Theory and Society (2026) 55:25
https://doi.org/10.1007/s11186-026-09690-2

The ideological orientation of academic social science research 1960–2024

James Manzi[1]

Received: 7 October 2025 / Accepted: 20 February 2026

Abstract

This study analyzes approximately 600,000 English-language social science abstracts published between 1960 and 2024 to estimate the long-run ideological orientation of disciplinary research output. Large language models (LLMs) were applied to each abstract using a fixed 2025 U.S. ideological spectrum, enabling consistent coding across six decades. Five key findings emerged. First, roughly 90 percent of politically relevant social science articles leaned left 1960–2024, and the mean political stance of every social science discipline was left-of-center every year during the period. Second, all disciplines showed leftward movement between 1990 and 2024. Third, policy-proximal disciplines generally showed limited rightward moderation between roughly 1970 and 1990, though policy-distal disciplines did not. Fourth, disciplines with greater leftward orientation generally displayed greater ideological homogeneity Fifth, sociocultural content was more consistently left-leaning than economic content, and that gap widened over time. Robustness checks using a wide assortment of alternative datasets and analytical methodologies indicated that these findings were unlikely to be artifacts of idiosyncratic assumptions. Methodologically, the study demonstrates the capacity of LLM-based text classification to deliver reliable, large-scale ideological measurement over time, a task previously impractical with human coding alone. Taken together, the analysis provides the first systematic, cross-disciplinary evidence of the long-run political orientation of anglophone social science scholarship, revealing both the persistence and the intensification of its leftward tendencies, particularly in sociocultural domains.

Keywords Ideology · Knowledge production · Academic discourse · Computational text analysis · Sociology of science · Large language models

✉ James Manzi
james.manzi@sociology.ox.ac.uk

[1] University of Oxford, Oxford, UK

Published online: 16 March 2026

Introduction

Numerous research methodologies have shown that U.S. college and university faculty tend to be on the political left. Over the past seventy-plus years, rigorous surveys have consistently shown that faculty disproportionately self-identify as left-of-center relative to the general public (Gross & Simmons, 2014a; Hamilton & Hargens, 1993; Honeycutt, 2024; Lazarsfeld & Thielens, 1958; Lipset, 1976; Stolzenberg et al., 2019). Klein and Stern provided specific survey evidence for the left orientation of the stated policy views of social science faculty (Klein & Stern, 2005b). Disproportionate representation of registered Democrats versus Republicans has also been observed among faculty in a wide variety of contexts: at eleven California universities (Cardiff & Klein, 2005); in economics, history, journalism, law, and psychology departments (Langbert et al., 2016); in law schools (Bonica et al., 2018); in 116 colleges across 30 states (Langbert & Stevens, 2021); and at a major state university system (Chin et al., 2025). Political donations records also show disproportionate donations to Democrats versus Republicans among university faculty in a range of contexts (Kaurov et al., 2022; Langbert, 2020; McGinnis et al., 2005). Recent analysis of faculty behavior on social media shows a similar ideological skew (Havey & Her, 2025). There is a scholarly consensus that U.S. university faculties are broadly to the left of the general public. Analyses positioned as critical responses to extreme claims of faculty radicalization have generally conceded this point (Mariani & Hewitt, 2008; Zipp & Fenwick, 2006). This overall body of research also generally agrees that humanities and social sciences faculties are further to the left than those in other disciplines, and that the skew is driven more by center-left than far-left faculty.

Many of these survey-based studies have reached conflicting conclusions about temporal trends, but most rely on survey series spanning fewer than twenty years. Honeycutt (2024) combined several such studies to evaluate long-term ideological trends from 1969–2022 and concluded that the ratio of left-of-center/right-of-center faculty self-identification had risen from 1.7 to 6.9 over that period. Kaurov et al. (2022) observed that faculty political donations have become increasingly dominated by contributions to Democrats since 2000. Changes to survey methods and the sharp growth in the volume of faculty political donations during these years, however, argue for caution when using these measures to infer long-term trends in the underlying ideology of the professorate as a whole.

The impacts of this ideological orientation on faculty behavior are widely contested and appear nuanced. There is evidence that direct biases in classroom instruction, grading, and graduate-level admissions are either unobserved or at most small effects; however, mechanisms by which ideology might influence research have recently been demonstrated in controlled experiments. A series of studies present evidence that that faculty ideology does not translate into student indoctrination (Boysen et al., 2009; Burmila, 2021; Gross & Simmons, 2014a; Mariani & Hewitt, 2008; Woessner & Kelly-Woessner, 2009) or grading (Bar & Zussman, 2012; Musgrave & Rom, 2014). Gross and Simmons used a randomized experiment to argue for the absence of ideological bias in graduate school admissions (Gross & Simmons,

2014b). The same authors, however, also provided survey evidence that 71 percent of scholars believe that personal views should guide selection of research topics (Gross & Simmons, 2014a). Observational studies have found that author ideology is directly related to topic selection by economists (Jelveh et al., 2024) and to ideological stances adopted in legal scholarship (Chilton & Posner, 2015). Recent randomized experiments have been used to delineate a causal relationship between faculty ideology and academic research, as well as some channels by which this impact could be operationalized. One experiment demonstrated that historians gave higher quality ratings to abstracts that are consistent with their own orientation (Schiekiera & Niemeyer, 2025). Another observed that, when presented with identical datasets and asked to estimate the effect of immigration on support for social programs, research teams composed of pro-immigration researchers estimated more positive impacts of immigration on public support for social programs, while anti-immigration research teams reported more negative estimates (Borjas & Breznau, 2024).

In sum, prior work has established that U.S. social science faculties have tended to be ideologically left-leaning for many decades and that this orientation could plausibly affect their research output. However, there has been relatively limited comprehensive empirical analysis of the ideological orientation of that output itself. This is partly because characterizing the political content of hundreds of thousands of academic texts, especially in a consistent and fine-grained way, has historically been prohibitively expensive and time-intensive when relying on human coders. Several studies have used natural language processing methods, including dictionary-based classifiers, supervised models, and word embeddings, to estimate ideological position in scholarly and political texts (e.g., Grimmer & Stewart, 2013; Jelveh et al., 2024; Slapin & Proksch, 2008). While informative, these approaches often require substantial domain-specific tuning and may struggle to capture contextual nuance across diverse texts. More recently, large language models (LLMs) have demonstrated an emerging capacity to evaluate ideological content in complex documents with greater semantic sensitivity and generalizability (Knudsen et al., 2024; Le Mens & Gallego, 2025; O'Hagan & Schein, 2023).

This study shifts attention from the political identities of scholars to the ideological character of their published outputs. Abstracts provide a compact record of which problems disciplines treat as significant and how those problems are framed, argued, or justified. Treating the abstract as the unit of analysis allows us to evaluate how scholarly work is positioned when interpreted relative to a stable contemporary benchmark. The aim is not to recover how authors in earlier periods conceived their own political views, but to establish a consistent frame under which disciplinary outputs can be compared across time and domains.

The analysis reveals several empirical regularities that have not previously been demonstrated in systematic form. Politically relevant social science abstracts are heavily concentrated at center-left positions, and right-leaning content is rare. This pattern is not uniform: policy-proximal disciplines show partial moderation in the 1970s–1980s before turning leftward thereafter, while disciplines focused on identity, culture, and social meaning move leftward more continuously across the period.

Abstracts also tend to be more left-leaning on sociocultural issues than on economic issues, and the gap widens over time. Disciplines that lean further left also tend to show lower internal dispersion. Taken together, these results speak directly to debates about long-run disciplinary asymmetry, the channels through which ideology shapes topic selection and interpretive framing in research, and the evidence that ideological orientation is more visible in scholarly output than in classroom instruction or admissions settings.

Data and methods

Inclusion criteria

The OpenAlex (Priem et al., 2022) database was selected as the primary source of social science works used in this analysis. OpenAlex indexes more than 260 million scholarly works by harvesting Crossref, PubMed, arXiv and other major sources. It provides extensive metadata by work, including title, publication journal name and date, work type, language, and abstract text. OpenAlex has been shown to be broadly comparable to proprietary databases such as SCOPUS (Scopus, n.d.) and Web of Science ('WoS'; Clarivate, n.d.), with somewhat broader coverage but more limited completion in specific metadata categories (Alonso-Álvarez & van Eck, 2025; Culbert et al., 2025). The complete OpenAlex February 27, 2025 database snapshot was downloaded and used throughout the analysis. 1960 was chosen as the starting point since earlier decades had sparse counts of articles with abstracts deposited in OpenAlex that were relevant to 2025 political debates.

Analysis was focused on eleven social science disciplines: anthropology, communications, criminology, economics, ethnic studies, gender studies, political science, psychology, public administration, public health, and sociology. Selecting scholarly works that represent the central intellectual content of any of these disciplines inherently requires judgement, in part because the boundaries of most social science disciplines are contested.

