

Umbrella-Word Inflation in Research-Project Documents: A Computational Text Analysis of diversity, sustainable, multi- and collaboration

Kohei Takeda

koh@keio.jp

Keio University: Keio Gijuku Daigaku <https://orcid.org/0009-0003-7897-1405>

Research Article

Keywords: buzzwords, research funding, text mining, decoupling, semantic change, institutional isomorphism

Posted Date: October 23rd, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11461965/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Umbrella-Word Inflation in Research-Project Documents: A Computational Text Analysis of *diversity*, *sustainable*, *multi-* and *collaboration*

Kohei Takeda

ZEN University, Keio University

Abstract

We measure how often vague, broadly applicable words, which we call umbrella words, appear in research-project documents, and we test whether their growth carries costs. From the October 2026 OpenAlex snapshot we built four corpora: 84,750 project objectives from the EU Framework Programmes (FP7, Horizon 2020 and Horizon Europe), 369,546 award abstracts from the US National Science Foundation (NSF), 186,647 abstracts from the US National Institutes of Health (NIH) and 1,290,000 journal-article abstracts. Participant data for 36,177 FP7 and Horizon 2020 projects were linked to their objectives. We distinguish normative, organisational and evaluative umbrella words. The share of NSF abstracts using *diversity* rose from 3.5% in 1987-91 to 26.1% in 2021-25 and the share using *sustainab-* from 0.09% to 10.2%, while the control word *experiment* went from 21.3% to 22.6%. Normative and evaluative families outgrew control words and frequency-matched ordinary words also after adjusting for the longer abstracts of recent years; in NSF abstracts *collaborat-* and organisational *multi-* grew no faster than the neutral noun *approach*. Most families spread across fields (for *sustainab-*, bias-corrected field concentration fell from 1.00 to 0.42), as boundary objects would. In NSF abstracts *diversity* also lost its objects: biological collocates fell from 76% to 42% and goal verbs such as *increase* and *promote* rose from 2% to 27%, a shift that word embeddings confirm; in articles *diversity* stayed biological. One SD of umbrella density went with 0.03 to 0.10 SD lower lexical specificity in the rest of the text, while word concreteness moved the other way in NSF abstracts. Collaboration vocabulary appeared in 13.6% of single-beneficiary EU projects, and within collaborative schemes it rose modestly with consortium size and international reach. Associations with outcomes were small next to budget and team size and differed in sign across corpora. Vocabulary jumped at EU programme boundaries within fixed scheme types. In 2025, DEI-related words in NSF abstracts fell by 53% to 88%, also within continuing subfields, while the topic mix of awards moved within its usual range. Since 2023, documents with LLM-typical words have carried more umbrella words. The clearest costs of umbrella-word inflation are its dependence on funders and politics and, for *diversity* in NSF abstracts, a shift from objects to goals.

Keywords: buzzwords; research funding; text mining; decoupling; semantic change; institutional isomorphism

1. Introduction

Project documents in international research and development cooperation are written to win funding, and they also tell other people what a project will do. Many readers of such documents believe that

words such as *diversity*, *sustainable*, *multi-stakeholder*, *multidisciplinary* and *collaboration* have become far more common over the past few decades. The size of this increase in research-funding texts, and what it costs, has rarely been measured on large corpora.

We call these words umbrella words. They are broad terms with a positive normative or institutional charge, used across many fields and contexts, and in a given document they need not point to a particular object, procedure or structure. The definition does not assume that the words are empty. Whether they still point to objects and structures is what several of our hypotheses test. Umbrella words can act as a shared vocabulary across disciplines and sectors. They can also circulate as markers of legitimacy with little specific content.

The words are not all of one kind, and we distinguish three types.

- Normative or substantive terms that began as concepts in particular fields or that name social goals: *sustainability*, *diversity*, *resilience*, *inclusion*, *empowerment*, *societal*. For these the main question is whether they keep their objects as they spread (H2).
- Organisational or process terms that describe how research is organised: *collaboration*, organisational *multi-*, *interdisciplinary*, *co-creation*, *partnership*, *cooperation*, *stakeholder*, *capacity building*. These make claims about structure that can be checked against data on consortia and authorship (H5).
- Evaluative policy terms of funding programmes: *innovation*, *impact*, *excellence*, *transformative*, *holistic*. They overlap with the promotional language studied by Vinkers et al. (2015) and Millar et al. (2022, 2023). We keep them in the index because funders use them as programme categories. Generic promotional adjectives (*novel*, *unprecedented*, *cutting-edge* and similar) are tracked separately as a reference group, since prior work has already measured them.

We ask two questions.

- RQ1 (scale): How much, and how fast, has the use of umbrella words grown? Does the growth exceed changes in vocabulary as a whole?
- RQ2 (costs): Is the growth accompanied by loss of concrete referents, more uniform documents, lower specificity, a gap between claims and project structure, a lack of association with outcomes, and dependence on funders' programmes, politics and LLM-assisted writing?

The study makes three contributions. It measures the growth of umbrella words against control words and against frequency-matched ordinary words in about two million documents that combine EU and US funding texts with journal articles. It turns costs discussed mainly in qualitative work, such as dilution and decoupling, into testable indicators, and it keeps apart indicators that a legitimate boundary object would also produce. It also treats the 2025 political change in the United States as a natural experiment, which indicates that the vocabulary can change much faster than the research portfolio.

2. Literature and theoretical framework

2.1 Buzzwords in development discourse

Cornwall (2007) analysed the central terms of development discourse, among them *participation*, *empowerment*, *partnership* and *sustainability*, as "buzzwords" and "fuzzwords". Because readers can attach different meanings to them, such terms create an appearance of consensus and work as a

"language of evasion" (see also Cornwall and Eade 2010). Bensaude-Vincent (2014) used *public engagement in science* to examine the politics of buzzwords where technoscience, markets and society meet. Redclift (2005) traced how the surface consensus around *sustainable development* broke into several parallel discourses.

2.2 Quantitative studies

Moretti and Pestre (2015) applied computational reading to World Bank annual reports from 1946 to 2012. They reported a decline in concrete referents and the rise of terms such as *governance*. In scientific writing, Vinkers et al. (2015) found a large increase in positive words (*novel, robust, innovative*) in PubMed abstracts. Millar et al. (2022) documented the growth of promotional language in abstracts of funded NIH grant applications, and Millar et al. (2023) found the same trend in the publications that NIH-funded research produced. Markowitz (2019) related the style of NSF abstracts to the dollar amount of the award: longer abstracts, abstracts with fewer common words and abstracts written with more verbal certainty received larger awards. Plavén-Sigra et al. (2017) reported that the readability of scientific abstracts has declined over time. Kobak et al. (2025) identified excess vocabulary typical of large language models (LLMs) in biomedical abstracts after 2023. Mehta and Jena (2025) found that diversity-related words in NIH award abstracts fell by about a quarter between January 2024 and June 2025. Renewed grants deleted these words at about ten times the rate of control words.

Two strands of computational linguistics measure meaning change directly. Hamilton et al. (2016) aligned word embeddings trained on successive periods and showed that frequent words change meaning more slowly (the law of conformity) and polysemous words faster (the law of innovation). Luo et al. (2019) modelled semantic bleaching of English intensifiers through falling similarity to the source meaning, rising similarity to a bleached prototype such as *very*, and rising productivity across the words they combine with.

These studies documented the growth of promotional and positive language and the drift of particular words. Few have tested, in large collections of research-funding documents, whether broad normative words lose their association with particular fields and objects, or whether organisational words correspond to how projects are actually organised. This study addresses that gap.

2.3 Theoretical tools

Vague words can connect different communities as boundary objects, which keep a common identity across fields while each field gives them a local meaning (Star and Griesemer 1989). They can also become empty signifiers that have lost specific content (Laclau 1996). The two readings predict different patterns. Both predict that a word spreads across fields. Only the empty-signifier reading predicts that the word also loses its concrete objects within the documents that use it. Brand and Jax (2007) argued that *resilience* drifted from a descriptive ecological concept towards a vague boundary object. Ahmed (2012) analysed university diversity statements as non-performative speech acts, statements that do not bring about what they name; a shift from describing diversity to promising to increase it fits that account. We therefore treat cross-field diffusion and loss of referents as separate indicators.

Organisational theory offers two further concepts. Decoupling describes formal structures adopted to fit institutional expectations but separated from actual activity (Meyer and Rowan 1977). Institutional isomorphism describes how organisations converge on the same norms (DiMaggio and Powell 1983).

The value of interdisciplinarity itself is contested. Interdisciplinary proposals have lower funding success (Bromham et al. 2016), and interdisciplinary researchers publish less but are cited more (Leahey et al. 2017). Jacobs and Frickel (2009) describe the evidence on its benefits as mixed. Leahey et al. imply that publication volume and citation impact can move in opposite directions, so we report both kinds of outcome. Finally, every grant corpus contains successful applications only. Suppose reviewers reward umbrella words. Then proposals with weaker underlying potential cross the funding line more often when they use them. Among funded projects, umbrella density would be associated with lower outcomes even if it had no effect of its own (selection on a collider).

3. Hypotheses

From this literature we derive seven hypotheses. H1 answers RQ1, and H2 to H7 answer RQ2.

Table 1. Hypotheses, the cost each would indicate, and main sources

Hypothesis	Statement	Cost it would indicate	Main sources
H1	The share of documents using umbrella words grows much faster than the share using control words and frequency-matched ordinary words, also after adjusting for document length	Inflation of the vocabulary	Vinkers et al. 2015; Millar et al. 2022, 2023
H2a	Umbrella words spread across fields (falling field concentration)	None by itself: also expected of a boundary object	Star and Griesemer 1989
H2b	Within documents, umbrella words lose their concrete objects and shift towards goal statements	Loss of information; non-performative use	Cornwall 2007; Brand and Jax 2007; Laclau 1996; Ahmed 2012
H3	Documents become more similar to one another and contain more boilerplate, and umbrella words carry part of the similarity	Homogenisation and templating	DiMaggio and Powell 1983
H4	Documents dense in umbrella words are less specific in the rest of their text	Thinner substantive content	Moretti and Pestre 2015
H5	Organisational vocabulary is common in single-actor documents and does not scale with the size of the collaborative structure once funding scheme and international participation are held constant	Decoupling (ceremonial language)	Meyer and Rowan 1977
H6	Among funded projects, associations between	No value as a	Markowitz

Hypothesis	Statement	Cost it would indicate	Main sources
	umbrella density and outcomes (publication volume, citation impact) are small next to structural inputs such as budget and team size, and are not consistent in sign across corpora	signal of quality	2019; Leahey et al. 2017; Millar et al. 2023
H7	The vocabulary changes discontinuously with funders' programmes and politics, and since 2023 it co-occurs with words typical of LLM-assisted writing	Dependence on outside actors; reversibility	Ahmed 2012; Mehta and Jena 2025; Kobak et al. 2025

H5 and H6 are stated in terms of magnitude and consistency. With several hundred thousand documents almost any coefficient is statistically different from zero, so a test against a strict null would say little. We judge them by the size of each association next to that of the structural variables in the same model, and by whether its sign agrees across corpora. For H6 the selection on funding described in Section 2.3 means that a negative association among funded projects cannot be read as a negative effect.

4. Data

4.1 Sources

The main source is the complete OpenAlex snapshot released on 14 October 2026 and published on Amazon S3 under CC0 (Priem et al. 2022). Its awards entity (17.9 million records) incorporates open data from national funders, including CORDIS, NSF Award Search and NIH ExPORTER. From the works entity (more than 460 million records) we drew 500 parquet row groups at random, with probability proportional to record count, which yielded 4.53 million works. We kept English-language journal articles with an abstract of at least 200 characters.

The works files are partitioned by the date a record was last updated, so a row group can contain records from the same ingestion batch, publisher or era. Within a publication year the sample is therefore a cluster sample and its field mix can differ from that of all OpenAlex articles. Section 5.3 describes how the field-concentration measures guard against shifts in the field mix.

Network restrictions in the computing environment blocked access to CORDIS itself (cordis.europa.eu and data.europa.eu) and to the World Bank API. Because the CORDIS records in OpenAlex begin in 2007, the EU Framework Programmes FP1 to FP6 (1984 to 2006) are not covered. We study long-run change with NSF abstracts from the 1980s onwards and with journal articles.

4.2 Corpora

Four corpora are analysed: EU project objectives, NSF award abstracts, NIH project abstracts and journal-article abstracts (Table 2). Participant data for FP7 and Horizon 2020 projects were linked to the EU objectives for H5.

Table 2. Corpora

Corpus	Document type	Documents	Years in the data	Years analysed	Median tokens
EU	Project objectives, FP7, Horizon 2020 and Horizon Europe	84,750	2007-2027	2008-2026	277
EU, collaborative schemes	RIA, IA, CSA, CP, ITN, RISE, DN and similar	30,515	2007-2027	2008-2026	277
EU, individual schemes	ERC, individual MSCA fellowships, SME Instrument, EIC Accelerator	49,650	2007-2027	2008-2026	278
EU, scheme not classified	Funding-scheme code missing or not recognised	4,585	2007-2027	2008-2026	258
NSF	Award abstracts, US National Science Foundation	369,546	1981-2026	1987-2025	275
NIH	Project abstracts, US National Institutes of Health (random sample, at most 12,000 a year)	186,647	1980-2026	1985-2025	83
Articles	Journal-article abstracts (at most 30,000 a year)	1,290,000	1980-2025	1980-2025	147
CORDIS structure	Participants and countries of FP7 projects and of Horizon 2020 projects listed by 2017, linked by grant number	36,177	2007-2017	2007-2017	not applicable

Notes: Years refer to the start year of a project or the publication year of an article. The EU data contain 1,158 projects with a scheduled start in 2027, signed before the snapshot; these and the 2007 projects of the first FP7 calls are outside the analysed years. Two further corpora were collected and are not analysed: 28,507 Gates Foundation grant purposes, which are single sentences (median 20 tokens), and 36,431 EU-funded articles, kept in the corpus files for later work.

EU duplicates were removed by grant number, with priority CORDIS over ERC over Crossref records. Schemes were classified from funding-scheme codes as collaborative, individual or not classified; the code is in `build_award_corpora.py`. From NSF abstracts we removed the sentence attached to every award from about 2014 ("This award reflects NSF's statutory mission ... broader impacts review criteria"). The NIH corpus is a random sample of at most 12,000 abstracts per start year. The NIH records in OpenAlex contain few abstracts for the early years, and the NIH corpus has between 318 and 661 abstracts a year in 1985-1994. For articles, the year-balanced sample keeps, within each publication year, the 30,000 articles with the lowest hash of their OpenAlex identifier, which works as a pseudo-

random draw. Years before 1990 have fewer articles than the cap, with a minimum of 17,082 in 1980. Abstracts that are identical within a corpus were kept once. The CORDIS structure files come from a public GitHub mirror of the CORDIS project files (spaghetti-open-data/eu-research-funding-network). The Horizon 2020 file is a 2017 snapshot, so it covers projects signed by then; later Horizon 2020 projects and all Horizon Europe projects have no participant data. OpenAlex lists only the coordinator for CORDIS awards and cannot fill the gap.

We also used the concreteness ratings of Brysbaert et al. (2014; 39,954 words) and the published excess-vocabulary list of Kobak et al. (2025), from which the LLM-typical words were taken (Section 5.1). All analyses of articles cover 1980-2025. The previous version of this paper started some article series in 1975.

5. Methods

5.1 Lexicon and tokenisation

Texts were lower-cased. Umbrella words were defined as regular-expression families; Appendix A lists every expression, and Appendix B lists the dictionaries used for collocates. Tokens are runs of letters that may contain hyphens or apostrophes, and runs of digits. The four core families are the subject of the research question: *diversity/diverse*, *sustainab-*, organisational *multi-* and *collaborat-*. With 15 further families from the literature on development discourse and research policy they form the 19-family umbrella index (UI), divided into the three types of Section 1: six normative or substantive families, eight organisational or process families and five evaluative families. In the text, tables and figures a trailing hyphen marks a truncated stem (*sustainab-*) and a slash joins alternative forms (*diversity/diverse*).

For *multi-* we excluded words such as *multiple*, *multiply* and *multiplex*, and we counted as organisational only compounds that describe organisation, scope or perspective (*multidisciplinary*, *multi-stakeholder*, *multi-actor*, *multi-sectoral*, *multi-institutional* and the others listed in Appendix A). The first version of this paper also counted *multifaceted* there. It describes a thing with many aspects, says nothing about organisation, and was removed.

