THE
ECONOMIC
JOURNAL

The Economic Journal, **134** (*August*), 2439–2469 https://doi.org/10.1093/ej/ueae026
Advance Access Publication Date: 29 April 2024

POLITICAL LANGUAGE IN ECONOMICS*

Zubin Jelveh, Bruce Kogut and Suresh Naidu

Does academic writing in economics reflect the political orientation of economists? We use machine learning to measure partisanship in academic economics articles. We predict the observed political behaviour of a subset of economists using phrases from their academic articles, show good out-of-sample predictive accuracy and then predict partisanship for all economists. We then use these predictions to examine patterns of political language in economics. We estimate journal-specific effects on predicted ideology, controlling for author and year fixed effects, that accord with existing survey-based measures. We show considerable sorting of economists into fields of research by predicted partisanship. We also show that partisanship is detectable even within fields, even across those estimating the same theoretical parameter. Using policy-relevant parameters collected from previous meta-analyses, we then show that imputed partisanship is correlated with estimated parameters, such that the implied policy prescription is consistent with partisan leaning. For example, we find that going from the most left-wing authored estimate of the taxable top income elasticity to the most right-wing authored estimate decreases the optimal tax rate from 77% to 60%.

Modern governments incorporate academic economists' research findings into policy analysis via a wide variety of formal and informal mechanisms. For example, economists inform central bank policy, antitrust policy and the design of taxes and regulation. The policy relevance of economics partially stems from its ability to combine economic theory (e.g., supply and demand) with parameter estimates (e.g., elasticities) to make prescriptions about optimal policies (e.g., taxes). Among social scientists, economists have a great deal of weight as government officials and public commentators. Their academic opinions and judgements are often expected to be non-partisan, but these experts may have partisan or political preferences of their own. This leads naturally to the question of how partisan is academic economics. Do the methodological conventions of academic economics, such as formal modelling, quantitative analysis and peer review successfully filter partisanship from academic economics research? We answer this question by applying tools from natural language processing (Gentzkow and Shapiro, 2010) to a comprehensive corpus of academic economics articles. We link academic economist political behaviour, measured from campaign contributions and political petition signing, with the plain text of academic articles. We then train a machine learning algorithm to predict political behaviour of authors within this linked sample, both unconditionally and within detailed fields of research. We show that our

* Corresponding author: Suresh Naidu, Department of Economics and SIPA, Columbia University, 420 w 118th Street, New York, NY 10025, USA. Email: sn2430@columbia.edu

This paper was received on 7 June 2019 and accepted on 28 March 2024. The Editor was Ekaterina Zhuravskaya.

The data and codes for this paper are available on the Journal repository. They were checked for their ability to reproduce the results presented in the paper. The replication package for this paper is available at the following address: https://doi.org/10.5281/zenodo.10884477.

We thank the editor, Rachel Kranton and two referees for a number of valuable suggestions. We also thank Daron Acemoglu, Jonathan Cantor, Genevieve Gregorich, Navin Kartik, Max Kasy, Ilyana Kuziemko, Rajiv Sethi, Cosma Shalizi, Laurence Wilse-Samson and seminar audiences at CIFAR, EMNLP, New York University, IAS, University of North Carolina, Columbia Business School, Harvard Business School, Princeton University, University of British Columbia, London School of Economics, Clemson University, the Santa Fe Institute and ETH-Zurich for valuable comments and feedback. Eugenie Dugoua, Nicolas Longuet-Marx, Natalie Carlson, Goran Laraveski, Natasha Plotkin and Qing Zhang provided excellent research assistance. We thank the AEA and JSTOR for the provision of data and are grateful to the Sanford C. Bernstein & Co. Center for Leadership and Ethics for partial funding. All errors are our own.

classifier achieves out-of-sample predictive performance comparable to many other social science applications (Zeng *et al.*, 2017; Berg *et al.*, 2020), and that the predicted ideologies (or predicted partisanship, which we use interchangeably in this paper) are correlated with responses from the Initiative for Global Markets survey of leading economists scored as liberal or conservative by Gordon and Dahl (2013).

We show that patterns of predicted ideology (or partisanship) across economics journals, measured as journal fixed effects, are consistent with measures from other work. We also show that economists exhibit substantial sorting on predicted ideology by field and department. Our main application of these predicted ideologies is to examine their relationship with published empirical papers. We draw policy-relevant elasticities from Fuchs *et al.* (1998) and locate available survey papers that compile estimates of these parameters. We collect estimates of the gender gap, returns to job training, labour supply elasticities, minimum wage elasticities and union productivity effects. We show that empirical results in several policy relevant fields in economics are correlated with the predicted political ideology of the author(s), with predicted liberals (conservatives) reporting elasticities that imply policies consistent with more interventionist (laissez-faire) ideology. While unable to rule out all sources of omitted variable bias, these specifications are robust to numerous alternative measures and sets of control variables, which we summarise using specification curves (Simonsohn *et al.*, 2020).

Our paper contributes a methodology for measuring ideology in academic economics that could be extended to other technical or putatively non-partisan domains of writing. Most research economists do not publicly announce any partisan position. Indeed, many of the professional practices and norms of economics are designed to eliminate partisanship from research. For example, the National Bureau of Economic Research (NBER) does not allow explicit endorsements of policy in its working paper series. In order to extract a measure of partisan ideology from academic research, we extend supervised learning methods from natural language processing (benchmarked against other methods in our companion short paper (Jelveh *et al.*, 2014)). Our approach is novel in that it allows the frequency of a phrase to have a different political valence, depending on the topic (e.g., JEL code) of the paper. This flexible and rich representation of academic language allows us to disentangle the partisanship of an author from the partisanship of their article's research field.

