper and in Ref [1], we have so far made choices of α based on inspection of the relevant histogram. A clear next step is to find ways to determine a scale-equalizing α for any given pair of distributions, and to do so only when sufficiently robust scaling is apparent. From a storytelling perspective, we are concerned with finding a scale-equalizing α that returns a ranked list of distinguishing types for two distributions such that the list comprises a balance of types from across the full range of observed probabilities [49]. We have performed some preliminary work for such an optimization, and note here that simple regression is made difficult by the overwhelming weight of rare types relative to common ones.

Table 3. List of supporting material sites and direct links for probability-turbulence divergence. Flipbooks show how probability-turbulence divergence behaves as α ranges from 0 to ∞ for the three studies in this paper: Pride and Prejudice, Twitter, and Barro Colorado Island.

Sites for supplementary material:
Allotaxonomy home: https://compstorylab.org/allotaxonomy/
GitLab repository: https://gitlab.com/compstorylab/allotaxonometer/
Zenodo record: https://doi.org/10.5281/zenodo.15006457/ [25]
Direct links to flipbooks on Zenodo:
S1 Flipbook—Pride and Prejudice, first half versus second half, 1-grams: https://zenodo.org/records/15006457/files/allotaxonometer-flipbook-1-probability-divergence-pride-and-prejudice-1-grams.pdf
S2 Flipbook—Pride and Prejudice, first half versus second half, 2-grams: https://zenodo.org/records/15006457/files/allotaxonometer-flipbook-2-probability-divergence-pride-and-prejudice-2-grams.pdf
S3 Flipbook—Pride and Prejudice, first half versus second half, 3-grams: https://zenodo.org/records/15006457/files/allotaxonometer-flipbook-3-probability-divergence-pride-and-prejudice-3-grams.pdf
S4 Flipbook—Twitter, 2020-03-12 versus 2020-05-30, 1-grams: https://zenodo.org/records/15006457/files/allotaxonometer-flipbook-4-probability-divergence-twitter-1-grams.pdf
S5 Flipbook—Twitter, 2020-03-12 versus 2020-05-30, 2-grams: https://zenodo.org/records/15006457/files/allotaxonometer-flipbook-5-probability-divergence-twitter-2-grams.pdf
S6 Flipbook—Twitter, 2020-03-12 versus 2020-05-30, 3-grams: https://zenodo.org/records/15006457/files/allotaxonometer-flipbook-6-probability-divergence-twitter-3-grams.pdf
S7 Flipbook—Barro Colorado Island, 1985 versus 2015 census, species counts: https://zenodo.org/records/15006457/files/allotaxonometer-flipbook-7-probability-turbulence-divergence-barro-colorado-island.pdf
https://doi.org/10.1371/journal.pcsy.0000077.t003

Supporting information

S1 Flipbook. Pride and Prejudice, first half versus second half, 1-grams, for $\vec{\alpha} = \{0, \frac{1}{12}, \frac{2}{12}, \dots, \frac{18}{12}, 2, 3, 5, 10, \infty\}$.
(PDF)

S2 Flipbook. Pride and Prejudice, first half versus second half, 2-grams, for $\vec{\alpha} = \{0, \frac{1}{12}, \frac{2}{12}, \dots, \frac{18}{12}, 2, 3, 5, 10, \infty\}$.
(PDF)

S3 Flipbook. Pride and Prejudice, first half versus second half, 3-grams, for $\vec{\alpha} = \{0, \frac{1}{12}, \frac{2}{12}, \dots, \frac{18}{12}, 2, 3, 5, 10, \infty\}$.
(PDF)

S4 Flipbook. Twitter, 2020-03-12 versus 2020-05-30, 1-grams, for $\vec{\alpha} = \{0, \frac{1}{12}, \frac{2}{12}, \dots, \frac{18}{12}, 2, 3, 5, 10, \infty\}$.
(PDF)

S5 Flipbook. Twitter, 2020-03-12 versus 2020-05-30, 2-grams, for $\vec{\alpha} = \{0, \frac{1}{12}, \frac{2}{12}, \dots, \frac{18}{12}, 2, 3, 5, 10, \infty\}$.
(PDF)

S6 Flipbook. Twitter, 2020-03-12 versus 2020-05-30, 3-grams, for $\vec{\alpha} = \{0, \frac{1}{12}, \frac{2}{12}, \dots, \frac{18}{12}, 2, 3, 5, 10, \infty\}$.
(PDF)

S7 Flipbook. Barro Colorado Island, 1985 versus 2015 census, species counts, for $\vec{\alpha} = \{0, \frac{1}{12}, \frac{2}{12}, \dots, \frac{18}{12}, 2, 3, 5, 10, \infty\}$.
(PDF)

Author contributions

Conceptualization: Peter Sheridan Dodds.

Data curation: Peter Sheridan Dodds, Joshua R. Minot, Michael V. Arnold, Thayer Alshaabi, Jane Lydia Adams, Andrew J. Reagan, Christopher M. Danforth.

Formal analysis: Peter Sheridan Dodds.