Control words fall into four groups: a function word (*the*), methodological nouns (*measure*, *data*, *method*, *experiment*, *sample*), a domain noun (*temperature*) and neutral abstract academic nouns that are not evaluative (*approach*, *analysis*, *process*, *structure*, *effect*). The last group was added in this revision, because umbrella words are mostly abstract nouns and adjectives and a baseline of concrete nouns alone could confuse a rise of evaluative framing with a general shift towards abstract vocabulary. As a further baseline, each family is compared with all ordinary words of similar initial frequency (Section 5.2).

LLM-typical words are 11 style words with a documented excess after 2023 in the list of Kobak et al. (2025): *delve*, *showcases* (*showcased*, *showcasing*), *underscore*, *intricate*, *intricacies*, *pivotal*, *meticulous*, *commendable*, *realm*, *notably* and *invaluable*, with their inflections. The code checks that none of them matches an umbrella family or a goal verb of Section 5.3. The full excess-vocabulary list of Kobak et al. was not used, because it contains umbrella words such as *diverse* and *innovative* and goal verbs such as *fostering*.

5.2 Growth (H1)

For each series s and year t we counted the documents n_{st} and the documents k_{st} that contain a family at least once, with Wilson 95% confidence intervals, and the occurrences per 100 tokens. The main growth estimate is a binomial generalised linear model on the yearly counts:

$$k_{st} \sim \text{Binomial}(n_{st}, \pi_{st}), \quad \text{logit}(\pi_{st}) = \alpha_s + \beta_s \cdot (t - 2000) / 10.$$

$\exp(\beta_s)$ is the odds ratio per decade. Standard errors are Newey-West heteroskedasticity- and autocorrelation-consistent (HAC) errors with $L = \lceil 4 (T/100)^{2/9} \rceil$ lags (2 or 3 for $T = 19$ to 46 years), which also absorb overdispersion. We report two checks. The first is weighted least squares of the empirical logit $\log((k_{st} + 0.5) / (n_{st} - k_{st} + 0.5))$, with weights $1/v_{st}$, where $v_{st} = 1/(k_{st} + 0.5) + 1/(n_{st} - k_{st} + 0.5)$, again with HAC errors. The second is the unweighted ordinary least squares fit of the first version.

Abstracts grew longer (the median NSF abstract had 158 tokens in 1987-91 and 369 in 2021-25), and a longer document is more likely to contain any given word. We therefore also report growth in occurrences per 100 tokens and an odds ratio per decade adjusted for length. For the latter, documents were grouped into cells of year by twenty quantile bins of log token count, and a binomial model with decade and the cell mean of log tokens was fitted with errors clustered by year.

The frequency-matched baseline uses, for each family and corpus, all single words other than stopwords and umbrella forms whose document share in the early window lies within a factor of 1.1 of the family's share. The factor was widened to 1.25, 1.5 or 2 until at least 100 words qualified; families with fewer than 20 matched words even at a factor of 2 were skipped. We report where the family's log2 growth ratio falls in the distribution of their ratios.

5.3 Diffusion across fields and loss of referents (H2)

Each article has one primary OpenAlex field (26 fields, with 245 subfields in the sample), taken from its primary topic; the 0.8% of articles without one were dropped. For each family w and period τ , $q_{\{f|w,\tau\}}$ is the field distribution of the articles that use the word and $p_{\{f|\tau\}}$ that of all articles. The bias-corrected concentration is

$$KL^*_{\{w,\tau\}} = \sum_f q_{\{f|w,\tau\}} \cdot \log(q_{\{f|w,\tau\}} / p_{\{f|\tau\}}) - E[KL \text{ of } n_{\{w,\tau\}} \text{ randomly drawn articles}],$$

where the expectation is the mean over 200 draws of as many articles as use the word, drawn from all articles of the period. Values near zero mean that the word is used without regard to field. To check that a change in the field mix of the sample does not drive the trend, we also computed KL^* with every period reweighted to the pooled 1980-2025 field distribution p^0 (each article in field f weighted by $p^0_f / p_{\{f|\tau\}}$). We further report the coefficient of variation of field-level usage rates (fields with at least 300 articles in the period) and the share of uses outside a word's home fields. Home fields were defined from the data: in the earliest period with at least 100 uses, the fields with the highest usage rates that together hold at least half of the uses. For *sustainab-*, *diversity* and *resilien-* we also report fields fixed in advance as the fields in which the words began as technical terms. KL^* was computed for all 19 families, the reference group and the control words.

Collocates were extracted with regular expressions, without part-of-speech tagging. In "sustainable X", X is the first word after *sustainable* or *sustainably*, skipping an inserted "and Y". In "X diversity", X is the

word immediately before *diversity* or *diversities*; in "diversity of X", X is the first word after *of*, skipping *the*, *their* or *its*. Stopwords were not counted as collocates. Each X was assigned to a class by hand-built dictionaries (Appendix B): environmental or resource, socio-economic or institutional, and technical objects for *sustainable*; biological, social, generic and goal-verb collocates for *diversity*. We report class shares over classified collocates, together with the share left unclassified. Rarefied type counts use 300 occurrences and 200 draws; periods with fewer than 300 occurrences are not reported, and class shares need at least 150.

In the "X diversity" slot a goal verb and a modifier compete for the same position: in "increase genetic diversity" the slot holds *genetic*. A second test therefore looks at every occurrence of *diversity*. It records whether one of the three preceding words of the same clause is a goal verb (*increase*, *promote*, *enhance*, *broaden*, *foster*, *support*, *ensure* and the others in Appendix B). Separately, it records whether the word before it, or the object of "diversity of", is biological.

As a check that does not depend on dictionaries, we trained skip-gram word2vec models (200 dimensions, window 5, minimum count 20, five epochs) on equal numbers of documents per period. For articles these were 120,000 documents in each of five periods; for NSF, as many abstracts in each of four periods as the smallest period holds (67,624). Following Hamilton et al. (2016), the models were aligned to the last period by orthogonal Procrustes on the shared vocabulary. We measured three things. The first is the cosine distance between a word's first and last vectors, compared with words of similar frequency through the residual of log distance on log frequency. The second is the overlap (Jaccard) of the 20 nearest neighbours in the first and last period, which needs no alignment. The third, computed within each period, is the cosine of *diversity* with eight biological and with eight social seed words, and of *sustainable* with eight environmental and eight socio-economic seed words (Appendix B).

5.4 Homogenisation (H3)

From each year we drew 1,500 documents at random, in five independent replicate draws (three for boilerplate). Each document became a TF-IDF vector. The vocabulary and IDF weights were fitted once per series on the pooled sample of all years, so a word that becomes common later keeps the same weight in every year. Settings were a minimum document frequency of 5, removal of words in more than half of all documents, sublinear term frequency and the English stopwords list of scikit-learn. Vectors were normalised to unit length, and the mean pairwise cosine similarity is

$$\text{mean cos}_t = (\|\sum_i v_i\|^2 - \sum_i \|v_i\|^2) / (N(N-1)), \quad N = 1,500.$$

For articles we also computed the mean over pairs from different fields. Because TF-IDF down-weights words that appear in many documents, deleting umbrella words is bound to change TF-IDF similarity little once they are common. We therefore measured the part of similarity carried by umbrella words in three ways. (i) We deleted them before vectorising and normalised the vectors again. (ii) We computed the exact share of the summed pairwise dot products that comes from umbrella-word dimensions, without normalising again. (iii) We computed that share for plain term-frequency vectors without IDF weighting. We also report the mean Jaccard overlap of the content-word sets of 2,000 random pairs, with and without umbrella words. The boilerplate rate is the share of a document's content 5-grams (at least two non-stopwords) that also occur in five or more other documents of the same year sample.

5.5 Specificity (H4)

We measured three properties of the text that remains after umbrella words are removed: (i) mean concreteness from the ratings of Brysbaert et al. (2014); (ii) mean inverse document frequency (IDF), $\log(N/df)$, against a reference table computed from 300,000 article abstracts (words in at least three of them), which we use as lexical specificity; and (iii) the share of numeric tokens. All three are means over tokens, so a repeated word counts each time. Concreteness ignores tokens that are not in the ratings, which cover on average 81% of the remaining tokens in grant abstracts and 73% in articles; tokens not in the IDF table get the maximum IDF. Because umbrella words are removed first, the association is not built into the measures. Each measure was standardised and regressed on standardised umbrella density and log token count, with year-by-scheme fixed effects (year-by-field for articles) and HC1 standard errors. Documents with fewer than 60 tokens were excluded. For the grant corpora we re-estimated the model on the awards that carry an OpenAlex topic, first with the same fixed effects and then with year by scheme by field and by subfield. For NSF we also dropped education-oriented awards: those in the subfield Education or with three or more words such as *education*, *teacher*, *school* or *curriculum*.

5.6 Decoupling (H5) and outcomes (H6)

For H5 we used the linked FP7 and Horizon 2020 projects. The number of participants counts the coordinator and the listed participants, and the number of countries counts their distinct countries. Collaboration vocabulary is *collaborat-*, *cooperat-* or *partnership*; interdisciplinarity vocabulary is organisational *multi-* or *inter-*, *trans-* or *cross-disciplinary*. The first version fitted logit models with log number of participants, an indicator for two or more countries, log token count and year fixed effects. We now add linear probability models with year-by-funding-scheme fixed effects, on all linked projects and within collaborative schemes only. To follow the vocabulary of single-beneficiary schemes into Horizon Europe, beyond the end of the participant data, we also compare collaborative and individual schemes across all EU objectives of 2008-2026. For articles we compared usage rates by authorship structure (single author, one institution, several institutions in one country, several countries), raw and standardised to the field mix of each period, and estimated a linear probability model with year-by-field fixed effects for 2010-2025. For articles we also tracked *international*, *transnational* and *cross-border*, the claim of cross-national work that co-authorship across countries can check directly.

For H6 the outcome for articles published in 2000-2021 is $\log(1 + \text{FWCI})$, with field-weighted citation impact capped at its 99th percentile. For EU projects starting in 2008-2019 and NSF awards from 1995-2019 the outcome is $\log(1 + \text{linked publications})$, the number of works OpenAlex links to the award up to the snapshot. Budgets are nominal amounts (EUR for the EU, USD for NSF). The year-by-scheme fixed effects absorb differences in price level and in time since the start between cohorts. Controls were log budget and log token count (for articles, log token count, log number of authors and international co-authorship). Fields for awards come from the OpenAlex primary topic of the award. Robustness checks re-estimated the baseline model on the awards that carry a topic, then added field or subfield fixed effects, replaced the linear term by density bins, and for NSF dropped education-oriented awards.

5.7 Dependence on funders, politics and LLMs (H7)

At the EU programme boundaries of 2014 and 2021 we computed the change in the share of documents between the two years before and the two years after, minus the change in NSF abstracts over the same

years. The 95% confidence intervals come from 500 bootstrap replicates that resample documents within each corpus-by-year cell. We estimated the jumps for all EU objectives and separately within collaborative and within individual schemes, so that a change in the scheme mix at a boundary cannot produce them.

For NSF in 2025 we computed relative changes against 2023-24 with 95% confidence intervals from the log risk ratio, and the same for EU and NIH abstracts. We compared the year-on-year total variation distance (TVD) of the topic mix of NSF awards that carry a topic with the movement in ordinary years (2016-2024), at the level of the 26 fields and of their subfields. Because TVD between finite samples is positive even when nothing changes, we also report the TVD expected from sampling alone (200 multinomial draws of the observed sizes from the pooled two-year mix). We recomputed the 2025 drops without education-oriented awards and with 2025 reweighted to the 2023-24 mix of subfields.

For LLM-typical words we measured the excess of the yearly share over a linear 2013-2022 trend, the approach of Kobak et al. (2025); "excess" always means the observed share minus the trend prediction. We compared umbrella density in documents with and without these words, as a ratio of means and as a difference adjusted for length within year-by-scheme cells (year-by-field for articles).

6. Results

6.1 H1: normative and evaluative umbrella words grew far faster than control words; organisational words grew about as fast as the fastest neutral nouns

The share of documents using umbrella words rose much faster than the share using most control words in NSF abstracts, articles and EU collaborative objectives (Table 3, Figures 1 and 2). Four in five NSF abstracts and more than nine in ten EU collaborative objectives of 2021-25 contain at least one umbrella word.

Table 3. Share of documents using each family, early and late window (multiple in parentheses)

Family	NSF, 1987-91 to 2021-25	Articles, 1980-84 to 2021-25	EU collaborative, 2008-10 to 2023-25	EU individual, 2008-10 to 2023-25
diversity/diverse	3.5% to 26.1% (×7.5)	0.75% to 4.7% (×6.3)	6.6% to 14.3% (×2.2)	8.1% to 10.9% (×1.4)
sustainab-	0.09% to 10.2% (×118)	0.02% to 4.3% (×171)	14.7% to 39.2% (×2.7)	2.6% to 12.1% (×4.7)
multi- (organisational)	1.1% to 5.3% (×4.9)	0.18% to 1.0% (×5.6)	9.3% to 15.0% (×1.6)	7.7% to 6.2% (×0.8)
collaborat-	8.5% to 27.3% (×3.2)	0.25% to 2.3% (×9.3)	16.5% to 25.0% (×1.5)	11.1% to 8.0% (×0.7)
innovat-	2.4% to 24.1% (×10.3)	0.46% to 3.8% (×8.3)	26.8% to 52.1% (×1.9)	9.2% to 24.4% (×2.7)

Family	NSF, 1987-91 to 2021-25	Articles, 1980-84 to 2021-25	EU collaborative, 2008-10 to 2023-25	EU individual, 2008-10 to 2023-25
resilien-	0.06% to 6.2% (×106)	0.02% to 1.3% (×62)	1.1% to 12.8% (×12.0)	0.6% to 3.5% (×6.2)
Umbrella index (any of 19)	25.8% to 82.1% (×3.2)	5.6% to 29.0% (×5.2)	74.6% to 92.7% (×1.2)	48.7% to 67.4% (×1.4)
Normative/substantive (any of 6)	4.4% to 42.4% (×9.7)	1.3% to 11.7% (×8.9)	24.7% to 60.3% (×2.4)	12.1% to 28.7% (×2.4)
Organisational/process (any of 8)	14.0% to 44.5% (×3.2)	1.2% to 6.0% (×5.2)	44.4% to 62.1% (×1.4)	26.1% to 26.1% (×1.0)
Evaluative (any of 5)	10.6% to 58.3% (×5.5)	3.4% to 18.1% (×5.3)	49.0% to 73.2% (×1.5)	26.2% to 46.9% (×1.8)
Promotional adjectives (reference)	10.8% to 43.4% (×4.0)	2.3% to 10.3% (×4.4)	32.3% to 44.2% (×1.4)	40.5% to 60.1% (×1.5)
Control: experiment-	21.3% to 22.6% (×1.06)	9.1% to 10.2% (×1.13)	11.2% to 7.3% (×0.65)	23.3% to 19.1% (×0.82)
Control: method	17.2% to 33.7% (×2.0)	11.3% to 33.2% (×2.9)	18.1% to 17.1% (×0.95)	23.7% to 27.1% (×1.15)
Control: data	21.3% to 43.1% (×2.0)	12.5% to 25.7% (×2.1)	21.4% to 33.8% (×1.6)	20.8% to 25.6% (×1.2)
Control: approach	10.9% to 34.5% (×3.2)	4.0% to 16.3% (×4.1)	29.2% to 37.8% (×1.3)	33.6% to 40.0% (×1.2)
Control: analysis	21.5% to 25.6% (×1.2)	10.2% to 27.3% (×2.7)	20.0% to 17.9% (×0.9)	25.3% to 23.0% (×0.9)
Control: the	99.5% to 100% (×1.00)	95.7% to 96.0% (×1.00)	100% to 100% (×1.00)	100% to 100% (×1.00)

Notes: Appendix Table C4 gives every family and control word for these four corpora; the tables for NIH and all EU objectives are in the repository. The multiples are ratios of the unrounded shares.

The scheme mix of the EU changed strongly. Collaborative schemes made up 52.8% of EU projects in 2008-10 and 34.6% in 2023-25, and individual schemes rose from 46.1% to 62.2%. Pooled over all EU objectives, *collaborat-* therefore stayed at 13.9% and organisational *multi-* rose only from 8.5% to 9.3%. Within collaborative schemes both rose (Table 3), while within individual schemes both fell. The flat pooled trend reported in the first version was a composition effect.