While models predicting ideology from text can show high predictive accuracy, they have not been applied in technical domains where partisanship is not immediately apparent. Importantly, detecting ideology in domains where institutions and norms are in place to maintain neutrality is different from predicting ideology in domains where it is overt, such as media or political speech.[1] Adjusting for topics may be particularly important in highly specialised domains, where language use is tailored to very narrow audiences of other experts.

If political preferences were irrelevant for academic research in economics, predicting political behaviour from academic writing should be very difficult. Furthermore, it is natural to hypothesise that while detecting partisanship in popular media or politician speech is reasonably easy, doing so in specialised, technical domains may be much harder. Nonetheless, our method generates good out-of-sample predictions of economist political behaviour based on academic writing alone. Furthermore, by using written language as the set of features for prediction, we can also produce article- and journal-specific predictions of ideology, and we show that the latter accord with other measures produced in the literature. Methods like ours may be useful for extracting ideology

[1] Vafa *et al.* (2020) showed that unsupervised methods of text classification work extremely well in measuring partisanship in a sufficiently rich text model.

from highly specialised, yet also partisan, fields like climate science, public health (particularly during COVID-19) and many engineering disciplines that are of immediate relevance to policy makers.

Why focus on economics to study political preferences in academic research? One reason is the simple lack of Republicans in other social sciences, reducing the power of statistical methods to detect partisan differences.[2] Economics also influences policy more than any other social science, with economists accounting for almost 70% of all PhD social scientists testifying before Congress (Maher *et al.*, 2020), and cited more than any other discipline in both the *New York Times* and The Congressional Record (Wolfers, 2015).[3] In the United States, the Council of Economic Advisors has no analogue in the other social sciences, and the representation of economists in institutions such as the Congressional Budget Office, the Federal Reserve, the Federal Trade Commission, the Department of Justice and other agencies is far larger than that of any other social science. Empirical work in economics informs policy proposals and evaluations, and economists often testify before Congress. More broadly, economic ideas are important for shaping economic policy by influencing the public debate and setting the range of expert opinion on various economic policy options (Rodrik, 2014).

Despite their importance in shaping policies, economists share a long-standing self-conception as apolitical. Stigler (1959) argued that while professional economics was averse to sudden, large, changes in its orientation, advances in economic science were non-partisan due to institutionalised incentives and norms for the dissemination of information. Stigler writes: 'The dominant influence upon the working range of economic theorists is the set of internal values and pressures of the discipline' (Stigler, 1960, p.40). Stigler believed that political and policy preferences do not drive economic research, and when they do, it is for the worse.[4] This belief that economics conforms with standard scientific norms[5] is the basis of a working consensus that is widely defended.[6]

Yet, the evidence for the view that scientific practices purge ideology from economics is surprisingly thin, relying upon surveys or subjective coding of political beliefs. The best evidence comes from a comprehensive survey undertaken by Fuchs *et al.* (1998), who asked a number of labour and public finance economists their views on parameters, policies and values. They concluded that 'one of the most important empirical results of this study is the strong correlation between economists' positions and their values, but an understanding of this relationship requires further research' (Fuchs *et al.*, 1998, p.1415). Closest to our paper is Gordon and Dahl (2013), who applied clustering techniques to the Institute for Global Markets (IGM) survey responses from prominent economists on a variety of policy questions to assess whether economists are divided over policy issues.

[2] Economics has more registered Republicans than any other social science, although they still are a minority. Cardiff and Klein (2005) used voter registration data in California to rank disciplines by Democrat to Republican ratios. They found that economics is the most conservative social science, with a Democrat to Republican ratio of 2.8 to 1. This can be contrasted with sociology (44 to 1), political science (6.5 to 1) and anthropology (10.5 to 1). Consequently, there is more ideological diversity in economics. Langbert (2020) found that the highest positions in the American Economics Association are overwhelmingly filled by registered Democrats and, among contributors, Democratic contributors.

[3] Fourcade *et al.* (2014) showed that the high status of economists is reflected in being the highest paid of the social scientists and the least likely to use interdisciplinary citations.

[4] Stigler (1960, p.43) continued 'Often, of course, the explicit policy desires of economists have had a deleterious effect upon the theory itself.... the effect of policy views on the general theory.... has stemmed from a feeling that the theory must adapt to widely held humanitarian impulses.'

[5] For example, norms as articulated, for example, by the sociologist Merton (1942).

[6] For example, see Chetty (2013).

Instead of survey-based methods, our paper uses the correlations between patterns of academic writing and observed political behaviour to forecast ideology.[7] Ideology extraction from text has received attention from multiple fields, including computer science, political science and economics. Gentzkow *et al.* (2018) provided overviews of many models used in the analysis of text, particularly in the domain of political behaviour. While our text-and-behaviour-based measure may mitigate some of the non-response and social desirability bias that may affect surveys, the selected nature of our political behaviour data may introduce other biases, which we discuss below.

Several papers investigate the determinants of economic publication and citation patterns (Ellison, 2011; 2013; Önder and Terviö, 2015; Card and DellaVigna, 2020). None of these papers look at predicted political ideology of economics articles, and none use the text of economics articles themselves as data. Instead, they analyse citation patterns or publication counts alone.[8]

Our paper is also the first to show correlations between predicted political ideologies and empirical results. We build on the policy-relevant classification of empirical estimates done by Fuchs *et al.* (1998), who classified a range of empirical parameters into implied liberal and conservative directions. Using collections of these estimates analysed by published meta-analyses, we show that there is a significant and robust correlation of our predicted ideology scores with empirical results. While we lack the data and the empirical design to establish causality, we think that these correlations are informative and worthy of further research.