One approach would be to rely on the OpenAlex categorization scheme. OpenAlex uses machine learning methods to score from 0 to 1 the strength of association of various 'concepts', including each of these eleven discipline names, with each work and with each journal as a whole. Article-level estimation appeared to be quite noisy as an indicator of social science discipline. As an illustration, all English-language articles with abstracts published 1960–2024 in the *American Sociological Review (ASR)*, considered to be one of the top two journals in the field, were collected. Approximately one-third of these articles had either no association or an association score of zero for the concept of sociology. Treating no associations score as zero, the mean sociology association score was 0.26. OpenAlex also generates concept association scores for each journal as a whole, and these were generally more stable. *ASR* had a journal-level sociology association score of 0.75, the highest association score for any of the eleven target disciplines with this journal. However, journals not generally considered core to these disciplines often had high association scores and quantitative indices of journal impact. As an example, *Academy of*

Management Review, which is typically categorized as a management journal, had a higher OpenAlex sociology association score and h-index impact score than did *ASR*.

To minimize researcher degrees of freedom while maintaining alignment with broadly-accepted views on discipline boundaries, the approach taken in the main analysis was to prompt GPT-4o to create lists of leading English language journals by discipline. This process identified 392 journals in total. Reviewing relevance and impact using Clarivate Analytics Journal Impact Factors (Journal citation reports, n.d.) and Scimago Journal Rank indicators (Scimago journal & country rank, n.d.) confirmed that these journals provided a reasonable representation of each discipline. 367 (94%) of the identified journal names could be matched with journal names appearing in OpenAlex through a combination of fuzzy matching and hand curation. The complete list of analyzed journals by discipline is provided in Supplementary Information (Exhibit S1). 599,194 works were identified as OpenAlex English-language articles with abstracts published 1960–2024 in any of these 367 social science journals (the 'Analyzed Abstracts'). Each work was assigned to a discipline based on the discipline of the journal in which it was published. Some journals were identified as relevant to more than one discipline (e.g., *Journal of Family Issues* was considered to be both a sociology and a gender studies journal), and in such cases, the works published by these journals were assigned fractionally to both disciplines when conducting discipline-level analysis.[1]

Later, analyses were executed to test the robustness of key findings to database selection and inclusion criteria, including journal ablation, systematically selecting works using OpenAlex discipline association scores at both the article and journal level without reference to the LLM-generated lists, and selecting works from the SCOPUS and WoS databases instead of OpenAlex.

Prompting method

In this study, ideology was operationalized as the relative positioning of a given text along a left–right spectrum of U.S. political thought. This includes economic and sociocultural dimensions, each characterized by their association with recognizable ideological actors and institutions. To ensure interpretive consistency over a 65-year period (1960 to 2024), political stance was evaluated against a fixed 2025 reference frame, derived from contemporary U.S. political categories.

This approach involved trade-offs. Using a static ideological scale, anchored to notions such as the 2025 political positions of Joe Biden, Elizabeth Warren, or the Heritage Foundation, ensured that the LLM's judgments were grounded in stable, well-defined anchor points. It avoided the problem of ideological drift, in which the meaning of 'left' or 'right' might evolve over time due to shifting partisan alignments or cultural contexts. However, this came at the cost of temporal anachronism. Texts written in earlier decades may have been judged by standards they would not have recognized.

While this may obscure how authors understood their own political commitments in historical context, it enabled direct comparability across time. That is, we could

[1] Calculation details are presented in Supplementary Information Exhibit S4.

assess not how 'left' a 1970s article was in reference to then-current ideological frameworks, but how it would have read to a 2025 observer applying today's ideological frameworks. This choice prioritized empirical consistency and interpretability over historical contextualization.

To minimize researcher degrees of freedom, GPT-4o was used to create a set of anchor points to define a 0–10 scale for the political spectrum in the United States in 2025 where 0 is far right, 10 is far left, and 5 is politically neutral. This scale is shown in Table 1.[2]

Table 1 Definition and anchors for political stance scores

Rating	Label	Politician Anchors	Media Outlet Anchors	Think Tank Anchors
0	Far Right	Steve King	American Free Press	VDARE Foundation
1	Hard Right	Lauren Boebert	The Epoch Times, Fox News (Opinion)	Center for Renewing America
2	Right Populist	Donald Trump	Washington Times	America First Policy Institute
3	Conservative	Ron DeSantis	National Review	Heritage Foundation
4	Center-Right	Tom Cotton	Wall Street Journal (Editorial)	American Enterprise Institute (AEI)
5	Centrist	Joe Manchin	Reuters/ USA Today, The Washington Post	Brookings Institution
6	Center-Left	Joe Biden	The Atlantic	Third Way
7	Progressive	Elizabeth Warren	The New Republic, The New York Times (Opinion)	Center for American Progress (CAP)
8	Leftist	Alexandria Ocasio-Cortez	The Nation	Roosevelt Institute
9	Radical left	Jamaal Bowman	Jacobin	People's Policy Project
10	Far Left	Chokwe Antar Lumumba	Monthly Review	Institute for policy Studies

[2] An expanded table with LLM rationales and confidence for anchor placement is provided as Exhibit S2 in Supplementary Information.

Throughout this paper, all terms denoting political stance – such as 'right' or 'left-leaning' – refer to LLM judgements based on this scale unless otherwise indicated.

The 'Prompting Method' used to query the LLM for ratings of political relevance and stance was defined as follows: (i) inject anchor points for each 0–10 rating value into the LLM system instructions and (ii) prompt the LLM to answer the following five questions about a text:

> Q1. Is the abstract directly relevant to any political or social debates in the United States in 2025?
>
> Q2. Based on the fixed 2025 U.S. political spectrum defined in your system instructions, where would you place the position taken in the abstract?
>
> Q3. According to the same 2025 political spectrum, would accepting the content of the abstract as true generally align with the political right or the political left in the United States?
>
> Q4. On the 2025 U.S. political spectrum, where would the abstract fall with respect to economic issues?
>
> Q5. On the 2025 U.S. political spectrum, where would the abstract fall with respect to social or cultural issues?

The complete combination of system instructions and prompt text is provided in Supplementary Information (Exhibit S3).

Using this approach, valid GPT-4o responses were obtained for 598,031 (99.8%) of the Analyzed Abstracts.

Separately, discipline-year publication volumes were estimated from the broader set of eligible works, irrespective of abstract availability. These volumes are used to compute discipline weights for aggregated multi-discipline analyses.[3]

Later, robustness checks were executed to evaluate the consistency of key findings when using alternative LLM models, alternative question wordings, alternative anchor points, full papers rather than abstracts, and prompts estimating political stance for each paper as it would have been understood in its year of publication rather than against a fixed 2025 scale.

Validation of prompting method

Before beginning the analysis, the reliability of the Prompting Method as a reasonable approach to measure political stance and relevance was validated using both external and internal methods.

Measurement of political stance was validated by applying the Prompting Method to material produced by the anchor think tanks. Recent substantive research articles were scraped from the homepages of 10 of the 11 anchor think tanks. The number of

[3] Calculation details are presented in Supplementary Information Exhibit S4.

articles by think tank ranged from 30 to 67. The exception was the VDARE Foundation (the most far-right think tank) whose website has been shut down; therefore, 20 prior VDARE articles were hand collected from the Internet Archive (Internet Archive, n.d.). GPT-4o was then used to create abstracts for all of these works in two steps. First, it was prompted to create a summary of each work, while removing all references to the think tank itself. Second, it was next presented with this list of sanitized summaries and prompted to modify each summary to an abstract of 150–200 words suitable for a contemporary social science journal without adding or omitting any information. The Prompting Method was then applied separately to each of these abstracts, and a mean overall political stance score was calculated for each think tank by averaging scores for Q2 for all abstracts derived from that think tank's works. These mean scores for blinded abstracts were highly predictive of the anchor point ratings for the think thanks that produced them ($r = 0.97$; $p < 0.01$).[4] That is, the method could correctly place the sanitized abstracts on the left–right spectrum.

The reliability of the LLM in determining political relevance was tested in part by applying this method to a 'control' discipline of theoretical mathematics, which presumptively should not be directly relevant to current political debates. The Prompting Method was applied to all English-language articles 1960–2024 with abstracts in the OpenAlex database published in *Annals of Mathematics*. Zero of these articles were identified as politically relevant. Further, when applied to the Analyzed Abstracts, the fraction of articles identified as directly relevant to 2025 U.S. political debates increased almost monotonically from fewer than 5 percent in 1960 to almost 50 percent in 2024 ($r = 0.99$; $p < 0.01$), despite the prompt not indicating the year in which the abstract was published.