Funding acquisition: Peter Sheridan Dodds, Christopher M. Danforth.

Investigation: Peter Sheridan Dodds.

Methodology: Peter Sheridan Dodds.

Project administration: Peter Sheridan Dodds.

Resources: Peter Sheridan Dodds.

Software: Peter Sheridan Dodds.

Supervision: Peter Sheridan Dodds, Christopher M. Danforth.

Validation: Peter Sheridan Dodds.

Visualization: Peter Sheridan Dodds.

Writing – original draft: Peter Sheridan Dodds.

Writing – review & editing: Peter Sheridan Dodds, Joshua R. Minot, Michael V. Arnold, Thayer Alshaabi, Jane Lydia Adams, Andrew J. Reagan, Christopher M. Danforth.

References

1. Dodds PS, Minot JR, Arnold MV, Alshaabi T, Adams JL, Dewhurst DR, et al. Allotaxonomy and rank-turbulence divergence: a universal instrument for comparing complex systems. *EPJ Data Sci.* 2023;12(1). <https://doi.org/10.1140/epjds/s13688-023-00400-x>
2. Zipf GK. Human behaviour and the principle of least-effort. Cambridge, MA: Addison-Wesley; 1949.
3. Simon HA. On a class of skew distribution functions. *Biometrika.* 1955;42(3/4):425. <https://doi.org/10.2307/2333389>
4. Newman M. Power laws, Pareto distributions and Zipf's law. *Contemporary Physics.* 2005;46(5):323–51. <https://doi.org/10.1080/00107510500052444>
5. Clauset A, Shalizi CR, Newman MEJ. Power-law distributions in empirical data. *SIAM Rev.* 2009;51(4):661–703. <https://doi.org/10.1137/070710111>
6. Barabasi A, Albert R. Emergence of scaling in random networks. *Science.* 1999;286(5439):509–12. <https://doi.org/10.1126/science.286.5439.509> PMID: 10521342
7. Mitchell M. Complexity: a guided tour. New York, NY: Oxford University Press; 2009.
8. Bettencourt LMA, Lobo J, Helbing D, Kühnert C, West GB. Growth, innovation, scaling, and the pace of life in cities. *Proc Natl Acad Sci.* 2007;104(17):7301–6.
9. Magurran AE. Measuring biological diversity. John Wiley & Sons; 2013.
10. Scholz CH. The mechanics of earthquakes and faulting. 3rd ed. Cambridge University Press; 2019.
11. Deza MM, Deza E. Dictionary of distances. Elsevier; 2006.
12. Cha SH. Comprehensive survey on distance/similarity measures between probability density functions. *International Journal of Mathematical Models and Methods in Applied Sciences.* 2007;1:300–7.
13. Cichocki A, Amari S. Families of Alpha- Beta- and Gamma- divergences: flexible and robust measures of similarities. *Entropy.* 2010;12(6):1532–68. <https://doi.org/10.3390/e12061532>
14. Monroe BL, Colaresi MP, Quinn KM. Fightin' words: lexical feature selection and evaluation for identifying the content of political conflict. *Polit anal.* 2008;16(4):372–403. <https://doi.org/10.1093/pan/mpn018>
15. Pechenick EA, Danforth CM, Dodds PS. Is language evolution grinding to a halt? The scaling of lexical turbulence in English fiction suggests it is not. *Journal of Computational Science.* 2017;21:24–37. <https://doi.org/10.1016/j.jocs.2017.04.020>
16. Dice LR. Measures of the amount of ecologic association between species. *Ecology.* 1945;26(3):297–302. <https://doi.org/10.2307/1932409>
17. Sorensen T. A method of establishing groups of equal amplitude in plant sociology based on similarity of species content and its application to analyses of the vegetation on Danish commons. *Videnski Selskab Biologiske Skrifter.* 1948;5:1–34.
18. Van Rijsbergen CJ. Information retrieval. 2nd ed. Butterworth-Heinemann; 1979.
19. Rényi A. On measures of entropy and information. In: *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability.* 1961.
20. Tsallis C. Nonextensive statistical mechanics and thermodynamics: historical background and present status. *Nonextensive statistical mechanics and its applications.* Springer; 2001. p. 3–98.
21. Keylock CJ. Simpson diversity and the Shannon–Wiener index as special cases of a generalized entropy. *Oikos.* 2005;109(1):203–7. <https://doi.org/10.1111/j.0030-1299.2005.13735.x>
22. Hill MO. Diversity and evenness: a unifying notation and its consequences. *Ecology.* 1973;54(2):427–32. <https://doi.org/10.2307/1934352>
23. Jost L. Entropy and diversity. *Oikos.* 2006;113(2):363–75. <https://doi.org/10.1111/j.2006.0030-1299.14714.x>
24. Condit R, et al. Complete data from the Barro Colorado 50-ha plot: 423617 trees, 35 years, v3. *DataONE Dash.* 2019.
25. Dodds PS. Flipbooks for probability turbulence divergence. 2025. <https://doi.org/10.5281/zenodo.15006457>
26. Looman J, Campbell JB. Adaptation of Sorensen's K (1948) for estimating unit affinities in prairie vegetation. *Ecology.* 1960;41(3):409–16. <https://doi.org/10.2307/1933315>
27. Stigler SM. Stigler's law of eponymy*. *Trans NY Acad Sci.* 1980;39(1 Series II):147–57. <https://doi.org/10.1111/j.2164-0947.1980.tb02775.x>
28. Sasaki Y. The truth of the f-measure. 2007.
29. Hellinger E. Neue Begründung der Theorie quadratischer Formen von unendlichvielen Veränderlichen. *Journal für die reine und angewandte Mathematik.* 1909;1909(136):210–71. <https://doi.org/10.1515/crll.1909.136.210>
30. Matusita K. Decision rules, based on the distance, for problems of fit, two samples, and estimation. *Ann Math Statist.* 1955;26(4):631–40. <https://doi.org/10.1214/aoms/1177728422>
31. Gavin DG, Oswald WW, Wahl ER, Williams JW. A statistical approach to evaluating distance metrics and analog assignments for pollen records. *Quat res.* 2003;60(3):356–67. [https://doi.org/10.1016/s0033-5894\(03\)00088-7](https://doi.org/10.1016/s0033-5894(03)00088-7)
32. Minkowski H. *Geometrie der Zahlen.* Leipzig: B. G. Teubner; 1896.
33. Newman DJ. *Normed linear spaces.* Springer-Verlag; 1973.
34. Krause EF. *Taxicab geometry: an adventure in non Euclidean geometry.* Dover Books on Mathematics: Dover Publications; 1987.