1. Data and Methodology

Our methodology is straightforward, and we preview it now. We begin by linking economists to two measures of political behaviour, campaign contributions and petition signings, to measure economists as conservative (+1) or liberal (−1) on a binary scale. Online Appendix A.5 discusses results from using each measure separately, and confirms that while they are correlated, there is independent information in each measure. We next link these authors to a corpus of academic economics articles obtained from JSTOR and NBER. We then use random forests to predict ideology from academic economics text, adjusting for unsupervised topics (via a correlated topic model) as well as imputed *Journal of Economic Literature* codes. We then show that our prediction varies primarily at the author level, and has good out-of-sample performance within the sample of authors for whom we measure behaviour. We detail each of these steps below.

1.1. *Linking Economists to Their Political Activity*

To define our set of economists, we obtained the member directory of the American Economics Association (AEA) for the years 1993, 1997 and 2002 to 2009. From these lists, we extracted over 53,000 potential authors, along with their name, location, email address, education, employer

[7] Fuchs *et al.* (1998) only surveyed economists at top forty schools, and had only a 50% response rate. The IGM survey only looks at a small sample of 'top' economists, and tends to be more Democratic than average by our measure, as we show below.

[8] Zingales (2014) looked at papers in managerial compensation, and found that top journals are more likely to publish papers that suggest that managerial pay increases are optimal and that IGM-surveyed economists who serve on boards are more likely to disagree with the statement that CEOs are paid more than their marginal productivity.

and occupation.[9] We then link the AEA member directory to two datasets with observed political behaviour: political campaign contributions and petition-signing activity.

We obtain campaign contribution data from the Federal Election Commission's website for the years 1979 to 2012. Campaign committees are required to publicly disclose information about individuals who have contributed more than \$200. These disclosures contain the contributor's name, employer, occupation, state, city, zip code, transaction date and transaction amount. We match the AEA roster to these individual contributions of which there are about twenty million. Since a person's information is often recorded differently across the AEA and FEC datasets, we apply a fuzzy string matching algorithm (Navarro, 2001; Tahamont *et al.*, 2021) to member and contributor attributes. We describe the methodology and the results in full detail in Online Appendix A.2. Summary statistics on the campaign contributions are provided in Online Appendix Table A.1.

Besides campaign contributions, we also proxy economist partisan behaviour with petition signings. Our data come from Hedengren *et al.* (2010), who collected thirty-five petitions signed principally by economists. We use fuzzy string matching and manual inspection to match the signatories to our economists. Hedengren *et al.* (2010) classified petitions according to whether they advocate for or against individual freedoms. Similarly, many of the petitions exhibit viewpoints that are aligned with the political left or right, particularly on economic issues. Examples include petitions for and against federal stimulus following the 2008 financial crisis and for and against tax increases. Online Appendix Table A.2 reproduces the list of petitions from Hedengren *et al.* (2010) that includes their classification on the liberty scale along with an additional column indicating our classification. We drop petitions classified as neutral.

We take a simple approach to assigning an ideology $\theta_{i,combined}$ to an economist based on their campaign contribution and petition signing behaviour. Let $pet_{k,i}$ be the number of petitions signed by economist i aligned with partisanship k taking on values d (left leaning), r (right leaning) or u (undetermined). A similar definition applies to $contrib_{k,i}$, which is the number of campaign contributions. The following logic is then applied to assigning ideologies.

For each economist i and ideology labels $x, y \in \{d, r\}, x \neq y$,

(*a*) if $pet_{x,i} > pet_{y,i}$ and $contrib_{x,i} > contrib_{y,i}$ then $\theta_{i,combined} = x$,
(*b*) if $pet_{x,i} > pet_{y,i}$ and $contrib_{x,i} = contrib_{y,i} = 0$ then $\theta_{i,combined} = x$,
(*c*) if $pet_{x,i} = pet_{y,i} = 0$ and $contrib_{x,i} > contrib_{y,i}$ then $\theta_{i,combined} = x$.
(*d*) Otherwise, $\theta_{i,combined} = u$.

If an economist has given more times to Democrats (Republicans) and signed more left-leaning (right-leaning) petitions, the assigned ideology is left leaning (right leaning). In the cases where the economist has zero contributions (or signed no petitions) then we only consider signed petitions (contributions). If there is disagreement between the signals, or one of them is indeterminate, but non-zero (e.g., same number of Republican and Democrat contributions), we treat the ideology as undetermined. For notational brevity, we drop reference to *combined* in $\theta_{i,combined}$ for the rest of the paper.

[9] Since AEA members are drawn, not only from academia, but government and the business world, not all of these individuals have produced academic research.

Table 1. *Petition Signing and Campaign Contribution Patterns.*

	Petitions		
Contributions	Left leaning (−1)	Undetermined	Right leaning (+1)
Left leaning (−1)	164	0	0
Undetermined	0	0	1
Right leaning (+1)	0	0	73

Notes: This table shows the overlap between our two 'groundtruth' measures of ideology for the sample of economists with papers in our corpus who both signed petitions and made campaign contributions.