The stability of LLM responses was checked by selecting a random sample of 1,000 of the Analyzed Abstracts, executing the Prompting Method twice on this sample, and then calculating several metrics for the inter-rater reliability between these two runs. First, categorical coding of political relevance ('yes' vs. 'no' in Q1) was evaluated using Cohen's Kappa, which adjusts for chance agreement (Cohen, 1960). The result ($\kappa = 0.949$) was 'almost perfect' (Landis & Koch, 1977). Second, the correlation of numeric stance scores between two independent runs was examined for the overall political stance score (Shrout & Fleiss, 1979). An extremely high Pearson correlation was observed ($r = 0.98$), indicating that the two scoring passes were nearly perfectly linearly related. Additionally, the two-way random effects intraclass correlation coefficient ('ICC') was also estimated for overall political stance scores. This measure likewise indicated excellent reliability between the two runs ($ICC = 0.98$). Taken together, these measures indicated that the Prompting Method produced stable outputs.

In addition to statistical methods, the validity of LLM scoring was sense-checked by directly examining a small number of prominent examples for facial plausibility. These examples were chosen systematically by selecting the ten politically relevant abstracts (those with an answer of 'yes' to Q1) with the highest OpenAlex citation counts. A summary of these ten works ranked by political stance score is displayed in Table 2.

[4] The p-value was estimated from a standard ordinary least squares regression (Virtanen et al., 2020), which tests the null hypothesis that the trend is zero.

Table 2 Example highly-cited works

Title	Journal	Publication Year	Political Stance Score
Hegemooic Masculinity	Gender & Society	2005	8
Language and Symbolic Power	Social Forces	1992	8
Asylums: Essays on the Social Situation of Mental Patients and Other Inmates	American Sociological Review	1962	7
Male–Female Wage Differentials in Urban Labor Markets	International Economic Review	1973	7
reasuring individual differences in implicit cognition: The implicit association test	Journal of Personality and Social Psychology	1998	7
The measurement of psychological androgyny	Journal of Consulting and Clinical Psychology	1974	7
Toward a New Conception of the Environment-Competitiveness Relationship	The Journal of Economic Perspectives	1995	7
Bank Runs, Deposit Insurance, and Liquidity	Journal of Political Economy	1983	6
reasuriog Economic Policy Uncertainty	The Quarterly Journal of Economics	2016	5
The Agenda-Setting Function of Mass Media	Public Opinion Quarterly	1972	5

It was not surprising to see 'Hegemonic Masculinity' published in *Gender & Society* and 'Language and Symbolic Power' in *Social Forces* rated as the furthest left of these highly-cited works or 'Measuring Economic Policy Uncertainty' in *The Quarterly Journal of Economics* and 'The Agenda-Setting Function of Mass Media' in *Public Opinion Quarterly* rated as the most centrist. All ten examples appeared facially plausible.

The LLM's rationale for political stance scores was also examined. The Prompting Method was modified slightly to require the LLM to explain its reasoning briefly, and this modified method was applied to each of these ten works. As an example of the results, the reasoning provided by the LLM for the 'Hegemonic Masculinity' overall political stance score was:

> *The abstract critiques traditional gender norms and hegemonic masculinity, advocating more fluid and egalitarian understandings of gender. It emphasizes social struggle, agency, and the reformulation of masculinities, which aligns with progressive views focused on structural change. The intersectional and multi-level approach to gender relations reflects left-leaning academic discourse. The stated aim of moving toward "gender democracy" signals a commitment to social justice and equality. Together, these elements place the article on the left side of the political spectrum.*

And for ‘Language and Symbolic Power’ it was:

> *The focus on linguistic exchanges, symbolic power, and political representation suggests a critical examination of social structures and power dynamics, which aligns with leftist academic discourse. The emphasis on concepts like "legitimate language," "symbolic power," and "political field" indicates an interest in how language and power intersect to shape social and political realities, a perspective often associated with left-leaning critical theory. This aligns with a progressive critique of societal institutions and power relations.*

The reasoning provided for the centrist-rated work ‘Measuring Economic Policy Uncertainty’ was:

> *The abstract does not explicitly take a political position; rather, it presents a research study on economic policy uncertainty and its effects on the economy. The study is empirical and focuses on developing an index to measure policy-related economic uncertainty, analyzing its impact on stock price volatility, investment, and employment. It does not advocate for specific policies or political ideologies but instead provides data-driven insights into how uncertainty affects economic variables. Given that the abstract is centered on economic analysis without a clear ideological stance, it would be placed in the center of the political spectrum.*

And for ‘The Agenda-Setting Function of Mass Media’, it was:

> *The abstract for "The Agenda-Setting Function of Mass Media" by McCombs and Shaw does not explicitly take a political position. Instead, it discusses the role of mass media in shaping public opinion by influencing the importance placed on topics. This concept is more analytical and descriptive rather than prescriptive or ideological. Therefore, it does not align with a specific point on the political spectrum. Given the focus on media influence without a clear political stance, it would be placed at a neutral position, which is 5 on the 0 (far right) to 10 (far left) scale.*

The reasoning approach used by the LLM appeared to be a realistic representation of how an informed observer would locate these works ideologically, rather than relying on simple heuristics such as word choice or explicit references to named ideologies or entities.

Main analysis methodology

The main analysis proceeded in four major steps. First was applying the Prompting Method to all Analyzed Abstracts and evaluating the distribution of overall political stance scores (Q2) by discipline over the period 1960–2024 as a whole. Second was charting the evolution of these scores by discipline from 1960 through 2024. Third was evaluating the ideological heterogeneity by discipline. Fourth was creating

weight-averaged social science scores across all disciplines for overall, economic, and sociocultural political stance scores (Q2, Q4, and Q5) and charting their evolution from 1960 to 2024.

Results

Upon application of the Prompting Method, 180,311 of the Analyzed Abstracts produced an answer of 'yes' to Q1. These were termed the 'Politically Relevant Abstracts', and they were the subject of subsequent analysis in Results. The political relevance filter in Q1 identifies explicit engagement with contemporary social or political debates as legible to a fixed 2025 ideological reference frame. This captures ideological expression as articulated through claims, arguments, and topics in the abstract text. It does not attempt to detect ideology embedded at deeper levels of disciplinary practice—such as theoretical commitments, methodological choices, problem framings, or epistemic vocabularies—even if such commitments may themselves be historically or politically conditioned. Later, robustness checks were executed to test the consistency of key findings when evaluating all Analyzed Abstracts.

Distributions of overall political stance scores (Q2) for the Politically Relevant Abstracts were calculated by discipline. Because the measurement is anchored to a contemporary U.S. 2025 ideological rubric, the patterns reported below describe changes in how social science abstracts are positioned relative to present-day categories of left, center, and right. They do not represent claims about how authors in earlier decades would have understood their own political commitments, nor about underlying ideological change in a neutral or absolute sense. The unit of analysis is the present-day legibility of texts, not their original intent. Figure 1 shows the mean

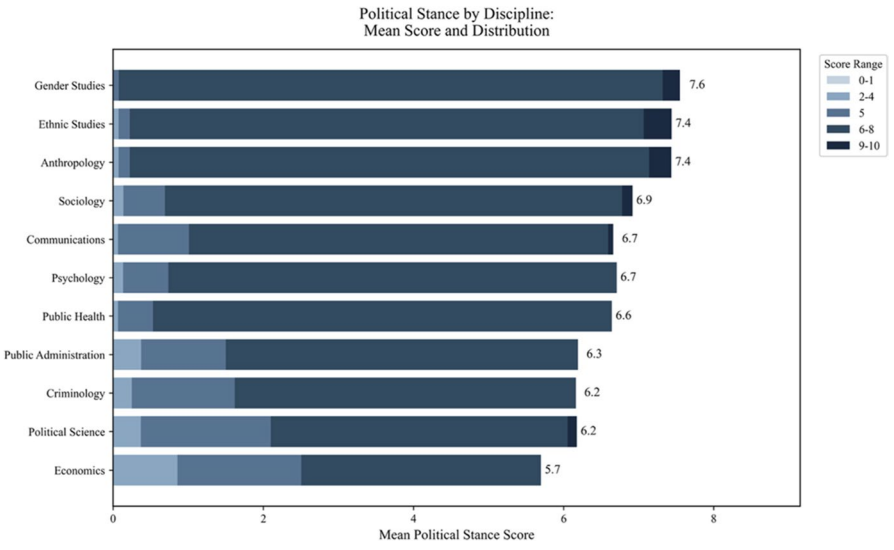

Fig. 1 Political stance by discipline: mean score and distribution

score by discipline. The percentage of articles by score range is shown by the shading of each bar.