Table 4. Growth per decade: odds ratios from the binomial model with HAC errors, the same adjusted for document length, and the position among frequency-matched ordinary words

Family	NSF: OR [95% CI]	NSF: OR, length-adjusted	NSF: percentile among matched words	Articles: OR [95% CI]	Articles: OR, length-adjusted	Articles: percentile among matched words
diversity/diverse	1.87 [1.67, 2.08]	1.51	98	1.56 [1.45, 1.68]	1.44	98
sustainab-	2.68 [2.29, 3.15]	2.28	100	2.68 [2.24, 3.21]	2.53	100
multi-(organisational)	1.39 [1.18, 1.64]	1.12	95	1.48 [1.40, 1.55]	1.39	98
collaborat-	1.43 [1.27, 1.62]	1.19	92	1.64 [1.50, 1.79]	1.52	99
innovat-	1.89 [1.74, 2.05]	1.74	100	1.73 [1.56, 1.91]	1.67	100
resilien-	4.22 [3.86, 4.61]	3.69	100	3.29 [2.49, 4.34]	2.97	100
Umbrella index	2.31 [1.99, 2.69]	1.78		1.62 [1.50, 1.75]	1.52	
Promotional adjectives	1.74 [1.62, 1.86]	1.38		1.45 [1.40, 1.49]	1.35	
Control: experiment-	1.03 [1.00, 1.05]	0.82	15	1.03 [1.02, 1.04]	0.98	46
Control: method	1.32 [1.29, 1.35]	1.09	73	1.42 [1.37, 1.47]	1.34	98
Control: data	1.38 [1.34, 1.42]	1.13	83	1.25 [1.17, 1.34]	1.17	88
Control: approach	1.54 [1.47, 1.62]	1.21	92	1.44 [1.34, 1.55]	1.35	99

Notes: The odds ratio per decade comes from a binomial model of the yearly counts with Newey-West errors (Section 5.2). The length-adjusted odds ratio adds log token count in a model of year by length-bin cells with errors clustered by year; its confidence intervals are in Appendix Table C4. The percentile


is the position of the family's early-to-late growth among all ordinary words with a similar early document share.

The binomial model with HAC errors gives wider intervals than the unweighted fit of the first version but the same ordering. For *diversity/diverse* in NSF abstracts the odds ratio per decade is 1.87 [1.67, 2.08], against 1.99 [1.86, 2.13] from the first version's fit. The weighted least-squares fit agrees closely with the binomial model (Appendix Table C4).

Document length matters for the comparison. The median NSF abstract had 158 tokens in 1987-91 and 369 in 2021-25, and the median article abstract 122 and 197; EU objectives stayed at about 275. Holding length fixed lowers every odds ratio for NSF abstracts, umbrella and control alike. The normative and evaluative families still grow far faster than any control word: *sustainab-* 2.28, *resilien-* 3.69, *innovat-* 1.74 and *diversity/diverse* 1.51 per decade, against 0.81 to 1.21 for the control words (leaving aside *the*, which is at its ceiling). The two organisational core families do not. Adjusted for length, *collaborat-* (1.19) and organisational *multi-* (1.12) grew in NSF abstracts about as fast as *approach* (1.21) and *data* (1.13). In articles they grew a little faster than the fastest controls (1.52 and 1.39 against 1.35 for *approach* and 1.34 for *method*). Growth in occurrences per 100 tokens tells the same story; for NSF the density of *collaborat-* rose by a factor of 1.7 and that of *approach* by 1.8 (Appendix Table C4). Against ordinary words of similar initial frequency, every normative and evaluative family except *excellence/excellent* ranks at or above the 97th percentile of growth in both NSF abstracts and articles. Among the organisational families, *collaborat-* and organisational *multi-* rank at the 92nd and 95th percentiles in NSF abstracts, where *approach* also ranks at the 92nd, and at the 99th and 98th in articles. *Cooperation* is the exception at the bottom: it fell in NSF abstracts and ranks at the 1st percentile there.

The control word *method* doubled in NSF abstracts and nearly tripled in articles. Much of this reflects longer abstracts and the spread of structured abstracts, which ask for a methods sentence: adjusted for length, its odds ratio per decade is 1.09 in NSF abstracts and 1.34 in articles. *Approach* and *data* grew for similar reasons, and *data* also tracks the real growth of data-intensive research.

Share of documents using each core family (95% Wilson CI)






In 2008-10, umbrella density was 0.86 occurrences per 100 tokens in collaborative EU projects, 0.37 in individual ones, 0.69 in NSF abstracts and 0.14 in articles. In collaborative EU projects it reached 1.90 in 2023-25, a factor of 2.2 (Figure 3). By type, normative and substantive words grew fastest in NSF abstracts, articles and EU objectives (Figure 3b). Averaged by programme, *sustainab-* rose from 9.5% of all EU objectives in FP7 to 13.4% in Horizon 2020 and 22.7% in Horizon Europe, and *resilien-* from 1.1% to 2.6% and 8.1%. NIH abstracts, which are short in OpenAlex (median 83 tokens), show much less growth: the umbrella index rose from 20.9% in 1985-94 to 27.4% in 2021-25.

H1 is supported for the normative and evaluative families in every corpus. For the organisational families it holds in articles and EU collaborative schemes, but in NSF abstracts their growth is no faster than that of the fastest-growing neutral nouns once document length is held constant.

Umbrella-word density by type (families listed in Appendix A)



6.2 H2: umbrella words spread across fields, and in NSF abstracts *diversity* lost its objects

Diffusion across fields (H2a). In the article sample, KL^* fell between the first and the last period for 15 of the 19 families (Table 5; all families are in Appendix Table C1). For *sustainab-* it fell from 1.00 in the 1990s to 0.42 in 2020-25, for *innovat-* from 0.62 to 0.20 and for *diversity/diverse* from 0.28 to 0.12. The fall does not come from a change in the field mix of the sample. With every period reweighted to the 1980-2025 field distribution, the same three values are 0.89 to 0.43, 0.57 to 0.20 and 0.25 to 0.13. The coefficient of variation of field usage rates fell in step (for *sustainab-*, from 1.70 to 1.01). Four families became more concentrated or stayed level: *inclusive/inclusion* (0.09 to 0.20, mostly in the social sciences and medicine), *excellence/excellent* (0.08 to 0.28), *cooperation* (about 0.3 throughout) and *co-creation/co-design* (0.43 to 0.45 since the 2010s).

Control words start from much lower concentration (KL^* between 0.04 and 0.17 in the 1980s, except *temperature* at 0.51), and seven of the eleven changed by less than 0.05. The largest fall among them was for *approach*, from 0.14 to 0.05, while *experiment* and *temperature* became more concentrated. Umbrella families started between 0.3 and 1.1 and most fell by 0.1 to 0.7. The share of uses outside a priori home fields rose for *sustainab-*, from 49% in the 1990s to 70% in 2020-25, and stayed level for *diversity* (about 70%) and *resilien-* (about 78%). The share outside data-defined home fields rose for control words too, so it mostly reflects the changing field mix of the sample; KL^* is the better guide. H2a is supported for most families. On its own it is also what a boundary object would show.

Table 5. Field concentration of article abstracts (bias-corrected KL^* , nats)

Family	1980-89	1990-99	2000-09	2010-19	2020-25	2020-25, fixed field mix
sustainab-	not estimated	1.004	0.711	0.480	0.417	0.425
stakeholder	not estimated	1.081	0.626	0.516	0.401	0.431
innovat-	0.624	0.529	0.424	0.291	0.204	0.200

Family	1980-89	1990-99	2000-09	2010-19	2020-25	2020-25, fixed field mix
multi- (organisational)	0.647	0.497	0.393	0.275	0.264	0.307
inter-/trans-/cross- disciplinary	0.510	0.545	0.337	0.329	0.249	0.264
resilien-	not estimated	0.344	0.379	0.355	0.291	0.315
collaborat-	0.332	0.361	0.262	0.239	0.221	0.241
diversity/diverse	0.276	0.265	0.226	0.189	0.123	0.133
inclusive/inclusion	0.088	0.085	0.089	0.139	0.200	0.201
Controls (median of 11)	0.085	0.087	0.123	0.101	0.067	0.062

Notes: The rows are the four core families and the families with the most uses in each type; Appendix Table C1 lists all 19 families, the reference group and every control word. "Not estimated" marks periods with fewer than 100 articles using the family. The last column holds the field mix of every period at its pooled 1980-2025 distribution. The control row excludes *the*.

Loss of referents (H2b). In NSF abstracts the objects of *sustainable* moved away from environmental and resource objects (Figure 4d). Among classified collocates their share was 83% in 1987-94 (from only 124 phrases), 68% in 1995-2004 and 55% in 2015-25. Socio-economic, institutional and durability objects (*growth, future, solutions, careers* and similar) rose from 15% to 31%, and technical objects (*computing, software, design* and similar) from 2% to 14%. The unclassified share rose from 15% to 26%, so the shift holds over classified phrases and does not come from a growing unclassified tail. *Sustainable development* fell from 40% to 3% of all "sustainable X" phrases. Articles moved the same way, less strongly: environmental or resource objects fell from 82% of classified collocates in the 1990s to 70% in 2020-25 (unclassified 21% to 23%). EU objectives stayed at about 59% throughout 2007-2026. Rarefied collocate counts at 300 occurrences grew from 141 to 171 types in NSF abstracts (1995-2004 to 2015-25) and from 108 to 149 in articles (1990s to 2020-25).


The clearest change is in *diversity* in NSF abstracts. Among classified collocates of "X diversity" and "diversity of X", biological ones (*genetic, species, microbial* and similar) fell from 76% in 1987-94 to 42% in 2015-25. Goal verbs rose from 2% to 27% and social collocates from 8% to 22%, while the unclassified share stayed near 37%. The window test confirms that this is not an artefact of goal verbs and modifiers competing for one slot. The share of occurrences of *diversity* with a goal verb among the three preceding words rose from 4.6% to 23.4%, and the share with a biological reference fell from 42.9% to 20.8%. Occurrences with both stayed at 1% to 2%, so goal phrases are mostly phrases such as "broaden the diversity" or "increase diversity", not biological phrases with a verb in front (Figure 4f). In articles the movement went the other way: biological collocates rose from 60% to 70% of classified collocates and goal verbs stayed below 6% of occurrences, in line with the growth of genomics and microbiome

research. EU objectives show a small move from biological collocates (69% to 62%) and a rise of goal verbs from 2.5% to 6.0% of occurrences.


The embedding check agrees for *diversity* and qualifies the result for *sustainable* (Figure 5). In NSF abstracts, *diversity* in 1987-94 was closer to the biological seed words (cosine 0.65) than to the social ones (0.41); by 2015-25 the order had reversed (0.50 against 0.58). Its nearest neighbours changed from *biodiversity*, *richness* and *invasions* to *inclusivity*, *diversification* and *inclusion*. In articles, *diversity* stayed closer to the biological seeds (0.63 against 0.56 in 2020-25), and its meaning changed less than that of 82% of words of similar frequency. *Sustainable* changed more than 84% of words of similar frequency in NSF abstracts and 77% in articles. Its 20 nearest neighbours in the first and last period had almost none in common (Jaccard 0.03, against a median of 0.11 to 0.14 for random shared words). Its neighbours moved from resource sectors (*fishery*, *forestry*, *agriculture*) to policy labels (*climate-smart*, *carbon-neutral*, *low-carbon*, *eco-friendly*). It did not move towards the socio-economic seed words, and in articles it moved slightly towards the environmental ones. *Inclusion*, *inclusive* and *resilience* (*resilient* in articles) were among the 4% of words that changed most in both corpora. Control words such as *method*, *approach* and *effect* changed less than most words of their frequency.

H2b is supported for *diversity* in NSF abstracts, where dictionaries, the window test and the embeddings agree. For *sustainable* it is supported in the narrower sense that the word lost its specific resource objects and came to modify general and institutional nouns, while its associations stayed environmental in a broad sense. It is not supported for *diversity* in articles, where the word became more biological as it spread across fields.

Cross-field diffusion (top) and referential dilution (bottom) of umbrella words



Distributional check: diachronic word embeddings (skip-gram, aligned by orthogonal Procrustes)



6.3 H3: NSF abstracts and articles became more similar, and umbrella words carry little of the similarity

Mean pairwise TF-IDF cosine similarity roughly doubled, from 0.018 in 1987 to 0.039 in 2025 for NSF abstracts and from 0.009 in 1980 to 0.017 in 2025 for articles (Figure 6a). For articles restricted to pairs from different fields it rose from 0.008 to 0.016, so the convergence is not a matter of more articles from the same field. Plain term-frequency vectors and the Jaccard overlap of content-word sets show the same doubling (Jaccard 0.025 to 0.049 for NSF, 0.009 to 0.019 for articles). EU objectives stayed between 0.025 and 0.027 on TF-IDF and between 0.030 and 0.032 on Jaccard. The five replicate draws per year differ by less than 0.001, which is why the bands in Figure 6 are barely visible.


Umbrella words carry little of this similarity under every weighting (Figure 6b). The exact share of the summed pairwise dot products that comes from umbrella-word dimensions rose for NSF from 0.3% in 1987 to 1.8% in 2025 (2.6% in 2024). For EU objectives it rose from 1.3% in 2008 to 3.7% in 2025, and for articles from 0.04% to 1.2%. Without IDF weighting the shares are smaller (NSF 0.8%, EU 3.2%, articles 0.7% in 2025), because frequent general words such as *research*, *project* and *students* then dominate the dot products. Deleting umbrella words before vectorising changes similarity by at most 2.2%, and Jaccard overlap barely moves. The small contribution is therefore not an artefact of IDF weighting. Most of the convergence comes from shared application and reporting vocabulary around the umbrella words.

The NSF boilerplate rate rose from 1.3% in 1987 to 2.3% in 2025 (replicate range 2.3% to 2.4%). The figure of 2.2% given for 2025 in the first version came from a single draw, and the drop shown there in 2025 did not recur in the replicate draws. The spikes have identifiable causes. In 2009 almost a third of the sampled awards carried the sentence on funding under the American Recovery and Reinvestment Act, and in 2022 about a tenth carried a sentence on the American Rescue Plan Act. Templated abstracts of Small Business Innovation Research awards ("This Small Business Innovation Research Phase I project...", "The broader impact/commercial potential of this...") produce the smaller peaks of 1998, 2003-04, 2012 and 2019-25. Apart from these, the most frequent shared 5-gram was "the goal of this project is", found in 39 of the 39 yearly samples. EU objectives have very little boilerplate (0.2% in 2008 to 0.1% in 2025). Until about 2010 their shared phrases were opening formulas such as "the main objective of the project is"; in 2024 they were "will pave the way for", "the state of the art" and "there is an urgent need". The article series is not interpretable. Its peaks come from clusters of records whose abstract field holds template text, such as Crossref and Altmetric notices ("attention score", "calculated by Crossref") or the item descriptions of a digitised archive, which the row-group sample draws together

because they were ingested together. Such records are 2.7% of the 1980s sample, 1.5% of the 1990s and under 1% from 2000. We show the article series dotted in Figure 6d for completeness.

H3 is partly supported: NSF abstracts and articles became more similar to one another, and NSF abstracts contain more boilerplate, but EU objectives did not converge, and umbrella words account for at most about 4% of pairwise similarity.

Homogenisation of project and article texts (N=1,500 documents per year; bands: range of replicate draws)



6.4 H4: documents dense in umbrella words use slightly less specific vocabulary

The coefficient for lexical specificity (mean IDF) was negative in every corpus (Table 6). Concreteness and the share of numbers were also negative, except in NSF abstracts, where concreteness had a small positive coefficient and the share of numbers none. Specificity fell most in the top quintile of umbrella density (Figure 7).