We choose a simple and interpretable binary measure because there seems to be no natural scale on which to measure intensity of partisanship in the data across the two measures. Both the frequency of petition signing and magnitude of contributions could be driven by professional networks and income/wealth, respectively, in addition to partisanship. Putting these very different continuous quantities on a single scale would require more assumptions. See Online Appendix A.5 for separate results for $\theta_{i,contributions}$ and $\theta_{i,petitions}$, as well as evidence that combining both sources produces at least weakly better predictions than using each separately.

There is an extremely high level of agreement across the two binary versions of these signals when considering authors who have signed petitions and made contributions. Prior to dropping authors who have undetermined ideology, there are 238 authors that made left- or right-leaning contributions *and* signed left- or right-leaning petitions. Table 1 shows the level of agreement between the two signals. We see that there are zero economists who are assigned opposing ideologies across the two measures, and only one economist who is assigned an undetermined ideology by the contribution measure and hence dropped from our sample of groundtruth authors.

A natural concern is that the two signals are picking up different dimensions of political ideology, for example cultural versus economic liberalism. When examining Online Appendix Table A.2, we see that the petitions are overwhelmingly about economic policies, except for two that are just for or against John Kerry for president. Campaign contributions, especially those to candidates or parties, are significantly harder to categorise as being motivated by particular social or fiscal concerns alone. However, the high degree of overlap between the petitions and the campaign contributions indicates that there are few partisan Democrat (Republican) economists who are conservative (liberal) on economic policy, so partisanship in this sample seems unlikely to be driven by social issues alone.

1.2. *Economic Papers Corpus*

To create our corpus of academic writings by economists, we obtained the full text of 62,888 research articles published in ninety-three journals in economics for the years 1991 to 2008 from JSTOR. We also collected 17,503 working papers from the website of the National Bureau of Economic Research covering June 1973 to October 2011, dropping any duplicates that also appear in JSTOR. These papers were downloaded in PDF format and optical character recognition software was applied to extract text.

We remove common words and capitalisation from the raw text and use a stemmer (Porter, 1980) to replace words with their morphological roots.[10] For example, a stemmer will resolve the words 'measures', 'measuring' and 'measured' to their common root 'measur'. After dropping

[10] These common words include terms not likely to be correlated with author partisanship such as 'a', 'the' and 'to'.

words or phrases that appear fewer than ten times and more than 100,000 times, we are left with 98,479 single- and multi-word phrases that will serve as predictors for our algorithm. We extract 33,579 one-word phrases (also referred to as unigrams), 56,807 two-word phrases (bigrams) and 8,093 phrases with three or more words.[11]

To further focus our attention on the phrase sequences that are most likely to contain ideological valence, we follow Gentzkow and Shapiro (2010) and rank phrases by Pearson's χ^2 statistic. Table 2 lists the phrases that are most consistently associated with left- or right-leaning ideology in our groundtruth sample of economists.[12] As we would expect from a technical corpus with peer review, the table exhibits none of the phrases often associated with partisanship by research looking at media or political text, suggesting that, for writing by academics, partisanship is likely to be encoded in much more specialised language. For example, right-leaning terms include stemmed variants of 'stock return', 'median voter' and 'rent seeking', which are typically associated with finance or political economy, and left-leaning terms include 'health insurance', 'welfare reform' and 'food stamps', which are related to health care and welfare.

These are clearly words associated with broad areas of research rather than particular policy stances or political ideologies. That they are predictive of author political behaviour is suggestive of sorting of researchers into fields on the basis of characteristics associated with partisan leanings. But, as the model in Online Appendix A.1 shows, if publications have to satisfy peer reviewers who are also sorted into fields based on partisan leanings then the partisanship revealed by a paper will be a combination of an author's partisanship and that of the audience for the paper (peer reviewers and editors). Fortunately, many economists write in a variety of research fields, allowing an individual's partisan leaning to be expressed independently of the research field.

1.3. *Accounting for Topics*

To investigate the extent to which sorting may explain the relationship between text and ideology, we construct measures of research areas, or 'topics'. Since we do not observe topics for all of the papers in our corpus, we use prediction methods from machine learning to predict topics for all papers. We map papers to topics using both unsupervised and supervised methods from machine learning, and then we predict authors' ideologies using phrase counts weighted by topic prevalence. For example, the correlation between political behaviour and the phrase 'transaction cost' is allowed to vary depending on whether the phrase is used in a labour economics or a macroeconomics topic.[13] These within-topic predictions are combined to form a final estimate of an author's political leaning. For robustness, we also predict author ideology without adjusting for topics, and show results with and without topic adjustment throughout.

If sorting into fields was driving the relationship between language and ideology, then it should be more difficult to predict ideology within fields. As we show below, not only are we able to predict ideology accurately within fields, but our topic-adjusted predicted ideologies (which are

[11] We extract multi-word phrases automatically using a modified version of the method from Mikolov *et al.* (2013) and implemented by the *gensim* module in the Python programming language. The method scores multi-word phrases by computing the normalised pointwise mutual information (NPMI), a measure of association ranging from -1 to 1. Multi-word phrases that have NPMI values closer to one are more likely to appear together than with other words.

[12] The method for ranking the phrases in Table 2 are further described in Section 1.4.

[13] For example, we see that the stemmed version of 'transaction cost' is the top right-leaning two-word phrase in *Journal of Economic Literature* (JEL) code J7 (labour discrimination) and the top left-leaning bigram in JEL code E6 (macroeconomic policy, macroeconomic aspects of public finance and general outlook). See the Online Appendix for a full list of top-leaning terms by topic.