Across all disciplines, abstracts were positioned left-of-center on the 2025 measurement frame, with mean scores ranging from 5.7 in economics to 7.6 in gender studies. These results were ordinally consistent with the self-reported faculty political affiliation for the disciplines of economics, political science, sociology, and anthropology in Klein and Stern (2005a) and for ideological skew for the disciplines of economics, political science, psychology, and sociology in Savolainen (2025). Most abstracts were rated between 6 and 8 in every discipline. Notably, even in the farthest left disciplines only a very small fraction was rated 9–10. The consistently left-of-center disciplinary averages were driven less by the prevalence of far-left viewpoints than they were by the near absence of right-of-center works in most disciplines. In economics, the least left-leaning discipline, 16% of works were rated 0–4; in all other disciplines, between 0 and 6% of works were rated in this range.

The mean overall political stance score was calculated by discipline for all Politically Relevant Abstracts for each year 1960–2024 in which at least 50 abstracts were available. In this analysis, disciplines were divided into two categories: policy-proximal and policy-distal. Disciplines for which methods, training, and institutional roles are routinely oriented toward the design, evaluation, or delivery of public programs were categorized as policy-proximal. Disciplines for which scholarly practice is primarily explanatory, interpretive, or critical in orientation were categorized as policy-distal. Policy-proximal disciplines were criminology, economics, political science, public administration, and public health. Policy-distal disciplines were anthropology, communications, ethnic studies, gender studies, psychology, and sociology.

Example trend analyses of a policy-proximal (economics) and a policy-distal (sociology) discipline are shown in Fig. 2. To visualize temporal trends in political

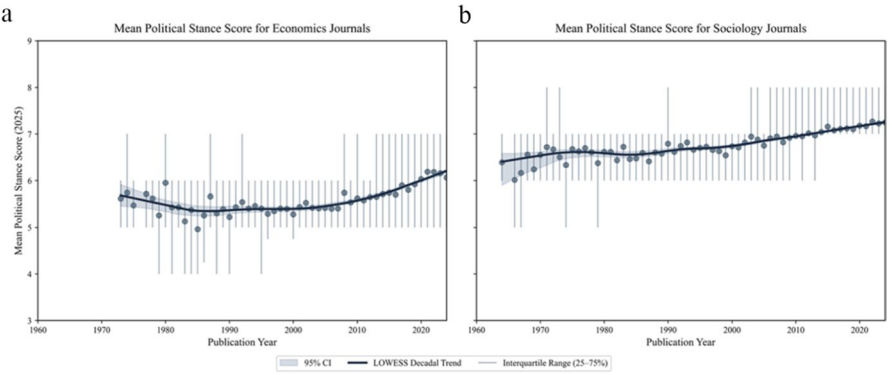

Fig. 2 **a** Mean political stance score for economics journals **b** Mean political stance score for sociology journals

stance, a 10-year moving average was estimated using the LOWESS[5] algorithm with a confidence band.[6] Annual Interquartile Range is also shown.

Both disciplines remained left-of-center every year. Economics demonstrated a 'U' pattern of rightward then leftward movement. Sociology was further left than economics every year and displayed generally leftward movement throughout the period.

These overall patterns tended to be true for policy-proximal and policy-distal disciplines in general. The results for all eleven disciplines are shown in Fig. 3a and b, where each line represents the LOWESS moving average of the mean political stance score for abstracts in that discipline. Separate plots are shown for policy-proximal and policy-distal disciplines.

Every discipline trendline was left-of-center every year 1960–2024 and every discipline trended leftward 1990–2024.[7] The policy-proximal disciplines of economics, political science, and public administration demonstrated rightward movement through sometime in the 1980s before reversing course and beginning leftward movement by about 1990. The policy-distal disciplines generally showed leftward movement throughout the period.

Visual inspection appeared to show that several disciplines displayed accelerated leftward movement beginning around 2010. Formal analysis[8] indicated that the most left-leaning disciplines – gender studies, anthropology, and ethnic studies – all had acceleration breakpoints between 2011 and 2014 ($p < 0.01$). In addition, psychology

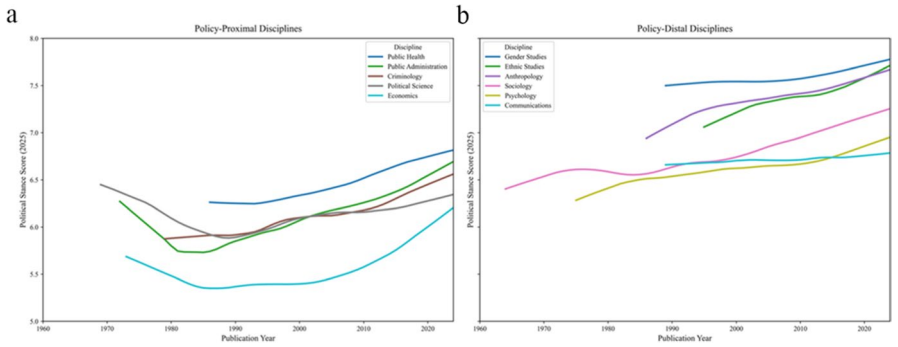

Fig. 3 **a** Policy-proximal disciplines **b** Policy-distal disciplines

[5] The LOWESS (Locally Weighted Scatterplot Smoothing) algorithm fits simple regressions to localized subsets of the data to produce a smooth curve that captures underlying patterns while minimizing noise (Cleveland, 1979).

[6] Confidence interval calculation methodology is specified in Supplementary Information Exhibit S4.

[7] Simple linear regressions of annual mean scores on publication year were fit for each discipline 1990—2024, retaining only discipline–year observations with a minimum of fifty annual papers to ensure stability of estimates. Every discipline demonstrated a statistically significant leftward trend ($p < 0.05$) from 1990–2024.

[8] The Quandt–Andrews sup-Wald test, which searches over unknown break dates and evaluates the maximum Wald statistic across candidate years (Andrews, 1993), was used to obtain p-values for an acceleration breakpoint 2005–2015 using Hansen's (1997) asymptotic approximations implemented via bootstrap.

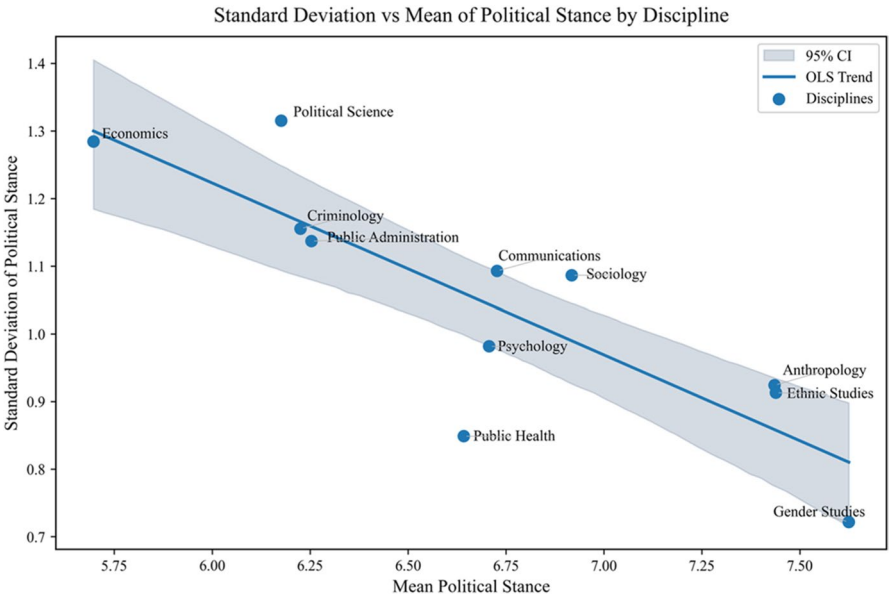

Fig. 4 Standard deviation vs mean of political stance by discipline

demonstrated an acceleration breakpoint in 2010 ($p=0.01$) as did economics in 2006 ($p<0.01$).

Next, ideological heterogeneity was considered by discipline. Figure 4 shows the overall mean and standard deviation of political stance scores for each discipline for all years 1960–2024, along with an ordinary least squares line and confidence band.

Disciplinary mean and standard deviation demonstrated a strongly negative correlation ($r=-0.84$; $p<0.01$). That is, the disciplines that had the most leftward political orientation also tended to be those that were most politically homogenous.

Political stance scores were grouped by discipline by decade, and sample standard deviations were calculated for all discipline–decade combinations with at least 100 abstracts. These standard deviations are displayed over time separately for policy-proximal and policy-distal disciplines in Fig. 5a and b.