35. Deza M, Laurent M. Geometry of cuts and metrics. Springer; 1996.
36. Austen J. Pride and prejudice. Broadview Press; 2001.
37. Alshaabi T, Dewhurst DR, Minot JR, Arnold MV, Adams JL, Danforth CM, et al. The growing amplification of social media: measuring temporal and social contagion dynamics for over 150 languages on Twitter for 2009-2020. EPJ Data Sci. 2021;10(1):15. <https://doi.org/10.1140/epjds/s13688-021-00271-0> PMID: 33816048
38. Alshaabi T, Adams JL, Arnold MV, Minot JR, Dewhurst DR, Reagan AJ, et al. Storywrangler: a massive exploratorium for sociolinguistic, cultural, socioeconomic, and political timelines using Twitter. Sci Adv. 2021;7(29):eabe6534. <https://doi.org/10.1126/sciadv.abe6534> PMID: 34272243
39. Bergstrom CT, West JD. Why scatter plots suggest causality, and what we can do about it. 2018. <https://arxiv.org/abs/1809.09328>
40. Liu Y, Heer J. Somewhere over the rainbow: an empirical assessment of quantitative colormaps. In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. ACM; 2018. p. 598.
41. Kessler JS. Scattertext: a browser-based tool for visualizing how corpora differ. arXiv preprint 2017. <https://doi.org/10.48550/arXiv.1703.00565>
42. Dodds PS, Minot JR, Arnold MV, Alshaabi T, Adams JL, Reagan AJ, et al. Computational timeline reconstruction of the stories surrounding Trump: Story turbulence, narrative control, and collective chronopathy. PLoS One. 2021;16(12):e0260592. <https://doi.org/10.1371/journal.pone.0260592> PMID: 34879105
43. Trelease W. The Piperaceae of Panama. Systematic plant studies: mainly tropical America. 1927.
44. Standley PC. The flora of Barro Colorado Island, Panama. Smithsonian Miscellaneous Collections. 1927.
45. Thies W, Kalko EKV. Phenology of neotropical pepper plants (Piperaceae) and their association with their main dispersers, two short-tailed fruit bats, *Carollia perspicillata* and *C. castanea* (Phyllostomidae). Oikos. 2004;104(2):362–76. <https://doi.org/10.1111/j.0030-1299.2004.12747.x>
46. Andrade TY, Thies W, Rogeri PK, Kalko EKV, Mello MAR. Hierarchical fruit selection by Neotropical leaf-nosed bats (Chiroptera: Phyllostomidae). J Mammal. 2013;94(5):1094–101. <https://doi.org/10.1644/12-mamm-a-244.1>
47. Shannon C. The bandwagon (Edtl.). IRE Trans Inf Theory. 1956;2(1):3. <https://doi.org/10.1109/tit.1956.1056774>
48. Haegeman B, Hamelin J, Moriarty J, Neal P, Dushoff J, Weitz JS. Robust estimation of microbial diversity in theory and in practice. ISME J. 2013;7(6):1092–101. <https://doi.org/10.1038/ismej.2013.10> PMID: 23407313
49. Williams JR, Bagrow JP, Danforth CM, Dodds PS. Text mixing shapes the anatomy of rank-frequency distributions. Phys Rev E Stat Nonlin Soft Matter Phys. 2015;91(5):052811. <https://doi.org/10.1103/PhysRevE.91.052811> PMID: 26066216