Table 6. Change in specificity of the remaining text per SD of umbrella density (SD units, [95% CI])

Corpus and model	n	Lexical specificity (mean IDF)	Concreteness (Brysbaert)	Share of numeric tokens
EU, year x scheme FE	83,286	-0.103 [-0.111, -0.096]	-0.100 [-0.108, -0.093]	-0.055 [-0.063, -0.048]
EU with topic, year x scheme x subfield FE	34,307	-0.092 [-0.111, -0.074]	-0.098 [-0.115, -0.081]	-0.074 [-0.092, -0.056]
NSF, year x scheme FE	363,341	-0.079 [-0.082, -0.076]	+0.023 [0.020, 0.026]	-0.002 [-0.005, 0.002]
NSF with topic, year x scheme FE	123,093	-0.066 [-0.072, -0.059]	+0.053 [0.047, 0.060]	-0.010 [-0.016, -0.004]
NSF with topic, year x scheme x field FE	123,093	-0.037 [-0.043, -0.030]	+0.024 [0.018, 0.031]	-0.012 [-0.018, -0.005]
NSF with topic, year x scheme x subfield FE	123,093	-0.029 [-0.036, -0.023]	+0.016 [0.009, 0.022]	-0.016 [-0.023, -0.010]
NSF without education-oriented awards, year x scheme FE	302,946	-0.077 [-0.080, -0.073]	+0.024 [0.020, 0.027]	-0.012 [-0.015, -0.008]


Corpus and model	n	Lexical specificity (mean IDF)	Concreteness (Brysbaert)	Share of numeric tokens
NIH, year x scheme FE	155,036	-0.040 [-0.045, -0.035]	-0.075 [-0.080, -0.070]	-0.051 [-0.055, -0.046]
NIH with topic, year x scheme x subfield FE	74,502	-0.050 [-0.059, -0.040]	-0.075 [-0.085, -0.065]	-0.054 [-0.064, -0.045]
Articles, year x field FE	1,125,725	-0.090 [-0.092, -0.088]	-0.064 [-0.065, -0.062]	-0.060 [-0.061, -0.058]
Articles, year x subfield FE	1,120,459	-0.081 [-0.083, -0.080]	-0.062 [-0.064, -0.060]	-0.052 [-0.054, -0.051]

Notes: Log token count is controlled; HC1 standard errors. All three measures use the same documents (at least 60 tokens, with both specificity measures defined). "With topic" are the awards to which OpenAlex assigned a primary topic. Education-oriented awards are those in the subfield Education or with three or more education words (60,395 NSF awards).

The NSF pattern is not an artefact of programme composition. The positive concreteness coefficient survived field and subfield fixed effects, shrinking to +0.016 within subfields, and the removal of 60,395 education-oriented awards left it at +0.024. Education and outreach awards therefore do not explain it on their own. A possible source is the broader-impacts paragraph that NSF requires in every field, which names people and places (*students, schools, communities*); we did not test this directly. On the awards with a topic, the negative coefficient for lexical specificity in NSF abstracts roughly halved once fields were held fixed (-0.066 to -0.037), so part of it reflects differences between fields; within subfields it stayed negative. In the other corpora the coefficients barely changed with finer fixed effects.

The effects are small. In raw units, one SD of umbrella density goes with a mean IDF lower by 0.03 nats in NSF abstracts, 0.04 in EU objectives and 0.06 in articles, against a standard deviation across documents of 0.37 to 0.67 nats. Concreteness falls by at most 0.016 points on the five-point scale. The difference is concentrated in the most umbrella-dense documents. Between articles without umbrella words and those in the top density quintile, mean IDF differs by 0.37 nats and the share of numeric tokens falls from 3.7% to 1.4%; in the grant corpora the IDF difference is 0.05 to 0.08 nats. H4 is supported for lexical specificity in every corpus and for concreteness and numbers outside NSF, with effects of at most a tenth of a standard deviation.

Specificity of the remaining text by umbrella-word density (within year x scheme, or year x field)



6.5 H5: organisational vocabulary is common without collaborative structure, but within collaborative schemes it follows structure loosely

EU projects. Table 7 shows the linked FP7 and Horizon 2020 projects ($n = 36,177$; 55.3% have a single participant). Pooled over both programmes, 13.6% of single-participant projects used collaboration vocabulary, and these projects supplied 38.5% of all documents that used it. Interdisciplinarity vocabulary was about as common in single-participant projects (16.3%) as in consortia of ten or more partners (17.9%). Most consortia did not use collaboration vocabulary at all: among projects with two or more participants, 27.0% did.

Table 7. Use of vocabulary by number of participating organisations, EU FP7 and Horizon 2020

Programme	Participants	n	Collaboration vocabulary	Interdisciplinarity vocabulary
FP7	1	13,281	14.7%	16.8%
FP7	2-4	2,335	29.1%	16.7%
FP7	5-9	4,590	23.5%	13.2%
FP7	10 or more	4,939	28.2%	16.6%
H2020	1	6,742	11.6%	15.2%
H2020	2-4	1,004	25.1%	14.8%
H2020	5-9	1,518	27.6%	21.9%
H2020	10 or more	1,768	30.4%	21.6%
Both, pooled	1	20,023	13.6%	16.3%
Both, pooled	2-4	3,339	27.9%	16.1%
Both, pooled	5-9	6,108	24.5%	15.3%
Both, pooled	10 or more	6,707	28.8%	17.9%

Notes: Collaboration vocabulary is *collaborat-*, *cooperat-* or *partnership*; interdisciplinarity vocabulary is organisational *multi-* or *inter-*, *trans-* or *cross-disciplinary*. Horizon 2020 projects are those in the 2017 snapshot of the participant files. The pooled rows are computed from the unrounded counts.

The models separate the effect of consortium size from that of the funding rules (Table 8). The first version fitted a logit with year fixed effects only. There, collaboration vocabulary was associated with participation of two or more countries (0.858 [0.727, 0.989]) and unrelated to the log number of participants (0.016 [-0.041, 0.072]). That model pools single-beneficiary fellowships with consortia, so the country indicator absorbs the difference between the two kinds of scheme. With year-by-scheme fixed effects, collaboration vocabulary rises with the number of participants: by 2.6 percentage points per log unit [1.2, 3.9] across all linked projects and by 5.3 [3.7, 6.8] among collaborative-scheme projects with two or more participants. In the latter group, projects spanning two or more countries used it 21.8 points [16.3, 27.2] more often than single-country projects. Interdisciplinarity vocabulary also rises with

the number of participants within schemes (3.4 points per log unit [2.2, 4.6] in the same group), but not with international participation. Going from 3 to 30 partners, a factor of ten, raises the predicted probability of collaboration vocabulary by about 12 points from a base of roughly 25%.

Table 8. Linear probability models of vocabulary use, linked FP7 and Horizon 2020 projects (percentage points, [95% CI])

Sample and fixed effects	n	Collaboration: per log participant	Collaboration: two or more countries	Interdisciplinary: per log participant	Interdisciplinary: two or more countries
All linked projects, year x scheme	36,177	+2.6 [1.2, 3.9]	+2.7 [0.1, 5.2]	+2.4 [1.4, 3.4]	-1.7 [-3.9, 0.5]
Collaborative schemes, year x scheme	15,698	+2.3 [0.9, 3.8]	+5.9 [1.4, 10.4]	+2.5 [1.4, 3.5]	-4.2 [-7.2, -1.2]
Collaborative schemes, 2+ participants, year x scheme	14,627	+5.3 [3.7, 6.8]	+21.8 [16.3, 27.2]	+3.4 [2.2, 4.6]	+0.5 [-3.1, 4.2]

Notes: Log token count is controlled; HC1 standard errors. The logit with year fixed effects only, reported in the first version, is in Appendix Table C2.


Beyond 2017 the participant data end, but funding rules fix the structure of individual schemes: 93.7% of linked individual-scheme projects have one beneficiary. Across all EU objectives, the share of individual-scheme projects using collaboration vocabulary fell from 12.6% in FP7 to 11.1% in Horizon 2020 and 9.7% in Horizon Europe; in collaborative schemes it was 29.2%, 28.3% and 32.4%. Interdisciplinarity vocabulary in individual schemes stayed at 16% to 18% throughout and rose in collaborative schemes from 15.2% to 26.9%. The rise of organisational vocabulary in Horizon Europe therefore took place mostly where the structure exists.

Articles. In raw rates, single-author papers in 2020-25 used collaboration vocabulary (4.2%) and *international, transnational* or *cross-border* (6.3%) more often than internationally co-authored papers (3.1% and 3.6%). Much of this gap comes from field: single-author papers are concentrated in the social sciences and humanities, where cooperation and international relations are topics of study. Standardised to the field mix of each period, the rates in 2020-25 were 3.6% and 5.3% for single-author papers against 4.1% and 4.9% for internationally co-authored ones (Figure 8b, c). In a linear probability model for 2010-2025 with year-by-field fixed effects, single-author papers used collaboration vocabulary as often as internationally co-authored papers (+0.13 points [-0.05, 0.31]) and international vocabulary 1.0 point [0.8, 1.2] more often (0.6 [0.4, 0.8] within subfields). Collaboration vocabulary rose in every authorship group at about the same pace, from about 1% of articles in the 1980s to 4% to 7% in 2025, single-author papers included. Its growth is therefore not explained by the growth of co-authorship, even though internationally co-authored papers also became more common (from 2% to 15% of articles with author data). Article vocabulary reflects the spread of a discourse more than growth in actual

collaboration, and claims of cross-national work are not more common where co-authorship is cross-national.

H5 is partly supported. Organisational vocabulary is common in documents without the corresponding structure: one in eight single-beneficiary EU projects uses collaboration vocabulary and one in six uses interdisciplinarity vocabulary, and in articles the vocabulary barely depends on authorship once field is held constant. Within collaborative EU schemes, however, it rises with consortium size and international reach, so it is loosely coupled to structure. The first version's statement that collaboration vocabulary was unrelated to the number of partners came from pooling scheme types and does not hold within schemes.

Collaboration vocabulary and actual collaborative structure



6.6 H6: associations between umbrella density and outcomes are small and differ in sign across corpora

Table 9. Change in outcome per SD of umbrella density ([95% CI])

Model	Coefficient	n
Articles, log(1 + FWCI), year by field FE	+0.024 [0.022, 0.026]	573,970
Articles, log(1 + FWCI), year by subfield FE	+0.022 [0.020, 0.024]	573,970
EU, log(1 + linked publications), year by scheme FE	+0.013 [0.004, 0.023]	50,750
EU, same model, awards with a topic	+0.026 [0.009, 0.044]	20,199
EU, awards with a topic, plus field FE	+0.040 [0.021, 0.058]	20,199
EU, awards with a topic, plus subfield FE	+0.045 [0.023, 0.068]	20,199
NSF, log(1 + linked publications), year by scheme FE	-0.067 [-0.070, -0.063]	242,778
NSF, same model, awards with a topic	-0.071 [-0.077, -0.064]	83,310
NSF, awards with a topic, plus field FE	-0.044 [-0.050, -0.038]	83,310
NSF, awards with a topic, plus subfield FE	-0.028 [-0.034, -0.022]	83,310
NSF, without education-oriented awards	-0.065 [-0.069, -0.061]	199,634


Notes: EU and NSF models control for log budget and log token count; article models control for log token count, log number of authors and international co-authorship. HC1 standard errors. Articles 2000-2021, EU projects starting 2008-2019, NSF awards 1995-2019.

The associations are small next to the structural variables in the same models. One SD of log budget goes with 0.71 more log points of linked publications in the EU and 0.33 more in NSF, against +0.013 and -0.067 for one SD of umbrella density. In articles, one SD of log team size goes with 0.20 more log points of FWCI and international co-authorship with 0.43, against +0.024 for umbrella density. The sign is positive for articles and EU projects and negative for NSF awards. The density bins show the shapes behind the linear terms (Figure 9). For NSF the outcome declines steadily across quintiles, to -0.20 log points in the top quintile against documents without umbrella words. For the EU it rises slightly, to +0.06. For articles it steps from -0.03 in the lowest quintile to a plateau of +0.07 to +0.10 in the higher ones.

Restricting the samples to awards with a topic does not select on positive outcomes. Awards with a topic had more zero-publication cases in the EU (48% against 41%) and the same share in NSF (49% against 50%), and fewer linked publications on average (EU 3.6 against 18.8; NSF 2.5 against 4.0). The change between the baseline and the field-controlled EU estimate in the first version (+0.013 to +0.040) is about half sample and half field: on the restricted sample the baseline is +0.026, and field fixed effects raise it to +0.040. For NSF the restricted baseline is -0.071, so the attenuation to -0.044 and -0.028 comes from the field controls. Dropping the 43,144 education-oriented NSF awards leaves the coefficient at -0.065, so education grants, which may produce fewer indexed journal articles, do not explain the negative sign. Within subfields the NSF association is a third of its pooled size, which suggests that much of it reflects differences between programmes in how many publications are linked to awards.

H6 is supported in the sense stated in Section 3. No association is large next to budget and team size, and the sign is not the same across corpora. Two cautions apply. The positive associations for articles and EU projects are statistically clear, so umbrella density is not unrelated to outcomes; it is weakly and inconsistently related. And all three corpora contain funded or published work only, so selection on funding or publication can bias these associations in either direction (Section 2.3).

Outcomes by umbrella-word density (residualised on fixed effects and controls; 95% CI)




6.7 H7: the vocabulary follows funders, politics and LLM-typical writing

Programme boundaries. When Horizon 2020 began in 2014, the share of EU objectives using *innovat-* rose from 24.9% in 2012-13 to 32.7% in 2014-15. Relative to the change in NSF abstracts, the difference-in-differences is +7.1 percentage points [5.5, 8.6] (Figure 10a). The jump appears within collaborative schemes (+10.9 [8.5, 13.3]) and within individual schemes (+11.1 [9.5, 12.7]), so it does not come from a

change in the scheme mix. Inter-/trans-/cross-disciplinary rose by +1.9 [0.8, 3.0] at the same boundary. The control words *experiment* and *method* fell by −3.6 [−5.0, −2.1] and −4.3 [−5.9, −2.9] relative to NSF, partly because NSF abstracts changed format in 2014. In that year NSF dropped the "Broader impacts" and "Intellectual merit" headings from its abstracts; the share of NSF abstracts that mention broader impacts fell from 30% to 16%. The apparent jump of *impact* in EU objectives at this boundary (+12.6 points relative to NSF, against a raw EU change of +0.7) is an artefact of that change, and we do not interpret it.

When Horizon Europe began in 2021, several families rose discontinuously. In all EU objectives, *co-creation/co-design* rose from 2.1% to 4.5%, a raw change of +2.3 points and a difference-in-differences of +1.6 [1.0, 2.2] relative to NSF. Inter-/trans-/cross-disciplinary rose by +2.9 [1.8, 4.2], organisational *multi-* by +1.8 [0.8, 2.7], *societal* by +1.5 [0.6, 2.4], *stakeholder* by +1.3 [0.4, 2.4], *sustainab-* by +1.3 [0.2, 2.5] and promotional adjectives by +3.1 [1.4, 4.6], while *innovat-* fell by −4.5 [−6.2, −2.7]. The control word *method* also rose (+2.5 [1.0, 4.0]) and the interval for *experiment* included zero (+1.1 [−0.3, 2.5]), so changes of one or two points at this boundary are within the range of the controls. The larger jumps are concentrated in collaborative schemes, for which Horizon Europe introduced new work programmes: *sustainab-* +7.5 [5.3, 9.7], *stakeholder* +6.1 [4.2, 8.3], *co-creation/co-design* +5.6 [4.2, 7.0], *societal* +3.5 [2.0, 5.1] and *resilien-* +2.8 [1.5, 4.2], against +3.0 [1.0, 5.0] for *method*. In individual schemes the jumps were in interdisciplinarity vocabulary (+3.8 [2.4, 5.2]) and organisational *multi-* (+2.0 [1.0, 3.0]), and *innovat-* fell (−6.5 [−8.3, −5.0]). Research content does not turn over in two years, so the vocabulary appears to follow the wording of programme documents. Appendix Table C3 gives every estimate with its confidence interval and the numbers of documents.