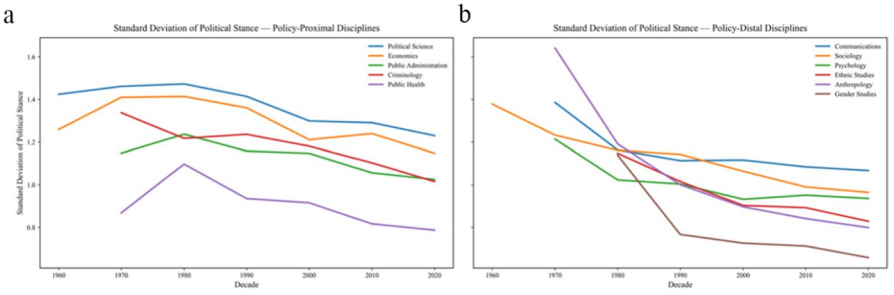

Fig. 5 **a** Standard deviation of political stance-policy-proximal discipline **b** Standard deviation of political stance-policy-distal discipline

All policy-proximal disciplines except criminology displayed rising heterogeneity of scores through the 1980 s, then declining heterogeneity through the 2020s. This pattern mirrors the rightward shift in these disciplines through roughly the 1980 s, followed by leftward movement since. The policy-distal disciplines and criminology generally showed movement to lower heterogeneity over the entire period. This also mirrors the movement leftward for these disciplines over the same period. Greater leftward orientation was associated with greater political stance homogeneity, whether considered cross-sectionally across disciplines or temporally within disciplines.

Finally, discipline-level trends in political stance scores were combined to create aggregate social science trends for overall, economic, and sociocultural political stance scores. This was accomplished by averaging discipline scores in each year weighted by the estimated fraction of all relevant articles published in each discipline across the entire OpenAlex database. The substantial majority of relevant articles were found to be in four disciplines: psychology, political science, economics, and sociology. Detailed weighting calculations are described in Supplementary Information (Exhibit S4). The weight-averaged total social science trend is shown in Fig. 6.

Since every discipline was left-of-center every year, their weighted-average was of course also left-of-center every year. The weight-averaged political stance score demonstrated leftward movement in the 1960 s, followed by stability through about 1990, followed by leftward movement through the 2020s. The same statistical tests applied to discipline-level trends identified an acceleration of leftward movement in 2010 ($p < 0.01$).

To estimate sources of temporal change in political stance scores, a three-way Mundlak decomposition was implemented (Mundlak, 1978; Bell & Jones, 2015), which partitions variation in a time-varying outcome into between-group, between-

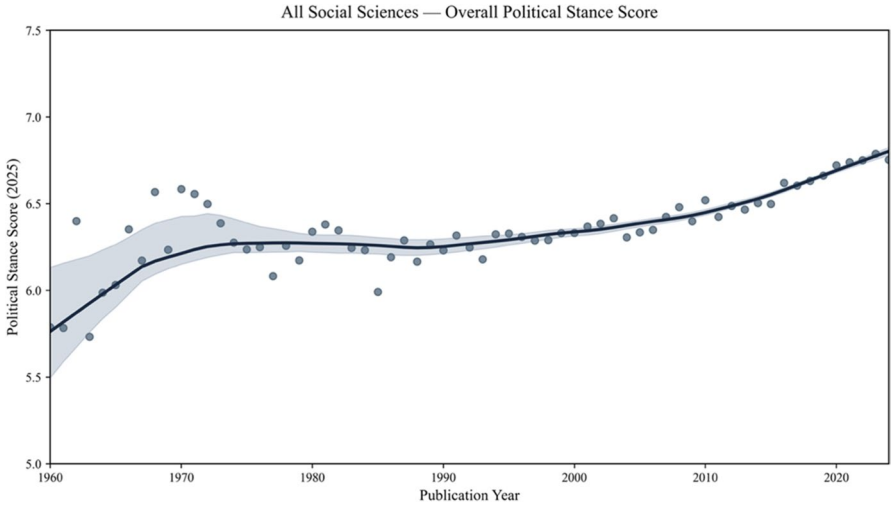

Fig. 6 All social sciences-overall political stance score

unit, and within-unit components.[9] The decomposition indicated that approximately 51–57% of the observed leftward trajectory is attributable to differences in disciplinary composition, 27–29% to variation between authors within disciplines, and 16–20% to changes within authors over the course of their careers. In practical terms, this means that positional change in stance scores is driven primarily by the kinds of abstracts entering the publication stream—new contributors and subfields whose work is classified farther to the left under the fixed 2025 rubric—while within-author shifts exist but play a materially smaller role under the same evaluative frame.

Figure 7a and b show mean annual scores and trends for political stance scores separately on economic issues and sociocultural issues (Q4 and Q5).

As with overall scores, the mean abstract was left-of-center every year 1960–2024 whether rated on economic or sociocultural issues. The differences, however, were striking. First, abstracts were consistently rated as further left on sociocultural issues than economic issues. Second, this disparity grew over time ($p<0.01$). In the early years of the period, sociocultural scores were generally about 10–15 percent higher than economic scores, but by the 2020 s, this difference had increased to about 25–30 percent. On economic issues, the LOWESS trend followed the same kind of 'U' pattern observed for several policy-proximal disciplines and by the 2020 s was approximately back to its prior late-sixties peak, while sociocultural scores were unprecedentedly left.

In summary, five key findings were observed in analysis of the Politically Relevant Abstracts. First, roughly 90 percent of politically relevant social science articles leaned left 1960–2024, and the mean political stance of every social science discipline was left-of-center every year during the period. Second, all disciplines showed leftward movement between 1990 and 2024. Third, policy-proximal disciplines generally showed limited rightward moderation between roughly 1970 and 1990, though

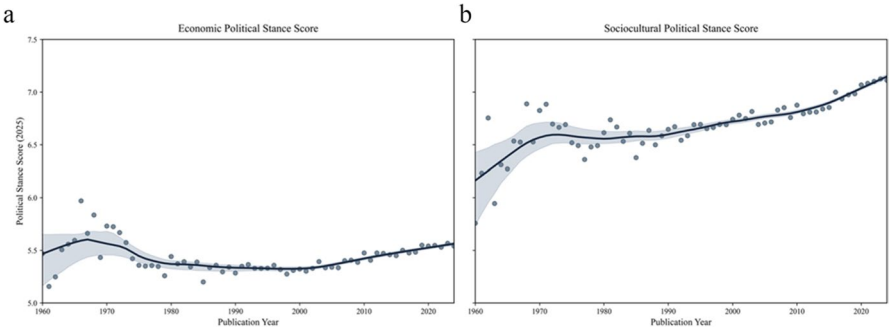

Fig. 7 **a** Economic political stance score **b** Sociocultural political stance score

[9] The standard two-level Mundlak formulation was extended to account for three hierarchical levels: (i) between disciplines, using each discipline's mean publication year values; (ii) between authors within disciplines, using each author's deviation from the discipline mean year value; and (iii) within authors, using the deviation of each article's year from the author's mean value. The analysis was weighted using the discipline–year publication weights described previously. Thresholds of 2, 3, and 10 total publications per author were tested.

policy-distal disciplines did not. Fourth, disciplines with greater leftward orientation generally displayed greater ideological homogeneity Fifth, sociocultural content was more consistently left-leaning than economic content, and that gap widened over time.

Robustness analyses

A wide variety of robustness checks was executed to test the generalizability of these findings. One set of robustness analyses tested variations to analytical methods. A second applied the base analytical methods to alternative datasets built using different inclusion criteria or data sources. A final analysis tested the effect of analyzing all abstracts, including those that the LLM considered politically irrelevant.

Analytical variations

Five analytical variant types were tested. Each defined an analytical pipeline that modified the base pipeline used to generate the findings in Results, applied both the base pipeline and this modified pipeline to a random subset of all abstracts, and then compared the results produced by each approach. Two types of metrics were developed for each comparison. Abstract-level alignment was evaluated using both Pearson correlation and ICC. Weighted values for mean overall political stance score and percent of abstracts rated as 'left' were calculated for both the base and modified samples to provide a measure of whether the modified pipeline would have substantially changed estimated overall ideological orientation.

Each of the five analytical variations is described below.

Alternative LLM models

1,000 abstracts were randomly selected from Analyzed Abstracts, and materially the same prompts and system instructions described in Data and Methods were presented for each of these abstracts to four widely-used LLMs created by multiple organizations: GPT-4o (baseline), GPT-4o-mini, mixtral-8 × 7b-instruct ('Mixtral'), and llama-v3p3-70b-instruct ('Llama').

Prompt variations

1,000 abstracts were randomly selected from Analyzed Abstracts, and six variations to the Prompting Method were applied: (i) automatically generated word variation using synonyms; (ii) randomly re-ordered questions for each prompt; (iii) anchor points added directly into prompt text as examples; (iv) liberal/conservative used instead of left/right; (v) 'chain-of-thought' prompt instruction added to request explicit description of LLM reasoning; and (vi) LLM instructed to rely strictly on the text in the abstract and ignore all other context.