Table 2. *Top Forty Unigrams and Bigrams Most Associated with Left-Leaning and Right-Leaning Ideologies as Measured by χ^2 Values.*

Unigram	Bigram	Other	Unigram	Bigram	Other
women	child_care	journal_post_keynesian_econom	vote	public_choic	close_end_fund
employ	post_keynesian	journal_econom_issu	insur	rent_seek	journal_polit_economi
work	labor_market	long_term_care_insur	disclosur	stock_return	blackwel_publish_ltd
wage	social_capit	canadian_public_polici_analys	advertis	brown_williamson	london_school_econom_polit
care	minimum_wage	labor_market_outcom	hayek	social_secur	journal_financi_econom
famili	singl_mother	public_polici_analys_politiqu	tariff	bank_japan	journal_monetari_econom
union	health_care	review_intern_polit_economi	senat	child_support	unit_root_test
hospit	health_insur	minimum_wage_increas	court	path_depend	bid_ask_spread
social	low_wage	long_term_care	voter	life_expect	public_choic_kluwer_academ
train	african_american	child_care_subsidi	incumb	median_voter	publish_print_netherland
industri	credit_union	capit_account_liber	tullock	time_seri	american_journal_econom_sociolog
plan	welfar_reform	live_wage_ordin	rule	cite_note	impuls_respons_function
poverti	tax_expenditur	feder_fund_rate	shock	unit_root	journal_law_econom
canada	food_stamp	industri_labor_relat	contract	human_capit	copyright_john_wilei_son
employe	labor_forc	foreign_direct_invest	cartel	stock_price	journal_risk_insur
mother	low_incom	singl_parent_famili	yeager	switch_cost	wall_street_journal
forest	new_orlean	low_incom_famili	politician	properti_right	digit_sic_industri
keyn	food_expenditur	labor_forc_particip	litig	network_extern	southern_econom_journal
occup	industri_relat	cambridg_journal_econom	liggett	self_insur	ltd_appl_econ
children	profit_share	intra_industri_trade	arbitr	life_insur	social_secur_benefit
global	high_perform	african_american_women	cigarett	genet_algorithm	monetari_polici_shock
china	work_forc	canada_unit_state	fraud	drug_enforc	american_polit_scienc_review
unemploy	poverti_line	earn_incom_tax_credit	microsoft	null_hypothesi	line_item_veto
caregiv	treatment_group	brook_trade_forum	candid	smoot_hawlei	test_unit_root
manag	worker_compens	health_insur_coverag	return	drug_arrest	springer_public_choic
poor	employe_ownership	food_stamp_program	grower	insid_trade	review_financi_studi
survei	black_women	journal_human_resourc	measur	journal_financ	fama_french_factor
need	non_profit	human_resourc_practic	legisl	black_death	secur_exchang_commiss
plant	black_men	treatment_control_group	price	school_district	capit_labor_ratio
site	suicid_attack	labor_relat_review	index	crime_rate	strongli_disagre_strongli_agre
provinc	live_arrang	brook_paper_econom_activ	elect	cross_hold	ludwig_von_mise
veblen	white_men	new_labor_forum	model	econom_freedom	cobb_dougla_product_function
gender	head_start	high_school_degre	polic	sampl_period	smoot_hawlei_tariff
interview	low_skill	foreign_tax_credit	bureaucrat	bond_rate	digit_sic_code
percent	collect_bargain	sourc_author_calcul	payoff	transact_cost	journal_legal_studi
arrear	labour_market	meet_associ_evolutionari	state	friedman_schwartz	feder_trade_commiss
sector	men_women	journal_human_resourcest	bond	public_good	major_leagu_basebal
invest	welfar_recipi	health_care_financ	contest	toll_road	balanc_growth_path
cohort	sexual_orient	public_us_microdata	beta	bond_price	india_sri_lanka
polici	visibl_minor	annal_the_american	issuer	firm_s	foreign_exchang_market
child	wage_inequ	low_birth_weight	market	new_zealand	journal_intern_monei_financ
ford	labor_suppli	work_hour_week	cattl	journal_law	child_support_enforc
respond	inform_sector	journal_post_keynesian	legislatur	law_enforc	overlap_gener_model
cent	white_women	nation_research_council	bank	district_court	spot_exchang_rate
stet	labor_relat	world_bank_washington	steel	properti_crime	commerci_financi_chronicl
canadian	work_er	new_york_citi	softwar	law_econ	grade_point_averag
girl	ration_choic	nest_logit_model	test	drug_us	journal_polit_economi_august
hour	rel_wage	author_associ_professor	appl	reserv_price	error_correct_term
skill	world_bank	long_term_contract	competitor	monetari_base	conceal_handgun_law
household	manag_care	univers_massachusett_amherst	antitrust	major_parti	journal_econom_dynam_control
adult	race_gender	washington_world_bank	period	sex_educ	abnorm_stock_return
liber	emploi_ment	dual_labor_market	volatil	toll_free	pareto_effici_alloc
black	job_search	instrument_variabl_estim	student	child_labor	cross_countri_variat
develop	cge_model	statutori_tax_rate	auction	standard_deviat	exchang_rate_volatil
profit	public_assist	labor_market_flexibl	attornei	supra_note	brigham_young_univers
health	new_keynesian	union_non_union	block	monei_growth	american_journal_polit_scienc
estim	high_skill	monthli_review_press	trade	broker_loan	stock_price_reaction
cost	statist_canada	labor_market_condit	consolid	liabil_rule	journal_polit_economi_june
migrat	north_korea	discret_choic_model	statist	pressur_group	bureau_censu_histor_statist
marx	travel_cost	fast_food_restaur	variabl	load_factor	univers_texa_dalla
woman	live_wage	politiqu_vol_xxix	polit	primari_elect	feder_reserv_bank_minneapoli
neoclass	south_korea	sub_saharan_africa	wealth	bad_type	ohio_state_univers
firm	critic_think	latin_america_caribbean	hawaii	asset_price	stock_market_reaction
nafta	women_ag	world_bank_econom_review	reput	confer_committe	journal_econom_educ