Vocabulary shifts at EU programme boundaries and, since 2023, with LLM-typical words



The 2025 reversal at NSF. In early 2025, after US executive orders on DEI, NSF was reported to have screened its funded projects for DEI-related keywords (Times Higher Education 2025). Court filings in a later lawsuit by University of California researchers showed that federal agencies used keyword screens when they terminated grants (CalMatters 2026). Compared with 2023-24 (18,943 abstracts), NSF abstracts from 2025 (7,362) used *equity/equitable* 88% less often, *minority/minorities* 86% less, *women* 84% less, *underrepresented* 82% less, *inclusive/inclusion* 74% less, *diversity/diverse* 53% less and *broadening participation* 53% less (Table 10, Figure 11). *Climate* fell by 63% and *AI/artificial intelligence* rose by 71%. Among the control words *method* rose by 14% and the others changed by less than 5%. Over the same years EU objectives used *diversity/diverse* 41% more often and *equity/equitable* 71% more often, from a low base. NIH abstracts moved in the same direction as NSF, less strongly: *equity/equitable* −58%, *minority/minorities* −69%, *underrepresented* −68%, *inclusive/inclusion* −37% and

diversity/diverse –22%. The NIH figures agree in direction with the 25% fall in diversity language that Mehta and Jena (2025) report for NIH awards from January 2024 to June 2025.

Table 10. Use of selected words in NSF abstracts, 2023-24 and 2025


Word	Share 2023-24	Share 2025	Relative change [95% CI]	Without education-oriented awards	2025 reweighted to the 2023-24 subfield mix
equity/equitable	6.4%	0.8%	–88% [–91, –84]	–87%	–89%
minority/minorities	8.1%	1.1%	–86% [–89, –83]	–87%	–85%
women	4.4%	0.7%	–84% [–88, –79]	–84%	–84%
underrepresented	19.2%	3.5%	–82% [–84, –79]	–82%	–80%
inclusive/inclusion	8.2%	2.2%	–74% [–78, –69]	–73%	–72%
diversity/diverse	29.2%	13.7%	–53% [–56, –50]	–51%	–49%
broadening participation	5.7%	2.7%	–53% [–60, –45]	–55%	–43%
climate	12.2%	4.6%	–63% [–67, –58]		
AI/artificial intelligence	11.4%	19.5%	+71% [61, 82]		
method (control)	32.9%	37.4%	+14% [10, 18]	+14%	+12%
experiment-(control)	21.4%	22.4%	+5% [–1, 10]	+6%	+3%
data (control)	42.4%	42.8%	+1% [–2, 4]	0%	+3%

Notes: n = 18,943 abstracts in 2023-24 and 7,362 in 2025. Confidence intervals from the log risk ratio. Education-oriented awards (4,624 in 2023-24 and 1,575 in 2025) are those in the subfield Education or with three or more education words. The last column uses the 57 subfields with at least 20 awards in 2023-24 and 10 in 2025, among awards with a topic.

The drops are not driven by the composition of the portfolio. They are almost unchanged without education-oriented awards, and within the 57 continuing subfields, with 2025 reweighted to the 2023-24 mix, they range from –43% (*broadening participation*) to –89% (*equity/equitable*). Topic coverage did not fall in 2025: the share of NSF awards with an OpenAlex topic was 28% in 2023, 30% in 2024 and 34% in 2025. The year-on-year total variation distance of the topic mix of those awards was 7.1% at the level of fields and 14.8% at the level of subfields. Sampling alone would produce about 4.4% and 11.6% for samples of these sizes, so the net movement was 2.7 and 3.2 points. In 2016-2024 the net movement

ranged from -0.7 to 4.6 points for fields and from 0.3 to 5.2 points for subfields (Figure 11b). The vocabulary changed together across many words while the topic mix moved within its usual range.

2025: flagged words fall sharply while the topic mix of NSF awards moves within its usual range



LLM-typical words. The share of documents with at least one of the 11 LLM-typical words departed sharply from its 2013-2022 trend (Figure 10b). In EU objectives it rose from 4.1% in 2022 to 9.8% in 2025, an excess over the trend of 5.5 points in 2025. In articles it rose from 2.2% to 9.5%, an excess of 7.2 points. In NSF abstracts it rose from 3.2% to 6.5% in 2024, an excess of 3.0 points, and fell back to 5.0% in 2025 (excess 1.5). NIH abstracts, which are short in OpenAlex, rose from 0.8% to 1.5% (excess 0.7).

In 2013-19, documents with these words had about as many umbrella words as other documents in the grant corpora, and slightly more in articles. The ratio of mean umbrella density in documents with LLM-typical words to that in documents without them was 0.93 in the EU, 1.00 in NSF and 0.83 in NIH in 2013-19, and 1.28 in articles. In 2024-25 the ratios were 1.06, 1.18, 1.02 and 1.78. Adjusted for length within year-by-scheme cells (year-by-field for articles), the difference in umbrella words per 100 tokens rose from +0.01 [-0.04, 0.06] to +0.13 [0.05, 0.20] in the EU, from +0.01 [-0.02, 0.04] to +0.15 [0.09, 0.22] in NSF and from +0.05 [0.03, 0.08] to +0.28 [0.26, 0.31] in articles; in NIH it stayed near zero. These estimates use a marker list from which every umbrella word and goal verb was removed. In the first version *multifaceted* was counted both as an LLM-typical word and as organisational *multi-*, which inflated the ratios (for NIH, from 1.15 to 1.45 then, against 0.83 to 1.02 now). LLM-assisted writing may therefore be adding to umbrella-word inflation in EU, NSF and article texts. The evidence is correlational and covers only two years.

H7 is supported. The vocabulary jumps at programme boundaries within fixed scheme types, collapses within a year after a political instruction while the research portfolio changes little, and has begun to cluster with LLM-typical words.

7. Discussion

7.1 Costs of umbrella-word inflation, and how large they are

The results support treating the growth of umbrella words as carrying costs, but the costs differ in size and in how firmly they are established. We discuss five, from the best to the least supported.

1. Dependence on outside actors and reversibility (H7). This is the clearest result. The vocabulary turned over within one or two years of programme changes, within collaborative and individual EU schemes alike, and within one year of a political instruction, while the topic mix of NSF awards moved within its usual range. Where the words come from funders more than from research communities, the values they are meant to express can be withdrawn along with the words. In 2025 the same vocabulary served federal agencies as keywords for screening grants. Justifications that rest on umbrella words did little to protect the practices they described.
2. Promises in place of objects (H2b). In NSF abstracts *diversity* moved from an object of study to a goal that applicants promise to increase: dictionaries, a window test and word embeddings agree. This is consistent at scale with Ahmed's (2012) account of non-performative statements. The shift is specific. In articles *diversity* stayed biological, and *sustainable* lost its resource objects without turning into a socio-economic word. Much of the spread of umbrella words across fields (H2a) is also what a boundary object would show, and we do not count it as a cost.
3. Loss of discriminating power (H1, H3, H4). About four in five NSF abstracts and more than nine in ten EU collaborative objectives now contain at least one umbrella word, so the mere presence of these words no longer distinguishes projects. The further costs to readers are small. Umbrella words carry at most about 4% of the pairwise similarity between documents, and documents dense in them are less specific in the rest of their text by at most a tenth of a standard deviation (0.03 to 0.06 nats of mean IDF per SD). The extra search cost for reviewers is therefore modest on average and concentrated in the most umbrella-dense documents.
4. A loose link between organisational claims and organisation (H5). Collaboration vocabulary appears in about one in eight single-beneficiary EU projects, and most consortia do not use it at all. Within collaborative schemes, however, its use rises with the number of partners and with international participation. The vocabulary is loosely coupled to formal structure. The test also measures organisations and countries, which are not disciplines, so it says little about whether claims of interdisciplinarity are met.
5. No value as a quality signal (H6). Umbrella density was weakly related to outcomes, positively in articles and EU projects and negatively in NSF awards, and always far less than budget or team size. If writing umbrella words helps applicants win funding, that advantage does not show up as later productivity or impact. Because only funded projects are observed, this cannot be read as an effect of the words themselves.

LLM-assisted writing may add to the trend: since 2023, documents with LLM-typical words carry more umbrella words than other documents, also after adjusting for length (Section 6.7). The evidence is correlational and the period is short.

7.2 Alternative explanations and caveats

Umbrella words also do useful work. As boundary objects they let people from different fields and sectors work together (Star and Griesemer 1989), and their growth partly reflects real interest in social

problems and new fields of application. Our field-concentration measures fall whenever a word comes into legitimate use in many fields, which is why we read them as evidence of diffusion only.

Three findings are harder to explain by legitimate new applications alone. In NSF abstracts the objects of *diversity* shifted from biological things to institutional goals. Some of the vocabulary appears in projects whose formal structure does not match it. And the vocabulary reversed within a year, independently of research content.

The growth of *diversity* in NSF abstracts also reflects a policy aim: the Broader Impacts review criterion and efforts to broaden the participation of under-represented groups. This study does not assess that aim. It shows that when the aim was required as vocabulary, the vocabulary became separable from practice and politically reversible.

7.3 Implications for practice, and their trade-offs

- Operationalisation for named commitments. Application forms can ask applicants to attach an indicator, an object and a responsible party to the few umbrella commitments that a project lists among its deliverables, such as a sustainability plan or a diversity target. Applying this to every occurrence would be impractical, since most documents contain several umbrella words, and it would add to the reporting burden and invite new boilerplate or synonyms, the templating that H3 measures. A short, structured section for named commitments avoids both.
- Monitoring with care. Funders can measure how far the vocabulary of their calls is copied into applications. In Horizon 2020 and Horizon Europe, applicants are told to address the "expected outcomes" and "scope" of a topic, so overlap with the call partly reflects required relevance. Monitoring tells funders about their own role as a source of vocabulary. It does not measure applicant quality.
- Reviewer guidance. Following Millar et al. (2022) on promotional language, reviewer guidance can state that umbrella density is not a signal of quality and that claims of collaboration or interdisciplinarity should be checked against the consortium, the fields involved and the budget.
- Practice over wording. Diversity and sustainability can be assessed as practices to be achieved and recorded, such as hiring, training or facilities. Records of practice are also harder to remove in a political reversal than wording is.

8. Limitations

1. Period and coverage. EU data start with FP7 in 2007, so FP1 to FP6 are missing. Long-run change is represented by NSF, which is mostly a domestic funder. Documents from international development (World Bank, USAID, OECD CRS) were not usable in the form available to us: USAID records are accounting summaries and Gates Foundation records are single sentences.
2. Participant data end in 2017. The CORDIS mirror covers FP7 and the Horizon 2020 projects signed by 2017, so the structural test of H5 misses the later Horizon 2020 projects and all of Horizon Europe, the period in which *co-creation*, interdisciplinarity vocabulary and organisational *multi-* rose most. The comparison of collaborative and individual schemes extends the test to 2026 only through funding rules, which fix the number of beneficiaries of individual schemes at one.

3. What the structural variables measure. The number of beneficiary organisations measures the size of a consortium, and the number of countries its international reach. Neither measures cognitive or disciplinary diversity: one university can host several disciplines in one project, and ten partners can belong to one discipline. Single-beneficiary grants can also involve real collaboration with partners who do not sign the grant, such as secondment hosts. The H5 results show how far the vocabulary follows formal structure. They do not show that projects that use it fail to collaborate, or that projects that avoid it do not.
4. Funded projects only. Without rejected applications we cannot estimate the effect of umbrella words on the probability of funding, and the selection on funding described in Section 2.3 can bias the outcome associations of H6 in either direction.
5. Dictionary method. Regular-expression families cannot fully resolve polysemy (biological *diversity*, physical *impact*, pharmacological *synergy*). Collocate classes come from hand-built dictionaries, and part of the collocates remain unclassified (Section 6.2 reports the shares). The embedding check uses static vectors per period. Contextual embeddings would allow sense-level measurement.
6. Article sample. Row groups were drawn at random from OpenAlex update-date partitions, so the sample within each year is a cluster sample and its field mix can differ from that of all articles. Records whose abstract field holds template text are 2.7% of the 1980s sample and under 1% from 2000; they lower early umbrella shares slightly and do not affect the growth ratios materially. Author country data are often missing for early years.
7. Outcome measures. Linked publication counts and FWCI depend on the accuracy of award-to-publication links in OpenAlex. Award topics are missing for about two thirds of NSF awards and three fifths of EU projects, and the subsample with topics has fewer linked publications than the rest (Section 6.6).
8. Causal inference. H4 to H6 and parts of H7 describe associations. The 2025 analysis combines two mechanisms we cannot separate. Some awards may have been cancelled, withheld or had their abstracts edited after keyword screening (ex post), and applicants and programme staff may have rewritten proposals to avoid flagged words before the award (ex ante). Mehta and Jena (2025) observed deletions in renewed NIH grants, which points to the second mechanism, but our NSF data do not link successive versions of the same project. The decline also continued within non-education subfields (Section 6.7), but finer programme categories that OpenAlex does not record, such as INCLUDES or ADVANCE, could still have shrunk.

9. Conclusion

Since the 1980s, the use of umbrella words in research-project documents has grown far faster than vocabulary in general. In NSF abstracts and articles most normative and evaluative families grew by multiples of about 5 to more than 100 (*excellence/excellent* is the exception) and outpaced control words and frequency-matched ordinary words, also after allowing for longer abstracts. The organisational families grew more modestly: in NSF abstracts, once length is held constant, no faster than the fastest-growing neutral nouns. In the EU the vocabulary grew within collaborative schemes, where *sustainab-* rose from 15% to 39% of objectives and *resilien-* from 1% to 13% in under twenty years. In individual schemes *collaborat-* and organisational *multi-* declined. The pooled EU series looked flat only because individual schemes took a growing share of projects.

The growth came with diffusion across fields, which boundary objects also show, and with a narrower set of costs that the data support more firmly. The clearest is dependence on outside actors: the vocabulary moves with funders' programmes and with politics, and in 2025 it was withdrawn from NSF abstracts within a year while the research portfolio barely changed. In NSF abstracts *diversity* moved from an object of study to a goal to be promised, as accounts of non-performative speech would expect. Other costs are small or conditional. Documents dense in umbrella words are slightly less specific; organisational claims follow formal structure only loosely; and umbrella density says little about outcomes.

The NSF case of 2025 comes close to a natural experiment and indicates that the vocabulary can change much faster than research content.

Four lines of further work follow. The analysis can be extended to EU programmes FP1 to FP6 and to World Bank project documents, and compared with Japanese KAKEN grants (with terms such as 多様性, 持続可能, 連携, 協働 and 共創). Participant data for later Horizon 2020 and Horizon Europe projects would allow the structural test to cover the period in which organisational vocabulary rose most. Data that include rejected applications would allow an estimate of effects on funding decisions. Contextual embeddings would allow sense-level measurement of meaning change.