Alternative LLM personas

1,000 abstracts were randomly selected from Analyzed Abstracts, and ten persona-conditioned prompting variants were constructed as five left–right paired personas, each pair representing parallel worldviews (e.g., progressive vs. conservative policy advisors; urban millennial Democrat vs. rural veteran Republican). These were injected into the LLM's system instructions. Several of these profiles represent the most ideologically extreme counterpositions relative to the model's known priors, thereby serving as explicitly adversarial prompts in the sense described by Bernardelle et al. (2025).

Anchor variations

1,000 abstracts were randomly selected from Analyzed Abstracts, and two approaches were used to test variations to anchors points defined in S3. First 'anchor regeneration' prompted the LLM to generate ten entirely new sets of 11 anchor points (positions 0–10). Second, 'anchor perturbation' created ten modified anchor sets by randomly shifting five of the original anchor positions by ± 1 in each modified anchor set. Together, these methods produced a total of 20 new anchor sets.

Contemporaneous political relevance and orientation measurement

The rationale for using a fixed 2025 rating scale for abstracts was described earlier. That said, one would expect that evaluations of political stance against the standards of the year of publication should not be completely unrelated to this fixed scale, and further, it would be troubling if using such a contemporaneous rating method produced a wholly different distribution of political stance scores. A random sample of 10,000 works was chosen from all Analyzed Abstracts and the LLM was used to obtain contemporaneous political relevance and stance scores for these abstracts. The Prompting Method was modified to add the year of publication to each prompt and request estimates of political relevance and political stance scores in that year for each abstract without reference to the fixed 2025 anchor points.

The results for all five analytical variations are shown in Table 3.

All tested analytical variations confirm that politically relevant academic social science abstracts have a mean left-of-center political stance score and that approximately 90% of abstracts are rated as 'left'. Inter-rater reliability metrics were all good

Table 3 Analytical variation robustness results

	Inter-Rater Reliability		Mean Political Stance Score		Percent Rated as 'Left'	
Analytical Variation	Pearson Correlation	ICC	Baseline	Variant	Baseline	Variant
Alternative LLMs	.77—.81	.76—.79	6.9	6.2–6.6	95	96–100
Prompt Text Variations	.93—.97	.93—.97	6.8	6.5–6.9	92	92–94
Alternative LLM Personas	.82—.90	.82—.90	6.3	6.0–6.7	96	86–96
Anchor Regeneration	.93—.94	.92—.94	6.3	6.1–6.2	96	92–96
Anchor Perturbation	.80—.93	.76—.93	6.3	6.1–6.7	96	87–96
Contemporaneous Scoring	.78	.78	6.4	6.7	98	92

to excellent. Those analytical variations that we should expect to align closely with the base case – small variations to prompt text and anchor regeneration – had ICC scores above 0.9. Other variations that we would expect to show some difference to base case because of materially changed prompting – anchor perturbation, alternative LLM personas, and contemporaneous scoring – showed ICC scores between 0.76 and 0.93. Alternative LLM models showed good agreement, and all LLMs showed mean political stance greater than 6 and left-of-center orientation for more than 90% of abstracts.

Dataset variations

Five dataset variant types were tested. Each applied the base analytical pipeline to these dataset variants and then compared the results to a relevant benchmark for results on the base dataset using the same metrics for overall ideological orientation that were used to evaluate analytical variations.

Each of the five dataset variations is described below.

Full papers

A random sample of DOIs for Analyzed Abstracts was submitted to the Unpaywall API (Unpaywall: Open database of scholarly articles., n.d.) and the first 2,459 works that were available as open access downloadable as PDFs were selected as a sample of full papers.

U.S. Authored papers

669 US universities and colleges named in the US News & World Report ('USNWR', Datasets | Andrew G. Reiter, n.d.) listings that could be matched with institutional affiliation names in the OpenAlex database were identified. 75,829 (approximately 42%) of the 181,311 politically relevant abstracts had at least one author affiliated with one of these institutions and were selected as a sample.

Alternative OpenAlex inclusion criteria

Three alternative approaches to selecting works from within the OpenAlex database were tested. First was 'ablation' in which a bootstrapping process was executed ten times to eliminate a randomly-selected half of the 367 journals used in the main analysis and construct a sample by applying the same methods used in the base analysis on this reduced set of journals. Second was 'journal-level concept-based selection'. All OpenAlex core journals with at least 1,000 cumulative works published in English 1960–2024 for which any of the eleven targeted disciplines had an association score of at least 0.7 was identified as a relevant journal for the discipline with the highest association score. 10,000 articles were randomly selected from the works in these journals. Third was 'article-level concept-based selection' in which a highly permissive sample of 10,000 works was randomly selected from all OpenAlex English-language articles with abstracts published 1960–2024 with any positive concept association score with any of the eleven target disciplines.

OpenAlex articles without deposited abstracts

In addition to the Analyzed Abstracts, 467,146 lacked abstracts in OpenAlex but satisfied all other inclusion criteria. A small fraction of these articles has abstracts deposited in the Semantic Scholar database (Semantic scholar | ai-powered research tool, n.d.). A random sample of 1,375 such works was collected from the Semantic Scholar API.

Alternative databases of academic works

Due to the limitations of the web interfaces provided by SCOPUS and WoS, systematic methods were implemented to select samples of sociology articles from both databases as test cases. All English-language articles with the subject of sociology published between 1960 and 2024 in any of the sociology journals used to create Analyzed Abstracts were extracted from both SCOPUS and WoS. The extracted SCOPUS works were sorted alphabetically by first author name, and the first 5,000 works were downloaded. 4,137 of these were found to have abstracts deposited in SCOPUS. 40 of the 406 pages of works returned for the WoS query were randomly selected, the 2,000 works referenced on these 40 pages were downloaded, and 1,073 of these works were found to have abstracts deposited in WoS. Results for these samples were compared to base results for sociology abstracts.

The results for all five dataset variations are shown in Table 4.

All tested dataset variations confirmed that politically relevant academic social science abstracts would have mean left-of-center political stance scores quite similar to those in the base dataset and that approximately 90% of the abstracts in these alternative datasets would be rated as ‘left’. A comparison of papers with and without US authors is provided as Supplementary Information Exhibit S5.

All OpenAlex analyzed abstracts

The Politically Relevant Abstracts that comprised the core set of works analyzed for political stance represented approximately 30% of all Analyzable Abstracts. This was a conscious analytical decision, but it would be troubling if including the roughly 70% of abstracts deemed politically irrelevant would have radically changed the key findings. 91% of the Analyzed Abstracts rated as irrelevant were rated as 5, or neutral, on a 0–10 political stance scale carefully fixed to 2025 U.S. political debates. That is, when forced to rate the political stance of an abstract judged to be politically irrelevant, the LLM made the sensible decision to assign a neutral rating in the vast majority of cases. Of course, this created a strong attraction to a mean score of 5 when evaluating all Analyzed Abstracts. Recall that only about 5% of abstracts published in 1960 were considered relevant, rising to about 50% in 2024; therefore, this attraction to a mean of 5 was even stronger in the earlier years of the period. The mean overall political stance score by discipline for all Analyzed Abstracts in shown in Fig. 8a and b.

Table 4 Dataset variation robustness results

		Mean Political Stance Score		Percent Rated as "Left"	
Dataset Variation	Sample Size	Baseline	Variant	Baseline	Vari-ant
Full Papers	2,4S9	6.7	6.9	93	99
US-Only	7S,829	6.7	6.6	93	92
Alternative OpenAlex Inclusion Criteria *Ablntion*	74,868–101,824	6.7	6.6–6.8	93	92–94
Journal-level Concept-Based Selection	10,000	6.7	6.7	93	93
Article-level Concept-Based Selection	10,000	6.7	6.6	93	89
OpenAiex Articles without Abstracts	1,373	6.7	6.8	93	94
Alternative databases *SCOPUS*	4,173	6.9	7.0	96	96
WoS	1,073	6.9	6.9	96	94

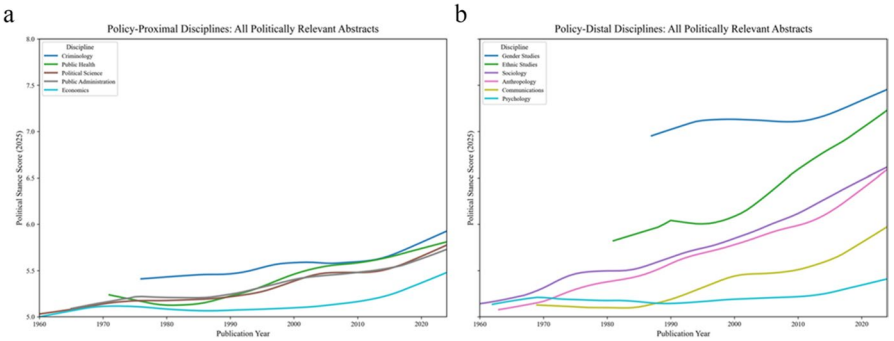

Fig. 8 **a** Policy-proximal disciplines: all politically relevant abstracts **b** Policy-distal disciplines: all politically relevant abstracts

When evaluating all Analyzed Abstracts, every discipline trendline remained left-of-center every year after 1960 and all disciplines demonstrated a statistically significant leftward trend 1990–2024, though the mean in all cases was shifted toward 5, especially in the earlier years of the period.