Table 2. *Continued*

Unigram	Bigram	Other	Unigram	Bigram	Other
local	famili_incom	tight_labor_market	confere	privat_properti	unit_state_coloni
workplac	natur_resourc	annal_the_american_academi	nyse	spot_rate	likelihood_ratio_statist
treatment	job_secur	labor_mar_ket	dissip	moral_hazard	toward_theori_rent_seek
ontario	middl_class	food_stamp_recipi	quot	financial_market	shall_issu_law
feminist	singl_parent	aldershot_edward_elgar	tournament	stock_exchang	fail_reject_null_hypothesi
institut	hour_work	ag_ag_squar	villag	growth_rate	american_econom_review
school	work_class	current_popul_survei	amend	bank_failur	exchang_rate_chang
categori	new_drug	congression_budget_offic	size	forward_rate	belslei_kuh_welsch
nation	marri_women	politiqu_vol_xxx	member	gener_elect	phillip_perron_test
inventor	loss_ratio	fair_poor_health	regress	test_statist	granger_causal_test
india	percentag_point	nuclear_power_plant	cointegr	cointegr_vector	real_exchang_rate
actor	effici_wage	imf_world_bank	station	forecast_error	dominion_bureau_statist
drug	marri_mother	high_school_graduat	rank	vote_share	busi_cycl_asymmetri
librari	east_asia	labor_market_experi	data	privat_label	royal_econom_societi
slave	nurs_home	north_american_free_trade	gasolin	tempor_aggreg	strateg_trade_polici
driver	women_men	social_scienc_human	stationari	market_valu	univers_chicago_all_right
race	radic_polit	black_non_hispan	crime	market_structur	statist_report_parenthes
hispan	non_hispan	nber_work_paper_cambridg	perform	patent_law	intern_trade_commiss
parent	random_assign	self_manag_team	parti	emot_intellig	dickei_fuller_test
fisheri	pollut_abat	secondari_school_enrol	congress	catastroph_loss	western_econom_journal
asia	tax_benefit	russel_sage_foundat	agent	cross_section	long_run_growth
practic	affirm_action	min_min_min_min	bondhold	valu_weight	springer_econ_growth
cooper	labor_manag	politiqu_vol_xxiv	investor	control_variabl	augment_dickei_fuller_test
provinci	bank_canada	substanc_abus_treatment	forfeitur	bond_issu	long_run_relationship
birth	financi_crisi	univers_british_columbia	announc	futur_contract	mortgag_back_secur
trip	wage_structur	jerom_levi_econom	game	lag_length	barro_sala_martin
japanes	high_wage	colleg_high_school	sport	monei_demand	per_capita_gdp
mexico	human_resourc	author_assist_professor	bidder	campaign_contribut	journal_busi_econom_statist
credit	unit_labor	research_council_canada	discount	first_amend	free_rider_problem
strike	race_ethnic	joint_profit_maxim	predat	small_busi	montana_state_univers
surplu	work_hour	labor_market_discrimin	broker	market_share	commerci_real_estat
nonprofit	long_term	offici_poverti_line	bankruptci	random_walk	cumul_abnorm_return
flexibl	develop_countri	high_school_dropout	segment	reserv_requir	feder_reserv_bank
uniform	structur_adjust	child_care_expens	loss	capit_good	resal_price_mainten
capitalist	reserv_wage	inter_american_develop	hoover	firm_level	russian_and_east_european_financ
immigr	dai_care	journal_transport_econom_polici	radio	breton_wintrob	rent_seek_game

Notes: To determine the directionality of a particular phrase, we computed the correlation between phrase counts and ideology. If this value was positive (negative), we defined that phrase to be right leaning (left leaning).

composed of weighted averages of the topic-specific predicted ideologies) have greater accuracy than an algorithm that does not take topics into account. This points to another utility for our topic adjustments: if the relationship between language and ideology changes across fields, then accounting for those shifts can lead to more accurate predictions of ideology.

Our first method for estimating topics takes advantage of classification codes maintained by the *Journal of Economic Literature*. These codes are hierarchical markers of an article's subject area. For example, the code C51 can be read, in increasing order of specificity, as mathematical and quantitative methods (C), econometric modelling (C5), model construction and estimation (C51). Our JSTOR dataset did not include JEL codes, so we obtain classifications for 539,572 published articles and the 1.4 million JEL codes assigned to them by the *Journal of Economic Literature*.[14] The per-topic model performances are listed Online Appendix A.3. We predict codes for the first and second levels and refer to these topic mappings as JEL1 and JEL2.