References

- Ahmed, S. (2012). *On Being Included: Racism and Diversity in Institutional Life*. Duke University Press.
- Bensaude-Vincent, B. (2014). The politics of buzzwords at the interface of technoscience, market and society: The case of 'public engagement in science'. *Public Understanding of Science*, 23(3), 238-253. <https://doi.org/10.1177/0963662513515371>
- Brand, F. S., & Jax, K. (2007). Focusing the meaning(s) of resilience: Resilience as a descriptive concept and a boundary object. *Ecology and Society*, 12(1), 23. <https://doi.org/10.5751/ES-02029-120123>
- Bromham, L., Dinnage, R., & Hua, X. (2016). Interdisciplinary research has consistently lower funding success. *Nature*, 534, 684-687. <https://doi.org/10.1038/nature18315>
- Brysbaert, M., Warriner, A. B., & Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods*, 46(3), 904-911. <https://doi.org/10.3758/s13428-013-0403-5>
- Cornwall, A. (2007). Buzzwords and fuzzwords: Deconstructing development discourse. *Development in Practice*, 17(4-5), 471-484. <https://doi.org/10.1080/09614520701469302>
- Cornwall, A., & Eade, D. (Eds.). (2010). *Deconstructing Development Discourse: Buzzwords and Fuzzwords*. Practical Action Publishing / Oxfam.
- DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48(2), 147-160. <https://doi.org/10.2307/2095101>
- Hamilton, W. L., Leskovec, J., & Jurafsky, D. (2016). Diachronic word embeddings reveal statistical laws of semantic change. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics* (pp. 1489-1501). <https://doi.org/10.18653/v1/P16-1141>
- Jacobs, J. A., & Frickel, S. (2009). Interdisciplinarity: A critical assessment. *Annual Review of Sociology*, 35, 43-65. <https://doi.org/10.1146/annurev-soc-070308-115954>

- Kobak, D., González-Márquez, R., Horvát, E.-Á., & Lause, J. (2025). Delving into LLM-assisted writing in biomedical publications through excess vocabulary. *Science Advances*, 11(27), eadt3813. <https://doi.org/10.1126/sciadv.adt3813>
- Laclau, E. (1996). Why do empty signifiers matter to politics? In *Emancipation(s)* (pp. 36-46). Verso.
- Leahey, E., Beckman, C. M., & Stanko, T. L. (2017). Prominent but less productive: The impact of interdisciplinarity on scientists' research. *Administrative Science Quarterly*, 62(1), 105-139. <https://doi.org/10.1177/0001839216665364>
- Luo, Y., Jurafsky, D., & Levin, B. (2019). From insanely jealous to insanely delicious: Computational models for the semantic bleaching of English intensifiers. In *Proceedings of the 1st International Workshop on Computational Approaches to Historical Language Change* (pp. 1-13). <https://doi.org/10.18653/v1/W19-4701>
- Markowitz, D. M. (2019). What words are worth: National Science Foundation grant abstracts indicate award funding. *Journal of Language and Social Psychology*, 38(3), 264-282.
- Mehta, N., & Jena, A. B. (2025). Changes in diversity language in National Institutes of Health grant awards: Observational study. *BMJ*, 391, e087222. <https://doi.org/10.1136/bmj-2025-087222>
- Meyer, J. W., & Rowan, B. (1977). Institutionalized organizations: Formal structure as myth and ceremony. *American Journal of Sociology*, 83(2), 340-363. <https://doi.org/10.1086/226550>
- Millar, N., Batalo, B., & Budgell, B. (2022). Trends in the use of promotional language (hype) in abstracts of successful National Institutes of Health grant applications, 1985-2020. *JAMA Network Open*, 5(8), e2228676. <https://doi.org/10.1001/jamanetworkopen.2022.28676>
- Millar, N., Batalo, B., & Budgell, B. (2023). Promotional language (hype) in abstracts of publications of National Institutes of Health-funded research, 1985-2020. *JAMA Network Open*, 6(12), e2348706. <https://doi.org/10.1001/jamanetworkopen.2023.48706>
- Moretti, F., & Pestre, D. (2015). Bankspeak: The language of World Bank reports. *New Left Review*, 92, 75-99. <https://newleftreview.org/issues/II92/articles/franco-moretti-dominique-pestre-bankspeak.pdf>
- Plavén-Sigray, P., Matheson, G. J., Schiffler, B. C., & Thompson, W. H. (2017). The readability of scientific texts is decreasing over time. *eLife*, 6, e27725. <https://doi.org/10.7554/eLife.27725>
- Priem, J., Piwowar, H., & Orr, R. (2022). OpenAlex: A fully-open index of scholarly works, authors, venues, institutions, and concepts. arXiv:2205.01833. <https://arxiv.org/abs/2205.01833>
- Redclift, M. (2005). Sustainable development (1987-2005): An oxymoron comes of age. *Sustainable Development*, 13(4), 212-227. <https://doi.org/10.1002/sd.281>
- Star, S. L., & Griesemer, J. R. (1989). Institutional ecology, 'translations' and boundary objects: Amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. *Social Studies of Science*, 19(3), 387-420. <https://doi.org/10.1177/030631289019003001>
- Vinkers, C. H., Tijdink, J. K., & Otte, W. M. (2015). Use of positive and negative words in scientific PubMed abstracts between 1974 and 2014: Retrospective analysis. *BMJ*, 351, h6467. <https://doi.org/10.1136/bmj.h6467>

News sources

- CalMatters (2026, 21 July). White House admits it used keywords to kill billions worth of California research grants. Republished by KPBS and LAist. <https://calmatters.org/education/higher-education/2026/07/research-grants-california-2/>
- Times Higher Education (2025, 5 February; updated 12 February). Trump's anti-DEI research witch-hunt 'absurd' and 'immoral'. <https://www.timeshighereducation.com/news/trumps-anti-dei-research-witch-hunt-absurd-and-immoral>

Appendix A. Lexicon

All expressions are applied to lower-cased text; the code is in `src/lexicon.py`. Counts are numbers of matches; a document uses a family if it has at least one match. The three types are those of Section 1.

Normative/substantive (6 families):

- diversity/diverse: `\bdivers(?:e|ely|ity|ities|ification|ified|ify|ifying)\b`
- sustainab-: `\bsustainab(?:le|ly|ility|ilities)\b`
- resilien-: `\bresilien(?:ce|t|cy)\b`
- inclusive/inclusion: `\binclusi(?:ve|vely|veness|vity)\b|\binclusion\b(?:!\s+(?:and\s+exclusion\s+)?(?:criteri|bod(?:y|ies)|complex|compound))\b`
- empower-: `\bempower(?:s|ed|ing|ment)?\b`
- societal: `\bsocietal\b`

Organisational/process (8 families):

- collaborat-: `\bcollaborat(?:e|es|ed|ing|ion|ions|ive|ively|or|ors)\b`
- multi- (organisational): tokens in the list below
- inter-/trans-/cross-disciplinary: `\b(?:inter|trans|cross)-?disciplin(?:ary|arity|e|es)\b`
- co-creation/co-design: `\bco-?(?:creat(?:e|es|ed|ing|ion)|design(?:s|ed|ing)?|produc(?:e|es|ed|ing|tion))\b`
- partnership: `\bpartnerships?\b`
- cooperation: `\bco-?operat(?:ion|ions|ive)\b`
- stakeholder: `\bstake-?holders?\b`
- capacity building: `\bcapacit(?:y|ies)[- ](?:building|development|strengthening)\b`

Evaluative (5 families):

- innovat-: `\binnovat(?:ion|ions|ive|ively|iveness|e|es|ed|ing|or|ors)\b`
- impact: `\bimpact(?:s|ed|ing|ful)?\b`
- excellence/excellent: `\bexcellen(?:ce|t)\b`
- transformative: `\btransformati(?:ve|onal)\b`
- holistic: `\bholistic(?:ally)?\b`

Reference group outside the index:

- promotional adjectives (reference): `\b(?:novel|unprecedented|ground-?breaking|cutting-?edge|breakthroughs?|pioneering|world-?class|game-?chang(?:er|ers|ing)|paradigm[-`

]shift(?:s|ing)?|revolutioni[sz](?:e|es|ed|ing)|revolutionary|state-of-the-art|unique)\b

- synerg-: \bsynerg(?:y|ies|istic|istically|etic)\b
- engag-: \bengag(?:e|es|ed|ing|ement|ements)\b

Control words:

- measur- (control): \bmeasur(?:e|es|ed|ing|ement|ements)\b
- data (control): \bdata\b
- method (control): \bmethods?\b
- experiment- (control): \bexperiment(?:s|al|ally)?\b
- sample (control): \bsamples?\b
- temperature (control): \btemperatures?\b
- the (control): \bthe\b
- approach (control): \bapproach(?:es)?\b
- analysis (control): \banalys(?:is|es)\b
- process (control): \bprocess(?:es)?\b
- structure/structural (control): \bstructur(?:e|es|al)\b
- effect (control): \beffects?\b

LLM-typical words:

\b(?:delv(?:e|es|ed|ing)|showcas(?:es|ed|ing)|underscor(?:e|es|ed|ing)|intricate|intricacies|pivotal|meticulous(?:ly)?|commendable|realms?|notably|invaluable)\b. None of these matches an umbrella family or a goal verb; the code checks this.

Organisational *multi-* tokens: multi-actor, multi-actors, multi-agency, multi-benefit, multi-country, multi-cultural, multi-disciplinarity, multi-disciplinary, multi-institutional, multi-lateral, multi-national, multi-partner, multi-perspective, multi-professional, multi-sector, multi-sectoral, multi-sectorial, multi-stakeholder, multi-stakeholders, multi-thematic, multiactor, multiagency, multicountry, multicultural, multiculturalism, multidisciplinarity, multidisciplinary, multiinstitutional, multilateral, multinational, multipartner, multiprofessional, multisector, multisectoral, multistakeholder. Excluded *multi-* words: multiple, multiples, multiplet, multiplets, multiplex, multiplexed, multiplexes, multiplexing, multiplication, multiplicative, multiplicities, multiplicity, multiplied, multiplier, multipliers, multiplies, multiply, multiplying, multipole, multipoles, multitude, multitudes. All other *multi-* words with at least eight letters count only in *multi-* (all), which is not part of the index.

Appendix B. Dictionaries for collocates, goal verbs and embedding seeds

The code is in `src/a2_dilution.py` and `src/a2b_embeddings.py`. A collocate of "X diversity" is classified as a goal verb first, then as biological, social or generic; a collocate of "sustainable X" as technical first, then as environmental or resource, then as socio-economic, institutional or durability. Collocates in none of the lists are unclassified.

- Biological collocates of *diversity*: alleles, allelic, alpha, animal, animals, antibody, antigenic, arthropod, avian, bacteria, bacterial, bat, beta, biodiversity, biological, bird, birds, brillouin, butterfly, cell, cellular, chemical, clonal, compositional, coral, crop, crops, cryptic, cultivars, diazotroph, ecological, ecosystem, epigenetic, eukaryotic, evolutionary, faunal, fish, floral,

floristic, forest, functional, fungal, fungi, gamma, gene, genes, genetic, genome, genomic, genotypes, genotypic, germplasm, gut, habitat, haplotype, host, ichthyofaunal, immune, insect, interspecific, intraspecific, invertebrate, isoform, isolates, landraces, landscape, latitudinal, lichen, livestock, mammal, mammalian, marine, metabolic, microalgae, microbes, microbial, microbial-community, microbiome, microbiota, microorganisms, molecular, morphological, natural, neuronal, nucleotide, organismal, organisms, parasite, pathogen, phenotypic, phylogenetic, phytoplankton, plant, plants, pollinator, population, predator, protein, proteome, proteomic, receptor, repertoire, sequence, shannon, shannon-wiener, soil, spatial, species, strain, strains, structural, substrate, taxa, taxonomic, tcr, trait, tree, trophic, varieties, vegetation, vertebrate, viral, virus.

- Social collocates of *diversity*: actors, audiences, backgrounds, board, communities, countries, country, cultural, culturally, cultures, demographic, disciplinary, economic, educational, epistemic, equity, ethnic, ethnically, ethnolinguistic, faculty, gender, geographic, geographical, groups, human, identities, inclusion, individuals, institutional, intellectual, language, languages, linguistic, members, minorities, minority, national, nationalities, ontological, opinions, participant, participants, partners, people, perspectives, racial, racially, regional, religious, sexual, social, socio-cultural, socioeconomic, staff, stakeholders, stem, student, students, team, teams, underrepresented, users, viewpoints, voices, women, workforce, workplace.
- Generic collocates of *diversity*: activities, amazing, applications, approaches, areas, array, aspects, astounding, broad, conditions, considerable, contexts, data, enormous, environments, exceptional, experiences, expertise, extraordinary, fields, forms, full, great, greater, high, huge, ideas, immense, increasing, incredible, interests, large, methods, needs, options, outcomes, products, range, remarkable, rich, scales, sectors, services, set, settings, sheer, significant, skills, solutions, sources, staggering, strategies, striking, surprising, tools, topics, tremendous, types, unexpected, uses, variety, vast, views, ways, wide.
- Goal verbs: achieve, achieving, address, addressing, advance, advancing, assure, broaden, broadening, celebrate, embrace, encourage, encouraging, enhance, enhancing, ensure, ensuring, expand, expanding, foster, fostering, fosters, improve, improving, increase, increasing, maintain, maintaining, promote, promotes, promoting, respecting, support, supporting, supports, value, valuing.
- Environmental or resource objects of *sustainable*: adaptation, agricultural, agriculture, alternative, alternatives, aquaculture, arctic, aviation, batteries, battery, bio-based, bioeconomy, bioenergy, biofuels, biomass, biomaterials, blue, building, buildings, built, carbon, catalysis, catalysts, catalytic, chemical, chemicals, chemistry, circular, cities, city, coastal, concrete, conservation, construction, consumption, cooling, crop, crops, demolition, development, diet, dietary, diets, ecological, ecosystem, ecosystems, electricity, energy, environment, environmental, farming, fashion, feedstock, feedstocks, fertilizers, fisheries, fishery, food, forest, forestry, forests, fuel, fuels, grazing, green, groundwater, harvest, harvesting, heating, housing, hydrogen, intensification, irrigation, land, landscape, landscapes, lifestyles, lighting, livelihood, livelihoods, management, manufacturing, material, materials, mobility, nuclear, ocean, packaging, pest, plastics, polymer, polymers, power, processes, production, rangeland, raw, recycling, resource, resources, soil, solar, sources, sourcing, supply, synthesis, synthetic, tourism, transport, transportation, urban, urbanization, use, utilization, waste, wastewater, water, yield, yields.

- Technical objects of *sustainable*: applications, architecture, computing, cyberinfrastructure, design, electronics, engineered, engineering, manufacturing, method, methods, nanotechnology, pathway, pipeline, process, product, route, smart, software, technologies, technology. (Checked before the environmental list, so *manufacturing* counts as technical.)
- Socio-economic, institutional or durability objects of *sustainable*: access, ai, approach, approaches, banking, benefits, business, businesses, capacity, career, careers, change, collaboration, collaborations, communities, community, competitive, competitiveness, cooperation, data, digital, economic, economy, education, educational, effect, effects, employment, engagement, entrepreneurship, eu, european, exploitation, finance, financial, financing, framework, funding, future, futures, global, governance, growth, health, healthcare, human, impact, impacts, implementation, improvement, improvements, infrastructure, innovation, innovations, institutional, institutions, integration, international, interventions, investment, investments, jobs, learning, local, logistics, long-term, manner, market, marketing, means, mechanism, mechanisms, model, models, national, network, networks, open, open-source, operation, operations, organisations, organization, organizations, outcomes, partnership, partnerships, pathways, peace, performance, platform, policies, policy, practices, procurement, products, professional, profitability, program, programme, programmes, programs, public, recovery, regional, research, results, security, services, social, societies, society, socio-economic, solution, solutions, strategies, strategy, structures, success, support, system, systems, training, transformation, transition, transitions, value, viable, way, ways, workforce, world.
- Embedding seed words for *diversity*: species, genetic, biodiversity, microbial, taxa, phylogenetic, ecological, habitat (first pole); gender, ethnic, racial, cultural, inclusion, equity, minority, underrepresented (second pole).
- Embedding seed words for *sustainable*: environmental, ecological, energy, conservation, resources, renewable, agricultural, water (first pole); economic, growth, business, financial, institutional, career, organizational, competitiveness (second pole).

Appendix C. Supplementary tables

Tables generated by `scripts/appendix_tables.py` from `results/tables`; the tables for all EU objectives and for NIH are in `results/tables/growth_summary.csv` and `growth_length_adjusted.csv`.