In summary, key findings appeared robust to a wide range of variations in analytical assumptions and dataset construction approaches.

Discussion

This study provides the first long-run (1960–2024), large-scale, cross-disciplinary analysis of the ideological orientation of anglophone academic social science research using LLM-based text classification. By applying a fixed 2025 U.S. ideological frame to historical abstracts, it produces a consistent time series that reveals how the work of past decades reads to a contemporary observer. This approach enables comparisons across six decades that would be impossible if coding relied on contemporaneous, shifting political categories.

Prior work has shown that U.S. university social science faculty have leaned left for decades. This study extends those findings to published research output: when evaluated under a stable 2025 ideological reference frame, anglophone social science abstracts have been persistently classified to the left of center for decades, and this asymmetry has intensified over the past thirty years, most sharply on sociocultural dimensions. These results should not be interpreted as dispositive evidence of ideological bias. They do not, in themselves, imply partisan distortion or suppression of alternative perspectives. One possibility is that sustained, disinterested inquiry into social phenomena has arrived at conclusions that happen to align more closely with liberal than conservative viewpoints. Another is that certain research topics have become more salient over time, and these topics are inherently more likely to be framed in ways that resonate with progressive values. While the data clearly show a political orientation in published output, they do not reveal the mechanisms by which that orientation is produced.

These results can, however, be situated within a sociological context, as illuminated by their dynamics along two dimensions: the evolution of overall ideological orientation over time and the large and growing division between the ideological orientation on economic versus sociocultural issues.

Following Bourdieu's *Homo Academicus* (Bourdieu, 1988), universities are at once embedded in broader political fields and partially insulated through their own internal logics. Field-environment interaction theories (Fligstein & McAdam, 2012; Meyer & Rowan, 1977) predict that ideological orientation in research will reflect exogenous political shifts during periods of high permeability. By contrast, Gross's (2013) occupational self-selection model suggests that as a field's political reputation becomes well known, ideological composition increasingly reproduces itself through recruitment, reducing responsiveness to short-term political changes. As a general statement about the entire period, the empirical observation that leftward movement has been dominated by new, more left-leaning academics entering the field rather than by academics moving leftward during their careers provides strong support for the occupational self-selection model. It is plausible, however, that if this self-selection has happened over more than six decades, then we would expect academic social science to have evolved from greater field-environment interaction to greater self-selection. That is, earlier decades should show greater responsiveness to political context, whereas later decades should display more stable, internally reinforced patterns.

This is precisely what we see in the data. The aggregate pattern—sharp leftward movement in the 1960 s, moderation in the 1970s–1980s, leftward drift from 1990

onward, and acceleration after 2010—fits a two-phase interpretation. The 1960 s and 1980 s shifts are consistent with higher field-environment permeability: the leftward swing aligns with civil rights, anti-war, and feminist mobilizations, while the subsequent moderation, more pronounced on economic issues, parallels the rise of market-oriented political regimes. The post-1990 period, by contrast, shows reduced volatility and sustained leftward drift, consistent with self-selection dynamics and a more insulated intellectual field. The 2010 acceleration, concentrated in certain disciplines, supports claims of ideological intensification in the last decade but appears as an extension of a long-running trajectory rather than a rupture.

The discipline-level trends are consistent with this view. Most policy-proximal disciplines generally exhibited a "U"-shaped pattern of leftward movement in the 1960 s, rightward moderation in the 1970s–1980s, and renewed leftward drift thereafter, suggesting greater permeability to the surrounding political climate in the earlier decades with growing left-oriented homogeneity post 1990. In contrast, policy-distal disciplines whose boundaries with external political institutions are more mediated by academic subcultures displayed steadier leftward trajectories, consistent with higher levels of internal insulation and self-reinforcement throughout the period.

But why has this leftward movement been so disproportionately intense for sociocultural versus economic issues? "Socially liberal and economically conservative" has been a well-known self-description among prosperous and highly educated Americans for decades (Page et al., 2013). This tendency has been described by Fraser (2019) as "progressive neoliberalism", and it extends well beyond academia – 'an alliance of […] mainstream liberal currents of the new social movements (feminism, antiracism, multiculturalism, environmentalism, and LGBTQ+rights)' with 'the most dynamic, high-end, "symbolic," and financial sectors of the US economy (Wall Street, Silicon Valley, and Hollywood)' (Fraser, 2019, p. 11). Social scientists moving substantially further left on sociocultural versus economic issues has allowed them to remain members in good standing of a coalition that includes some of the wealthiest people in an increasingly economically unequal society. For academics, the structure of this coalition is reflected precisely in the combination of moderately left economics and aggressive sociocultural progressivism shown in Figs. 7a and b.

An important limitation of this analysis is that it does not investigate the causes of the observed ideological patterns. Topic salience shifts, such as the rise of climate change, racial justice, and gender/sexuality research, may have increased the proportion of left-coded content. Editorial and funding priorities could have shifted toward projects with certain ideological framings. Growth of subfields with inherently left-coded topics, especially in the last three decades, could also play a role.

Empirically, the analysis supports a transition from a field responsive to external political cycles toward one in which ideological orientation is reproduced internally. The methodological pipeline demonstrates that LLMs can deliver consistent large-scale ideological coding across decades, enabling new lines of research in the sociology of knowledge. Future work could extend this approach cross-nationally, to other languages, or to other academic domains such as the humanities, STEM, or legal scholarship. Incorporating qualitative coding of sample texts could also illuminate the substantive content behind ideological scores, deepening the link between measurement and theory.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s11186-026-09690-2.

Funding Not applicable.

Data availability The bibliometric and text data analyzed in this study are derived from the public OpenAlex dataset (snapshot dated February 27 2025), which is freely available under the CC BY license. All processing scripts and model prompts used for data extraction, discipline classification, and LLM-based text coding are available from the author upon reasonable request. No proprietary or confidential data were used.

Declarations

Ethical approval Not applicable.

Conflict of interest Author declares that he has no conflict of interest.

References

Alonso-Álvarez, P., & van Eck, N. J. (2025). Coverage and metadata completeness and accuracy of African research publications in OpenAlex: A comparative analysis. *ArXiv preprint*. https://doi.org/10.48550/arXiv.2409.01120

Andrews, D. W. K. (1993). Tests for parameter instability and structural change with unknown change point. *Econometrica, 61*(4), 821.

Bar, T., & Zussman, A. (2012). Partisan grading. *American Economic Journal, Applied Economics, 4*(1), 30–48. https://doi.org/10.1257/app.4.1.30

Bernardelle, P., Fröhling, L., Civelli, S., Lunardi, R., Roitero, K., & Demartini, G. (2025). Mapping and influencing the political ideology of large language models using synthetic personas. *ArXiv* Article 2412.14843. https://doi.org/10.48550/arXiv.2412.14843

Bonica, A., Chilton, A., Rozema, K., & Sen, M. (2018). The legal academy's ideological uniformity. *Journal of Legal Studies, 47*(1), 1–43. https://doi.org/10.1086/698435

Borjas, G. J., & Breznau, N. (2024). Ideological bias in estimates of the impact of immigration. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.5069989

Bourdieu, P. (1988). *Homo academicus*. Polity Press.