[14] We were able to match and assign JEL codes to 37,364 of our JSTOR articles. The average paper was assigned to 1.90, 2.31 and 2.68 first-, second- and third-level JEL codes, respectively. We then predict codes for the set of papers that fall outside of the EconLit data. To do so, we take a 'one-versus-all' (Bishop, 2006) approach to construct a series of binary classification models, in this case gradient boosting (Friedman, 2002), a decision-tree-based classifier. For each JEL code, we take the set of papers for which we know the actual JEL codes and construct a training set where we define outcome $y_{p,j}$ as one if paper p was assigned code j and zero otherwise. We also construct a matrix of predictive features $\mathbf{C}$ where the (p, w)th element is the count of the number of times word w appeared in paper p. We estimate a series

In our second method, we use a variant of the well-known latent Dirichlet allocation (LDA) topic model, which provides an unsupervised classification of documents into latent factors, so that each document is given a probability of being in each of a number of latent 'topics'. One consequence of the Dirichlet prior used in LDA is that topic proportions are assumed independent, which is unlikely to hold in our context. To permit dependence, we use a related algorithm, the correlated topic model (CTM; Lafferty and Blei, 2006) that allows for the presence of one topic to be predictive of the presence of another, thus capturing more realistic latent topic distributions. Topic mappings were created with thirty, fifty and one hundred topics (CTM30, CTM50 and CTM100).

For each topic, we rank the words or phrases most relevant to that topic. These rankings can be used to qualitatively assess a real-world analogue to the algorithm-generated topics. We can similarly rank phrases within JEL topics by estimating the conditional probability that a word appears in a JEL topic. Online Appendix Tables A.9 to A.11 display the education topics for each mapping; note that some mappings have more than one topic that refers to education. The left-most column in each table shows the top twenty words associated with that topic, while the next two columns show the top left-leaning and right-leaning bigrams for papers in that topic, respectively.

1.4. *Predicting Ideology from Phrases*

In this section, we describe our algorithm for predicting political leanings. To recap, we have created a dataset that contains 2,471 economists who have both known groundtruth ideology as well papers in our corpus. These authors have written 20,029 papers from which we have extracted 98,479 phrases and associated counts for each paper. We have also created six mappings from papers to topics: JEL1, JEL2, CTM30, CTM50, CTM100 and NoTopic. The NoTopic mapping refers to pooling all papers without regard to topic.

The steps for our prediction algorithm proceed as follows. We first split our sample of 2,471 groundtruth authors into five partitions, or folds. We iteratively hold out one fold, which we call the *test set*, and build models on the dataset that is created by combining the four other folds, which we refer to as the *training set*. To avoid obtaining an optimistic measure of out-of-sample predictive performance, we remove co-authored papers from the training set where at least one of the co-authors is also in the test set. We then create F^{train}, a matrix where the rows represent each paper written by a groundtruth author in the training set and the columns represent phrases. The (r, p)th element in F^{train} is the number of times phrase p was used in the paper associated with row r. The mapping of rows to papers is referred to as $g(r)$. We also construct F^{test} in a similar fashion, but for test set authors.

For a topic mapping m, we iterate through each topic t, and, within a topic, we multiply each row in F^{train} by the probability that $g(r)$ was about topic t.[15] We then aggregate to the author level by summing the weighted phrase counts within author and call the resulting matrix E^{train}.[16]

of prediction models for each JEL code that generates $\hat{y}_{p,j}$, the probability that paper p is about topic t. The models perform well with an average area under the curve (AUC) of 0.96. We provide further details on AUC below.

[15] As a reminder, these probabilities are generated from the three unsupervised correlated topic models and the two supervised JEL prediction models.

[16] If a paper is written by multiple authors then that paper's phrase counts are repeated in $F^{(\cdot)}$ in as many rows as there are co-authors.

Specifically, for topic t and author i, we set

$$E_{i,\cdot}^{train} = \sum_{\{r|g(r)\in G(i)\}} \omega_{r,t,m} \cdot F_{r,\cdot}^{train},$$

where $G(i)$ is the set of papers written by author i and $\omega_{r,t,m}$ is the probability that paper $g(r)$ is about topic t under mapping m. The resulting training matrix has each row indicating a groundtruth author in the training set, and each column is a weighted sum of phrase counts, summed over the papers written by each author. We construct E^{test} in a similar fashion.

Next, we follow Gentzkow and Shapiro and filter out phrases in E^{train} that are not likely to be predictive of ideology. We create a ranking of phrases by partisanship by computing Pearson's χ^2 statistic for each phrase:

$$\chi^2_{p,t,m} = \frac{(c_{p,t,m,r}\, c_{\sim p,t,m,d} - c_{p,t,m,d} c_{\sim p,t,m,r})^2}{(c_{p,t,m,r} + c_{p,t,m,d})(c_{p,t,m,r} + c_{\sim p,t,m,r})(c_{p,t,m,d} + c_{\sim p,t,m,d})(c_{\sim p,t,m,r} + c_{\sim p,t,m,d})}.$$

Here $c_{p,t,m,\cdot}$ is the weighted counts of the number of times phrase p in topic t of mapping m was used by all economists with particular partisan behaviour (d or r) and $c_{\sim p,t,m,\cdot}$ is the number of times phrases in topic t that are not p were used. We calculate p-values from the χ^2 statistics and keep only those phrases where this value is ≤ 0.05.

To limit further, the noise that may exist in the predictors, we only keep phrases that are consistently associated with partisan behaviour. We partition the training set into five folds and hold out one fold at a time. We apply the χ^2 filter to the other four folds to identify significantly slanted phrases. We repeat this process for each possible holdout fold, which produces five sets of significant phrases. We then take the intersection across the five sets and the resulting phrases are used as input into the ideology prediction model. In other words, if a phrase is not significantly predictive of partisanship in any of these folds, it is not used in predicting the ideology of an author within a particular topic and topic mapping.