Table C1. Field concentration of article abstracts, all families (bias-corrected KL*, nats; last two columns: KL* with the field mix fixed at 1980-2025, and the coefficient of variation of field usage rates, first and last period)

Family	1980-89	1990-99	2000-09	2010-19	2020-25	2020-25, fixed mix	CV first to last
diversity/diverse	0.276	0.265	0.226	0.189	0.123	0.133	0.72 to 0.53

Family	1980-89	1990-99	2000-09	2010-19	2020-25	2020-25, fixed mix	CV first to last
sustainab-	not estimated	1.004	0.711	0.480	0.417	0.425	1.70 to 1.01
resilien-	not estimated	0.344	0.379	0.355	0.291	0.315	1.12 to 0.95
inclusive/inclusi on	0.088	0.085	0.089	0.139	0.200	0.201	0.55 to 0.67
empower-	not estimated	1.105	0.735	0.557	0.431	0.493	1.59 to 1.10
societal	0.931	0.872	0.568	0.418	0.438	0.492	1.52 to 1.03
collaborat-	0.332	0.361	0.262	0.239	0.221	0.241	1.15 to 0.83
multi- (organisational)	0.647	0.497	0.393	0.275	0.264	0.307	1.48 to 0.88
inter-/trans- /cross- disciplinary	0.510	0.545	0.337	0.329	0.249	0.264	1.34 to 0.90
co-creation/co- design	not estimated	not estimated	not estimated	0.433	0.451	0.515	1.38 to 1.30
partnership	0.907	0.918	0.665	0.511	0.437	0.471	1.85 to 1.19
cooperation	0.257	0.295	0.279	0.277	0.311	0.319	0.73 to 0.92
stakeholder	not estimated	1.081	0.626	0.516	0.401	0.431	1.98 to 1.20
capacity building	not estimated	not estimated	0.655	0.554	0.479	0.569	1.30 to 1.23
innovat-	0.624	0.529	0.424	0.291	0.204	0.200	1.31 to 0.83
impact	0.283	0.239	0.129	0.098	0.083	0.090	0.85 to 0.50
excellence/exce llent	0.077	0.081	0.085	0.188	0.276	0.279	0.46 to 1.01
transformative	not estimated	1.012	0.847	0.582	0.327	0.348	2.21 to 1.00

Family	1980-89	1990-99	2000-09	2010-19	2020-25	2020-25, fixed mix	CV first to last
holistic	not estimated	0.538	0.340	0.268	0.207	0.240	1.23 to 0.79
promotional adjectives (reference)	0.062	0.049	0.070	0.068	0.063	0.060	0.40 to 0.38
measur- (control)	0.114	0.085	0.083	0.085	0.070	0.066	0.54 to 0.35
data (control)	0.050	0.045	0.050	0.047	0.058	0.061	0.37 to 0.41
method (control)	0.085	0.081	0.124	0.103	0.052	0.050	0.39 to 0.30
experiment- (control)	0.169	0.186	0.222	0.240	0.260	0.228	0.59 to 0.61
sample (control)	0.085	0.087	0.138	0.105	0.067	0.062	0.42 to 0.41
temperature (control)	0.513	0.604	0.653	0.689	0.746	0.634	1.24 to 1.23
approach (control)	0.136	0.130	0.065	0.051	0.046	0.044	0.64 to 0.33
analysis (control)	0.036	0.020	0.016	0.014	0.015	0.016	0.32 to 0.23
process (control)	0.076	0.086	0.067	0.065	0.061	0.061	0.38 to 0.38
structure/struct ural (control)	0.129	0.132	0.141	0.165	0.186	0.182	0.54 to 0.67
effect (control)	0.077	0.092	0.123	0.101	0.081	0.076	0.44 to 0.39

Table C2. Logit models with year fixed effects only (first version), linked FP7 and Horizon 2020 projects (n = 36,177; coefficients in log odds, [95% CI])

Outcome	log participants	two or more countries	log tokens
Collaboration vocabulary	0.016 [-0.041, 0.072]	0.858 [0.727, 0.989]	0.518 [0.398, 0.638]
Interdisciplinarity	-0.051 [-0.110, 0.009]	0.181 [0.048, 0.314]	0.621 [0.506, 0.736]

Outcome	log participants	two or more countries	log tokens
vocabulary			

Table C3. Difference-in-differences at EU programme boundaries relative to NSF abstracts (percentage points, 95% bootstrap CI). Raw EU change in parentheses.

Family	All EU, 2014	Collaborative, 2014	Individual, 2014	All EU, 2021	Collaborative, 2021	Individual, 2021
diversity/diverse	−1.2 [−2.3, 0.0] (−0.2)	−0.4 [−2.0, 1.0] (+0.6)	−2.3 [−3.7, −1.1] (−1.3)	−0.7 [−1.9, 0.5] (+2.8)	−0.1 [−1.9, 1.5] (+3.4)	−0.8 [−2.2, 0.5] (+2.7)
sustainable	+0.6 [−0.5, 1.7] (+0.6)	+2.3 [0.3, 4.5] (+2.3)	+2.7 [1.7, 3.8] (+2.8)	+1.3 [0.2, 2.5] (+3.4)	+7.5 [5.3, 9.7] (+9.6)	−1.9 [−2.9, −0.7] (+0.3)
multi-organisational)	−0.4 [−1.4, 0.5] (−0.4)	+0.4 [−1.1, 2.0] (+0.4)	−0.7 [−2.0, 0.5] (−0.7)	+1.8 [0.8, 2.7] (+1.6)	+1.9 [0.1, 3.4] (+1.6)	+2.0 [1.0, 3.0] (+1.7)
collaborative	+0.7 [−0.7, 2.1] (−1.0)	+2.8 [0.5, 4.8] (+1.2)	+1.6 [0.0, 3.1] (−0.1)	−1.1 [−2.3, 0.1] (+0.3)	−0.4 [−2.6, 2.0] (+0.9)	−1.2 [−2.5, 0.0] (+0.2)
innovative	+7.1 [5.5, 8.6] (+7.9)	+10.9 [8.5, 13.3] (+11.7)	+11.1 [9.5, 12.7] (+11.9)	−4.5 [−6.2, −2.7] (−2.2)	−0.3 [−2.7, 2.2] (+2.0)	−6.5 [−8.3, −4.9] (−4.2)
impact	+12.6 [10.9, 14.1] (+0.7)	+14.1 [11.9, 16.5] (+2.2)	+13.2 [11.5, 15.0] (+1.4)	+1.4 [−0.1, 3.0] (+3.4)	+4.0 [1.7, 6.5] (+6.0)	+0.3 [−1.4, 2.0] (+2.3)
stakeholder	+0.0 [−0.8, 0.9] (+0.4)	+4.2 [2.4, 6.0] (+4.6)	+0.7 [0.3, 1.2] (+1.1)	+1.3 [0.4, 2.4] (+2.0)	+6.1 [4.2, 8.3] (+6.8)	−0.9 [−1.7, −0.2] (−0.2)
inclusive/inclusion	−0.1 [−0.6, 0.4] (+0.1)	+0.5 [−0.3, 1.4] (+0.6)	−0.1 [−0.5, 0.4] (+0.1)	−0.7 [−1.5, 0.1] (+1.9)	+1.9 [0.6, 3.1] (+4.4)	−2.0 [−2.7, −1.4] (+0.5)
resilient	−0.2 [−0.6, 0.3] (+0.6)	+0.6 [−0.2, 1.4] (+1.4)	−0.4 [−0.9, 0.1] (+0.4)	+0.4 [−0.3, 1.2] (+2.1)	+2.8 [1.5, 4.2] (+4.5)	−0.7 [−1.4, −0.1] (+0.9)
inter-/trans-/cross-disciplinary	+1.9 [0.8, 3.0] (+1.1)	+3.5 [2.0, 5.0] (+2.7)	+0.3 [−0.9, 1.6] (−0.5)	+2.9 [1.8, 4.2] (+3.3)	+1.2 [−0.5, 2.9] (+1.7)	+3.8 [2.4, 5.2] (+4.2)
societal	+0.6 [−0.1, 1.3] (+0.8)	+1.7 [0.4, 2.9] (+1.9)	+1.0 [0.4, 1.6] (+1.2)	+1.5 [0.6, 2.4] (+1.5)	+3.5 [2.0, 5.1] (+3.6)	+0.4 [−0.5, 1.2] (+0.4)
excellence/excellent	+0.1 [−0.7, 0.9] (−0.4)	+1.7 [0.3, 3.0] (+1.3)	−0.5 [−1.5, 0.5] (−1.0)	+0.6 [−0.0, 1.4] (+0.5)	+2.3 [0.7, 3.7] (+2.1)	−0.1 [−1.0, 0.6] (−0.3)
co-creation/co-	+0.1 [−0.1, 0.2] (0.0)	+0.5 [0.1, 1.0] (+0.5)	+0.0 [−0.1, 0.1] (0.0)	+1.6 [1.0, 2.2] (+1.6)	+5.6 [4.2, 7.0] (+5.6)	−0.6 [−0.9, −0.3] (−0.3)

Family	All EU, 2014	Collaborative, 2014	Individual, 2014	All EU, 2021	Collaborative, 2021	Individual, 2021
design	0.4] (+0.2)	1.0] (+0.6)	0.2] (+0.1)	2.2] (+2.3)	7.0] (+6.4)	-0.2] (+0.2)
transformative	+0.6 [0.1, 1.0] (+0.2)	+0.7 [0.1, 1.2] (+0.3)	+0.6 [0.1, 1.0] (+0.2)	-0.5 [-1.0, 0.2] (+0.8)	+0.0 [-0.8, 0.7] (+1.3)	-0.8 [-1.4, -0.2] (+0.5)
promotional adjectives (reference)	+0.5 [-1.2, 2.2] (+3.8)	-5.3 [-7.9, -3.0] (-2.1)	+1.1 [-1.1, 3.2] (+4.4)	+3.1 [1.4, 4.6] (+3.1)	+3.8 [1.4, 6.5] (+3.8)	+2.0 [-0.1, 4.1] (+2.0)
experiment-(control)	-3.6 [-5.0, -2.1] (-2.1)	-3.0 [-4.6, -1.4] (-1.5)	-6.9 [-8.5, -5.3] (-5.3)	+1.1 [-0.3, 2.5] (+2.0)	-2.3 [-3.9, -0.7] (-1.4)	+3.0 [1.5, 4.5] (+3.9)
method (control)	-4.3 [-5.9, -2.9] (-1.0)	-6.7 [-8.7, -4.6] (-3.4)	-4.4 [-6.3, -2.6] (-1.1)	+2.5 [1.0, 4.0] (+2.0)	+3.0 [1.0, 5.0] (+2.5)	+2.3 [0.5, 4.1] (+1.8)

Notes: documents before and after each boundary: all EU H2020 (2014): 8,820 and 8,274 EU documents, 19,820 and 20,250 NSF; all EU Horizon Europe (2021): 10,034 and 8,602 EU documents, 19,588 and 18,771 NSF; EU collaborative H2020 (2014): 4,235 and 2,777 EU documents, 19,820 and 20,250 NSF; EU collaborative Horizon Europe (2021): 3,609 and 2,986 EU documents, 19,588 and 18,771 NSF; EU individual H2020 (2014): 4,381 and 5,420 EU documents, 19,820 and 20,250 NSF; EU individual Horizon Europe (2021): 6,359 and 5,456 EU documents, 19,588 and 18,771 NSF. The 2014 estimate for *impact* reflects the removal of the "Broader impacts" heading from NSF abstracts in 2014 (Section 6.7).

Table C4 (NSF abstracts, 1987-91 to 2021-25). Growth of every family

Family	Share early to late	Density ratio	OR per decade, binomial + HAC	OR, WLS + HAC	OR, length-adjusted	Percentile among matched words
diversity/diverse	3.48% to 26.12%	4.52	1.87 [1.67, 2.08]	1.87	1.51 [1.38, 1.64]	98
sustainable	0.09% to 10.18%	90.97	2.68 [2.29, 3.15]	2.52	2.28 [2.10, 2.47]	100
resilient	0.06% to 6.20%	77.80	4.22 [3.86, 4.61]	4.12	3.69 [3.47, 3.93]	100
inclusive/inclusion	0.54% to 6.79%	7.38	2.33 [1.96, 2.75]	2.38	1.95 [1.62, 2.34]	100
empower-	0.03% to 3.23%	48.66	3.18 [2.91, 3.48]	3.09	2.72 [2.40, 3.09]	100
societal	0.22% to 6.41%	13.77	2.36 [1.92,	2.18	1.93 [1.75,	100

Family	Share early to late	Density ratio	OR per decade, binomial + HAC	OR, WLS + HAC	OR, length-adjusted	Percentile among matched words
			2.90]		2.13]	
collaborat-	8.53% to 27.32%	1.71	1.43 [1.27, 1.62]	1.42	1.19 [1.13, 1.26]	92
multi-(organisational)	1.07% to 5.30%	2.43	1.39 [1.18, 1.64]	1.34	1.12 [1.04, 1.21]	95
inter-/trans-/cross-disciplinary	1.75% to 13.06%	4.03	1.62 [1.38, 1.91]	1.57	1.33 [1.23, 1.43]	98
co-creation/co-design	0.01% to 2.51%	134.98	5.87 [4.55, 7.58]	5.18	5.25 [4.22, 6.53]	100
partnership	0.40% to 8.88%	11.60	1.91 [1.67, 2.17]	1.82	1.60 [1.50, 1.70]	100
cooperation	4.80% to 1.83%	0.18	0.72 [0.68, 0.77]	0.72	0.63 [0.61, 0.65]	1
stakeholder	0.01% to 5.40%	210.43	3.22 [2.76, 3.76]	3.13	2.78 [2.55, 3.03]	100
capacity building	0.00% to 0.75%		2.65 [2.18, 3.22]	2.33	2.36 [2.07, 2.69]	
innovat-	2.35% to 24.13%	6.31	1.89 [1.74, 2.05]	1.86	1.74 [1.66, 1.82]	100
impact	5.97% to 44.51%	4.49	2.02 [1.58, 2.59]	1.91	1.59 [1.39, 1.81]	100
excellence/excellent	2.75% to 2.83%	0.44	1.07 [0.97, 1.18]	1.07	0.83 [0.79, 0.87]	19
transformative	0.02% to 5.65%	106.90	3.02 [2.02, 4.51]	1.76	2.43 [1.93, 3.06]	100
holistic	0.04% to 2.01%	21.53	2.88 [2.56, 3.23]	2.83	2.43 [2.24, 2.64]	100
promotional	10.83% to	2.56	1.74 [1.62,	1.73	1.38 [1.35,	

Family	Share early to late	Density ratio	OR per decade, binomial + HAC	OR, WLS + HAC	OR, length-adjusted	Percentile among matched words
adjectives (reference)	43.36%		1.86]		1.41]	
measur- (control)	15.27% to 21.12%	0.73	1.16 [1.11, 1.21]	1.16	0.94 [0.92, 0.96]	38
data (control)	21.26% to 43.10%	1.53	1.38 [1.34, 1.42]	1.38	1.13 [1.10, 1.16]	83
method (control)	17.19% to 33.68%	1.07	1.32 [1.29, 1.35]	1.32	1.09 [1.06, 1.11]	73
experiment- (control)	21.31% to 22.57%	0.58	1.03 [1.00, 1.05]	1.03	0.82 [0.80, 0.84]	15
sample (control)	5.41% to 6.81%	0.67	1.08 [1.03, 1.14]	1.08	0.87 [0.84, 0.89]	27
temperature (control)	7.08% to 9.01%	0.77	1.07 [1.05, 1.09]	1.07	0.87 [0.84, 0.89]	30
approach (control)	10.91% to 34.50%	1.76	1.54 [1.47, 1.62]	1.54	1.21 [1.20, 1.23]	92
analysis (control)	21.45% to 25.59%	0.59	1.09 [1.05, 1.12]	1.09	0.87 [0.86, 0.89]	31
process (control)	20.36% to 32.38%	0.81	1.21 [1.16, 1.26]	1.21	0.93 [0.91, 0.94]	54
structure/structural (control)	21.89% to 25.55%	0.66	1.07 [1.04, 1.09]	1.07	0.84 [0.83, 0.86]	30
effect (control)	13.95% to 16.22%	0.56	1.07 [1.03, 1.11]	1.07	0.81 [0.79, 0.83]	24
umbrella index (19 families)	25.76% to 82.14%	3.48	2.31 [1.99, 2.69]	2.27	1.78 [1.65, 1.93]	
normative/substantive	4.37% to 42.37%	8.21	2.21 [2.02, 2.41]	2.20		
organisational/process	14.04% to 44.47%	1.97	1.54 [1.38, 1.73]	1.53		

Family	Share early to late	Density ratio	OR per decade, binomial + HAC	OR, WLS + HAC	OR, length-adjusted	Percentile among matched words
evaluative	10.60% to 58.25%	4.15	2.08 [1.72, 2.51]	2.03		

Table C4 (Articles, 1980-84 to 2021-25). Growth of every family

Family	Share early to late	Density ratio	OR per decade, binomial + HAC	OR, WLS + HAC	OR, length-adjusted	Percentile among matched words
diversity/diverse	0.75% to 4.73%	4.81	1.56 [1.45, 1.68]	1.57	1.44 [1.37, 1.52]	98
sustainable	0.02% to 4.27%	184.35	2.68 [2.24, 3.21]	2.66	2.53 [2.18, 2.94]	100
resilient	0.02% to 1.29%	45.78	3.29 [2.49, 4.34]	3.16	2.97 [2.35, 3.77]	100
inclusive/inclusion	0.42% to 1.91%	4.47	1.48 [1.33, 1.66]	1.49	1.37 [1.27, 1.47]	97
empowerment	0.02% to 0.81%	28.58	1.98 [1.67, 2.34]	1.90	1.80 [1.62, 2.01]	100
societal	0.09% to 0.54%	4.00	1.68 [1.45, 1.94]	1.67	1.52 [1.37, 1.69]	97
collaborative	0.25% to 2.31%	7.43	1.64 [1.50, 1.79]	1.64	1.52 [1.42, 1.63]	99
multi-organisational)	0.18% to 1.03%	4.07	1.48 [1.40, 1.55]	1.47	1.39 [1.33, 1.45]	98
inter-/trans-/cross-disciplinary	0.13% to 0.70%	4.33	1.63 [1.49, 1.78]	1.61	1.52 [1.42, 1.62]	97
co-creation/co-design	0.01% to 0.18%	16.28	2.80 [2.54, 3.09]	2.46	2.53 [2.24, 2.86]	100
partnership	0.05% to	6.27	1.47 [1.30,	1.41	1.38 [1.29,	99