Boysen, G. A., Vogel, D. L., Cope, M. A., & Hubbard, A. (2009). Incidents of bias in college classrooms: Instructor and student perceptions. *Journal of Diversity in Higher Education, 2*(4), 219–231. https://doi.org/10.1037/a0017538

Burmila, E. (2021). Liberal bias in the college classroom: A review of the evidence (or lack thereof). *PS Political Science & Politics, 54*(3), 598–602. https://doi.org/10.1017/s1049096521000354

Cardiff, C. F., & Klein, D. B. (2005). Faculty partisan affiliations in all disciplines: A voter-registration study. *Critical Review, 17*(3–4), 237–255. https://doi.org/10.1080/08913810508443639

Chilton, A. S., & Posner, E. A. (2015). An empirical study of political bias in legal scholarship. *Journal of Legal Studies, 44*(2), 277–314. https://doi.org/10.1086/684302

Chin, M. J., Ortega, A., Shaw, M. P., & Yoo, D. (2025). *Politics of the professoriate: Longitudinal evidence from a state public university system's universe of faculty* (EdWorkingPaper No. 25–1235). Annenberg Institute, Brown University. Retrieved July 2025, from https://edworkingpapers.com/sites/default/files/ai25-1235.pdf

Clarivate. (n.d.). Web of Science. Retrieved July 2025, from https://www.webofscience.com

Cleveland, W. S. (1979). Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association, 74*(368), 829–836. https://doi.org/10.1080/01621459.1979.10481038

Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educational and Psychological Measurement, 20*(1), 37–46. https://doi.org/10.1177/001316446002000104

Culbert, J. H., Hobert, A., Jahn, N., Haupka, N., Schmidt, M., Donner, P., & Mayr, P. (2025). Reference coverage analysis of OpenAlex compared to Web of Science and Scopus. *Scientometrics*. https://doi.org/10.1007/s11192-025-05293-3

Datasets | Andrew G. Reiter. (n.d.). Retrieved July 2025, from https://andyreiter.com/datasets/

Ellen Bara Stolzenberg, E. B., Eagan, K., Zimmerman, H. B., Lozano, J. B., Cesar-Davis, N. M., Aragon, M. C., & Rios-Aguilar, C. (2019). *Undergraduate teaching faculty: The HERI faculty survey 2016–2017*. Higher Education Research Institute Graduate School of Education & Information Studies University of California.

Fligstein, N., & McAdam, D. (2012). *A theory of fields*. Oxford University Press.

Fraser, N. (2019). *The old is dying and the new cannot be born: From progressive neoliberalism to trump and beyond*. Verso.

Grimmer, J., & Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis, 21*(3), 267–297. https://doi.org/10.1093/pan/mps028

Gross, N. (2013). *Why are professors liberal and why do conservatives care?* Harvard University Press.

Gross, N., & Simmons, S. (2014a). Political bias in the graduate school admissions process. In N. Gross & S. Simmons (Eds.), *Professors and their politics*. Johns Hopkins University Press.

Gross, N., & Simmons, S. (2014b). The social and political views of american college and university professors. In N. Gross & S. Simmons (Eds.), *Professors and their politics*. Johns Hopkins University Press.

Hamilton, R. F., & Hargens, L. L. (1993). The politics of the professors: Self-identifications, 1969–1984. *Social Forces, 71*(3), 603. https://doi.org/10.2307/2579887

Hansen, B. E. (1997). Approximate asymptotic P values for structural-change tests. *Journal of Business & Economic Statistics, 15*(1), 60. https://doi.org/10.2307/1392074

Havey, N., & Her, X. (2025). The politics of the professoriate: A social media approach. *The Review of Higher Education, 48*(4), 583–610. https://doi.org/10.1353/rhe.2025.a962245

Honeycutt, N. (2024). The politics of university faculty [preprint]. *PsyArXiv*. https://doi.org/10.31234/osf.io/dnxqh

Internet archive: Digital library of free & borrowable texts, movies, music & wayback machine. (n.d.). Retrieved July 2025, from https://archive.org/

Jelveh, Z., Kogut, B., & Naidu, S. (2024). Political language in economics. *The Economic Journal*. https://doi.org/10.1093/ej/ueae026

Journal citation reports . (n.d.). Clarivate. Retrieved July 2025, https://clarivate.com/academia-government/scientific-and-academic-research/research-funding-analytics/journal-citation-reports/

Kaurov, A. A., Cologna, V., Tyson, C., & Oreskes, N. (2022). Trends in American scientists' political donations and implications for trust in science. *Humanities and Social Sciences Communications*. https://doi.org/10.1057/s41599-022-01382-3

Klein, D. B., & Stern, C. (2005a). How politically diverse are the social sciences and humanities? Survey evidence from six fields. *Academic Questions*. Retrieved July 2025, from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=664042

Klein, D. B., & Stern, C. (2005b). Professors and their politics: The policy views of social scientists. *Critical Review, 17*(3–4), 257–303. https://doi.org/10.1080/08913810508443640

Knudsen, B., Madkour, A., Cholli, P., Haslam, A., & Prasad, V. (2024). Analysis of the political viewpoint of policy statements from professional medical organizations using chatgpt 4.0 (preprint). *JMIR Formative Research*. https://doi.org/10.2196/66204

Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics, 33*(1), 159. https://doi.org/10.2307/2529310

Langbert, M. (2020). Republicans need not apply: An investigation of the American economic association using voter registration and political contributions. *Econ Journal Watch, 17*(2), 392.

Langbert, M., Quain, A. J., & Klein, D. B. (2016). Faculty voter registration in economics, history, journalism, law, and psychology. *Econ Journal Watch, 13*(3), 422–451.

Langbert, M., & Stevens, S. (2021). Partisan registration of faculty in flagship colleges. *Studies in Higher Education*. https://doi.org/10.1080/03075079.2021.1957815
Lazarsfeld, P. F., & Thielens, W. (1958). *The academic mind: Social scientists in a time of crisis*. Free Press.
Le Mens, G., & Gallego, A. (2025). Positioning political texts with large language models by asking and averaging. *Political Analysis*. https://doi.org/10.1017/pan.2024.29
Lipset, S. M. (1976). *The divided academy: Professors and politics* (Carnegie Commission on Higher Education., Ed.). Norton.
Mariani, M. D., & Hewitt, G. J. (2008). Indoctrination U.? Faculty ideology and changes in student political orientation. *PS Political Science & Politics, 41*, 773–783. https://doi.org/10.1017/s1049096508081031
McGinnis, J. 0, Schwartz, M. A., & Tisdel, B. (2005). The patterns and implications of political contributions by Elite Law School Faculty. *Georgetown Law Journal, 931167*, Article 1167.
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
Mundlak, Y. (1978). On the pooling of time series and cross section data. *Econometrica, 46*(1), 69. https://doi.org/10.2307/1913646
Musgrave, P., & Rom, M. (2014). Fair and balanced? Experimental evidence on partisan bias in grading. *American Politics Research, 43*(3), 536–554. https://doi.org/10.1177/1532673x14561655
O'Hagan, S., & Schein, A. (2023). Measurement in the age of llms: An application to ideological scaling. *ArXiv preprint*, arXiv:2312.09203.
Page, B. I., Bartels, L. M., & Seawright, J. (2013). Democracy and the policy preferences of wealthy Americans. *Perspectives on Politics, 11*(1), 51–73. https://doi.org/10.1017/s153759271200360x
Priem, J., Piwowar, H., & Orr, R. (2022). OpenAlex: A fully-open index of scholarly works, authors, venues, institutions, and concepts. *ArXiv*
Savolainen, J. (2025). The methodological stagnation of sociology is related to its left-wing skew. *Theoretical Sociology*. https://doi.org/10.1007/s11186-025-09640-4
Schiekiera, L., & Niemeyer, H. (2025). Political bias in historiography - An experimental investigation of preferences for publication as a function of political orientation. *F1000Research, 14*, Article 320. https://doi.org/10.12688/f1000research.160170.1
Scimago journal & country rank. (n.d.). Scimago Journal & Country Rank. Retrieved July 2025, from https://www.scimagojr.com/
Scopus. (n.d.). Elsevier. Retrieved July 2025, from https://www.scopus.com
Semantic scholar | ai-powered research tool. (n.d.). Allen Institute for AI. Retrieved July 2025, from https://www.semanticscholar.org/
Shrout, P. E., & Fleiss, J. L. (1979). Intraclass correlations: Uses in assessing rater reliability. *Psychological Bulletin, 86*(2), 420–428. https://doi.org/10.1037/0033-2909.86.2.420
Slapin, J. B., & Proksch, S.-O. (2008). A scaling model for estimating time-series party positions from texts. *American Journal of Political Science, 52*(3), 705–722. https://doi.org/10.1111/j.1540-5907.2008.00338.x
Unpaywall: Open database of scholarly articles. (n.d.). OurResearch. Retrieved July 2025, from https://unpaywall.org/
Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., … van Mulbregt, P. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods, 17*(3), 261–272. https://doi.org/10.1038/s41592-019-0686-2
Woessner, M., & Kelly-Woessner, A. (2009). I think my professor is a democrat: Considering whether students recognize and react to faculty politics. *PS: Political Science & Politics, 42*(02), 343–352. https://doi.org/10.1017/s1049096509090453
Zipp, J. F., & Fenwick, R. (2006). Is the academy a liberal hegemony?: The political orientations and educational values of professors. *Public Opinion Quarterly, 70*(3), 304–326. https://doi.org/10.1093/poq/nfj009
Bell, A., & Jones, K. (2015). Explaining fixed effects: Random effects modeling of time-series cross-sectional and panel data. Political Science Research and Methods, 3(1), 133–153 https://doi.org/10.1017/psrm.2014.7

James Manzi is a doctoral student in Sociology at the University of Oxford.