Family	Share early to late	Density ratio	OR per decade, binomial + HAC	OR, WLS + HAC	OR, length-adjusted	Percentile among matched words
	0.48%		1.66]		1.49]	
cooperation	0.59% to 1.20%	1.74	1.18 [1.15, 1.20]	1.18	1.12 [1.10, 1.15]	81
stakeholder	0.00% to 1.26%	154.76	2.65 [2.41, 2.91]	2.54	2.40 [2.25, 2.56]	100
capacity building	0.00% to 0.12%	47.82	2.50 [2.11, 2.96]	2.29	2.25 [1.92, 2.63]	99
innovat-	0.46% to 3.79%	7.60	1.73 [1.56, 1.91]	1.74	1.67 [1.54, 1.80]	100
impact	1.93% to 12.70%	4.72	1.71 [1.55, 1.89]	1.71	1.58 [1.49, 1.67]	100
excellence/excellent	1.04% to 2.05%	1.41	1.17 [1.13, 1.21]	1.17	1.10 [1.09, 1.12]	78
transformative	0.02% to 0.55%	25.69	2.83 [2.10, 3.83]	2.75	2.62 [2.02, 3.41]	100
holistic	0.03% to 0.61%	11.69	2.07 [1.80, 2.38]	2.05	1.88 [1.70, 2.08]	100
promotional adjectives (reference)	2.33% to 10.33%	3.27	1.45 [1.40, 1.49]	1.45	1.35 [1.33, 1.38]	
measur- (control)	15.00% to 12.48%	0.64	0.94 [0.92, 0.96]	0.94	0.87 [0.86, 0.87]	20
data (control)	12.51% to 25.66%	1.71	1.25 [1.17, 1.34]	1.25	1.17 [1.14, 1.21]	88
method (control)	11.26% to 33.22%	1.99	1.42 [1.37, 1.47]	1.42	1.34 [1.31, 1.38]	98
experiment- (control)	9.05% to 10.22%	0.80	1.03 [1.02, 1.04]	1.03	0.98 [0.97, 0.99]	46
sample (control)	4.30% to 9.30%	1.48	1.22 [1.17, 1.28]	1.22	1.13 [1.12, 1.14]	90

Family	Share early to late	Density ratio	OR per decade, binomial + HAC	OR, WLS + HAC	OR, length-adjusted	Percentile among matched words
temperature (control)	6.12% to 4.68%	0.58	0.93 [0.92, 0.94]	0.93	0.90 [0.89, 0.91]	27
approach (control)	3.98% to 16.30%	3.00	1.44 [1.34, 1.55]	1.44	1.35 [1.29, 1.42]	99
analysis (control)	10.22% to 27.32%	2.11	1.33 [1.27, 1.39]	1.33	1.24 [1.22, 1.27]	96
process (control)	6.98% to 14.57%	1.49	1.22 [1.20, 1.23]	1.22	1.16 [1.14, 1.17]	93
structure/structural (control)	8.26% to 11.65%	1.06	1.08 [1.06, 1.10]	1.08	1.04 [1.01, 1.06]	70
effect (control)	14.34% to 18.33%	0.95	1.06 [1.04, 1.08]	1.06	0.98 [0.97, 0.99]	45
umbrella index (19 families)	5.63% to 29.04%	5.47	1.62 [1.50, 1.75]	1.62	1.52 [1.45, 1.60]	
normative/substantive	1.32% to 11.68%	9.09	1.77 [1.59, 1.97]	1.78		
organisational/process	1.16% to 6.04%	4.73	1.50 [1.43, 1.57]	1.50		
evaluative	3.42% to 18.09%	4.40	1.61 [1.47, 1.77]	1.62		

Table C4 (EU collaborative schemes, 2008-10 to 2023-25). Growth of every family

Family	Share early to late	Density ratio	OR per decade, binomial + HAC	OR, WLS + HAC	OR, length-adjusted	Percentile among matched words
diversity/diverse	6.60% to 14.34%	2.25	1.88 [1.41, 2.51]	1.90	1.89 [1.45, 2.46]	99
sustainab-	14.73% to 39.19%	3.14	2.47 [2.09, 2.92]	2.47	2.50 [2.17, 2.87]	99
resilien-	1.06% to 12.75%	11.94	6.53 [5.64, 7.56]	6.33	6.57 [5.54, 7.80]	100

Family	Share early to late	Density ratio	OR per decade, binomial + HAC	OR, WLS + HAC	OR, length-adjusted	Percentile among matched words
inclusive/inclusion	1.81% to 9.16%	5.98	3.31 [2.95, 3.72]	3.30	3.35 [2.89, 3.89]	99
empower-	1.20% to 7.81%	6.97	3.92 [3.25, 4.72]	3.82	3.95 [3.14, 4.95]	100
societal	2.87% to 10.49%	3.79	2.22 [1.80, 2.73]	2.17	2.24 [1.91, 2.63]	99
collaborat-	16.48% to 24.97%	1.51	1.34 [1.09, 1.66]	1.36	1.35 [1.13, 1.60]	96
multi-(organisational)	9.27% to 15.00%	1.67	1.46 [1.41, 1.52]	1.46	1.47 [1.39, 1.55]	91
inter-/trans-/cross-disciplinary	5.79% to 15.07%	2.88	2.19 [1.82, 2.64]	2.20	2.20 [1.82, 2.66]	100
co-creation/co-design	0.24% to 11.66%	47.41	10.38 [6.44, 16.72]	8.96	10.53 [7.48, 14.82]	100
partnership	5.58% to 6.69%	1.23	1.05 [0.92, 1.20]	1.05	1.05 [0.91, 1.22]	71
cooperation	12.65% to 6.69%	0.37	0.63 [0.58, 0.68]	0.63	0.63 [0.57, 0.69]	5
stakeholder	12.67% to 26.45%	2.16	1.88 [1.72, 2.06]	1.89	1.91 [1.70, 2.14]	99
capacity building	1.39% to 4.39%	3.16	2.39 [1.97, 2.90]	2.32	2.40 [1.91, 3.01]	98
innovat-	26.81% to 52.10%	2.54	2.10 [1.60, 2.76]	2.08	2.11 [1.73, 2.58]	100
impact	25.24% to 37.74%	1.58	1.51 [1.39, 1.65]	1.51	1.54 [1.38, 1.71]	96
excellence/excellent	7.31% to 10.74%	1.77	1.37 [1.25, 1.51]	1.36	1.38 [1.21, 1.57]	91

Family	Share early to late	Density ratio	OR per decade, binomial + HAC	OR, WLS + HAC	OR, length-adjusted	Percentile among matched words
transformative	0.10% to 3.57%	45.16	8.03 [6.64, 9.71]	7.36	8.03 [6.75, 9.56]	100
holistic	2.02% to 6.91%	3.42	2.29 [2.01, 2.60]	2.27	2.30 [2.05, 2.59]	98
promotional adjectives (reference)	32.33% to 44.23%	1.47	1.36 [1.27, 1.47]	1.36	1.37 [1.28, 1.47]	
measur-(control)	12.16% to 11.32%	0.77	0.94 [0.90, 0.99]	0.94	0.95 [0.90, 1.00]	39
data (control)	21.39% to 33.82%	1.73	1.53 [1.45, 1.62]	1.53	1.55 [1.43, 1.67]	98
method (control)	18.09% to 17.10%	0.85	0.94 [0.86, 1.03]	0.94	0.94 [0.88, 1.01]	44
experiment-(control)	11.19% to 7.27%	0.58	0.77 [0.70, 0.84]	0.77	0.77 [0.71, 0.83]	12
sample (control)	3.42% to 2.09%	0.54	0.65 [0.59, 0.73]	0.67	0.66 [0.55, 0.79]	14
temperature (control)	4.63% to 3.06%	0.64	0.75 [0.69, 0.82]	0.76	0.76 [0.68, 0.86]	23
approach (control)	29.18% to 37.82%	1.34	1.31 [1.23, 1.40]	1.31	1.32 [1.21, 1.45]	87
analysis (control)	19.97% to 17.87%	0.82	0.90 [0.88, 0.93]	0.90	0.91 [0.87, 0.95]	37
process (control)	28.55% to 28.52%	0.93	0.98 [0.92, 1.04]	0.98	0.98 [0.93, 1.04]	60
structure/structural (control)	15.89% to 9.55%	0.52	0.71 [0.65, 0.78]	0.72	0.72 [0.66, 0.78]	9
effect (control)	10.94% to 8.38%	0.76	0.82 [0.76, 0.89]	0.82	0.83 [0.75, 0.92]	20
umbrella index (19 families)	74.63% to 92.65%	2.21	2.62 [2.28, 3.02]	2.58	2.64 [2.30, 3.04]	

Family	Share early to late	Density ratio	OR per decade, binomial + HAC	OR, WLS + HAC	OR, length-adjusted	Percentile among matched words
normative/substantive	24.71% to 60.26%	3.69	2.80 [2.31, 3.39]	2.78		
organisational/process	44.38% to 62.07%	1.65	1.68 [1.50, 1.87]	1.67		
evaluative	49.02% to 73.19%	2.14	2.13 [1.86, 2.44]	2.12		

Table C4 (EU individual schemes, 2008-10 to 2023-25). Growth of every family

Family	Share early to late	Density ratio	OR per decade, binomial + HAC	OR, WLS + HAC	OR, length-adjusted	Percentile among matched words
diversity/diverse	8.05% to 10.94%	1.29	1.39 [1.13, 1.71]	1.38	1.39 [1.16, 1.66]	85
sustainab-	2.59% to 12.05%	4.81	2.41 [1.99, 2.92]	2.36	2.41 [2.06, 2.83]	100
resilien-	0.56% to 3.48%	6.10	3.97 [3.53, 4.46]	3.80	3.98 [3.52, 4.50]	100
inclusive/inclusion	0.68% to 2.23%	3.11	2.08 [1.86, 2.31]	2.00	2.09 [1.74, 2.50]	98
empower-	0.16% to 1.41%	7.76	2.82 [2.08, 3.81]	2.55	2.82 [2.21, 3.61]	99
societal	1.00% to 4.51%	4.38	2.48 [1.93, 3.18]	2.39	2.49 [2.06, 3.00]	99
collaborat-	11.10% to 8.01%	0.70	0.85 [0.75, 0.98]	0.85	0.84 [0.74, 0.97]	24
multi-(organisational)	7.65% to 6.24%	0.74	0.80 [0.72, 0.90]	0.80	0.80 [0.73, 0.87]	34
inter-/trans-/cross-disciplinary	8.98% to 12.39%	1.36	1.40 [1.23, 1.59]	1.39	1.39 [1.24, 1.57]	88

Family	Share early to late	Density ratio	OR per decade, binomial + HAC	OR, WLS + HAC	OR, length-adjusted	Percentile among matched words
co-creation/co-design	0.05% to 0.77%	21.45	6.22 [4.74, 8.15]	5.06	6.40 [4.50, 9.09]	100
partnership	0.33% to 0.83%	1.99	1.38 [0.83, 2.28]	1.22	1.37 [0.97, 1.94]	95
cooperation	2.85% to 1.46%	0.53	0.58 [0.50, 0.69]	0.59	0.58 [0.50, 0.67]	9
stakeholder	0.44% to 2.45%	4.32	2.33 [1.53, 3.55]	2.05	2.34 [1.76, 3.12]	100
capacity building	0.07% to 0.03%	0.45	0.82 [0.42, 1.59]	0.69		14
innovat-	9.17% to 24.40%	2.54	1.70 [1.16, 2.49]	1.60	1.69 [1.31, 2.18]	100
impact	14.53% to 25.99%	1.86	1.56 [1.45, 1.67]	1.55	1.55 [1.46, 1.65]	100
excellence/excellent	5.57% to 3.02%	0.52	0.63 [0.53, 0.75]	0.63	0.62 [0.54, 0.71]	2
transformative	0.21% to 2.67%	11.94	8.23 [5.89, 11.49]	7.57	8.24 [5.74, 11.83]	100
holistic	0.44% to 2.06%	4.58	2.87 [2.56, 3.21]	2.79	2.90 [2.49, 3.38]	99
promotional adjectives (reference)	40.45% to 60.13%	1.72	1.67 [1.57, 1.77]	1.67	1.66 [1.59, 1.72]	
measur-(control)	14.21% to 13.06%	0.84	0.93 [0.89, 0.96]	0.93	0.92 [0.87, 0.97]	59
data (control)	20.78% to 25.58%	1.36	1.25 [1.19, 1.32]	1.25	1.24 [1.16, 1.33]	91
method (control)	23.65% to 27.13%	1.05	1.19 [1.07, 1.32]	1.19	1.18 [1.08, 1.29]	87
experiment-	23.30% to	0.76	0.88 [0.72,	0.88	0.87 [0.74,	39

Family	Share early to late	Density ratio	OR per decade, binomial + HAC	OR, WLS + HAC	OR, length-adjusted	Percentile among matched words
(control)	19.05%		1.08]		1.02]	
sample (control)	4.95% to 5.37%	1.06	1.02 [0.95, 1.08]	1.01	1.01 [0.93, 1.09]	61
temperature (control)	4.69% to 5.81%	1.02	1.12 [1.05, 1.18]	1.11	1.11 [1.01, 1.23]	73
approach (control)	33.57% to 40.00%	1.15	1.22 [1.07, 1.40]	1.22	1.21 [1.09, 1.35]	84
analysis (control)	25.33% to 23.01%	0.85	0.97 [0.88, 1.08]	0.97	0.96 [0.87, 1.07]	59
process (control)	31.63% to 28.71%	0.84	0.90 [0.87, 0.93]	0.90	0.89 [0.85, 0.94]	60
structure/structural (control)	26.38% to 20.80%	0.72	0.82 [0.70, 0.97]	0.82	0.82 [0.73, 0.92]	35
effect (control)	20.53% to 18.46%	0.82	0.93 [0.81, 1.07]	0.93	0.92 [0.83, 1.02]	58
umbrella index (19 families)	48.73% to 67.42%	1.69	1.69 [1.50, 1.89]	1.68	1.67 [1.53, 1.81]	
normative/substantive	12.13% to 28.72%	2.60	2.06 [1.93, 2.20]	2.05		
organisational/process	26.13% to 26.10%	0.96	1.05 [0.99, 1.12]	1.05		
evaluative	26.17% to 46.93%	1.95	1.69 [1.38, 2.08]	1.68		

Appendix D. Data availability and reproduction

- OpenAlex snapshot of 14 October 2026 (CC0): <s3://openalex/data/parquet/awards/> and <s3://openalex/data/parquet/works/>
- Mirror of the CORDIS FP7 and Horizon 2020 project files: github.com/spaghetti-open-data/eu-research-funding-network
- Concreteness ratings as distributed with the R package [doc2concrete](#)
- LLM excess-vocabulary list: github.com/berenslab/llm-excess-vocab

All code, tables and figures are in github.com/xuoz/claude, branch `claude/pensive-mccarthy-1x5og8`. The script `scripts/run_all.sh` reproduces every step from download to the appendix tables; it needs about 15 GB of disk and 15 GB of memory, and the analysis scripts must run one at a time.

Step	Scripts
Article sample	<code>src/fetch_works_sample.py</code> , <code>src/build_works_corpus.py</code>
Funding corpora	<code>src/build_award_corpora.py</code>
Lexicon and features	<code>src/lexicon.py</code> , <code>src/features.py</code>
H1	<code>src/a1_trends.py</code>
H2	<code>src/a2_dilution.py</code> , <code>src/a2b_embeddings.py</code>
H3 to H7	<code>src/a3_homogenization.py</code> to <code>src/a8_nsf_2025.py</code>
Appendix C	<code>scripts/appendix_tables.py</code>

Use of AI assistance. The analysis code and the text of this draft, including this revision, were prepared with an AI assistant (Claude) under the author's direction. Journals differ in how they require such assistance to be disclosed, and the disclosure should be adapted to the target venue.