The phrases that are good predictors are intuitively plausible. In Table 2, we show the phrases that are most predictive without any topic adjustment. We keep proper names because they convey information about intellectual influences (e.g., Friedman, Keynes) and schools of thought; these are also quite a small share of our tokens (e.g., among the top hundred left-leaning bigrams only four are proper names, and among the top hundred right-leaning bigrams only seven are proper names). The top left-wing predicting bigrams are stemmed versions of child care, post Keynesian and labour market, while the top right-wing predicting bigrams are stemmed versions of public choice, rent seeking, stock returns. These are intuitively the patterns of sorting into field by predicted ideology that we would expect. But even predictive phrases within topic are intuitive. For example, the first table in Online Appendix A.6 shows phrases within Topic 19 of the CTM30 topic-adjusted prediction, which clearly corresponds to education. Within that topic, left-wing phrases are oriented towards interventionist policies such as Head Start (i.e., the federal program for children), affirmative action and the minimum wage, while right-wing phrases are associated with ability, such as human capital, cognitive skill and school attainment. This basic pattern shows up in all the topics associated with education, regardless of which specific topic adjustment is used, as can be seen in the other four tables in Online Appendix A.6.

After the phrases have been selected, we then build a model to predict authors' ideologies. Specifically, we use decision trees, a non-parametric machine learning algorithm that recursively partitions the input space into regions that seek to maximise the homogeneity of the outcome variable in each region. Partitioning is executed at each step in the tree by finding the variable that locally maximises the increase in homogeneity, as measured by the Gini impurity.[17] The advantage of decision trees is that they can model interactions without pre-specification by the analyst. A shortcoming of decision trees is that they can overfit data, i.e., find a signal where there is actually noise. To overcome this, we apply gradient boosting (Friedman, 2002), a model averaging algorithm that combines the output of a large number of trees.[18]

For a topic mapping m and an economist i, the procedure above produces a series of probabilities we call $\zeta_{i,t,m}$ that are the topic-specific probabilities that author i is a right-leaning economist. To produce a final prediction for an author, we aggregate across these topic-specific probabilities by taking a weighted average:

$$\widehat{\theta}_i = \sum_t \frac{\zeta_{i,t,m}\ P_m(Topic = t \mid author = i)}{\sum_t P_m(Topic = t \mid author = i)}.$$

Here the weights are $P_m(Topic = t \mid author = i)$, or the probability that author i writes about topic t under topic mapping m. We estimate

$$P_m(Topic = t \mid author = i) = \frac{1}{|G(i)|} \sum_{\{q \mid q \in G(i)\}} P_m(Topic = t \mid Paper = q,\ author = i),$$

averaging over all papers written by an author.

Predicted ideology values closer to zero are associated with a left-leaning ideology and values closer to one are associated with a rightward lean. To get ideologies in the $[-1, 1]$ range, we transform $\widehat{\theta}_i$ by multiplying by 2 and subtracting 1. For example, if $\widehat{\theta}_i = 0.5$, we multiple this number by 2 and subtract 1, returning the value of 0. Thus, our ideology scores are centred at 0 with a maximum value of 1 and minimum value of -1.

2. Validation

We assess the performance of our prediction model by computing the area under the receiver operating curve or the AUC (Fawcett, 2006) that can be interpreted as the probability that our classifier will rank a randomly chosen right-leaning author higher on our partisan scale than a randomly chosen left-leaning author. An AUC of one indicates that the classifier can perfectly separate left- from right-leaning authors, an AUC of 0.5 means that the classifier does no better than random guessing, and AUCs below 0.5 imply that the model actually does worse than random guessing.

Table 3 shows that all topic adjustment specifications are able to predict ideology better than random chance with AUCs ranging from 0.718 (JEL1) to 0.690 (NoTopic), comparable to other machine learning applications in social science. We also find that topic adjustments improve predictive accuracy, likely due to the ability to pickup changes in the sign of the correlation between language and ideology across fields.[19] The maximum correlation between predicted and groundtruth ideologies is 0.368. For comparison, the out-of-sample correlation reported by

[17] The Gini impurity is computed as $1 - \sum_j p_j^2$, where p is the proportion of economists of ideology j. The index is minimised when a variable perfectly splits economists into two different subspaces.

[18] We use the *lightgbm* package in the Python programming language and tune the following hyperparameters: number of trees, learning rate and maximum depth.

[19] Across all topic mappings and topics, there were 24,390 phrases that made it past the χ^2 significance filter. Of these, 15,672 appeared in multiple topics. When we look at these multi-topic phrases, we see that 32.3% (5,070) were correlated with right-leaning ideology in one topic and correlated with left-leaning ideology in another topic.

Table 3. *Predictive Performance of the Topic-Adjusted Prediction Algorithm.*

Topic map	No. topics	AUC	95% CI	Correlation	95% CI
JEL1	19	0.718	(0.697, 0.736)	0.368	(0.333, 0.400)
JEL2	99	0.698	(0.677, 0.720)	0.332	(0.294, 0.367)
CTM30	30	0.714	(0.694, 0.734)	0.364	(0.330, 0.396)
CTM50	50	0.707	(0.688, 0.729)	0.354	(0.322, 0.390)
CTM100	100	0.704	(0.683, 0.723)	0.347	(0.312, 0.378)
NoTopic	1	0.690	(0.671, 0.712)	0.326	(0.293, 0.362)

Notes: This table presents the predictive performance of various topic mappings. Listed are (1) the topic mapping, (2) the number of topics in the mapping used for prediction, (3) the area under the curve, (4) the bootstrapped confidence interval for (3), (5) the correlation between groundtruth and predicted ideologies and (6) the bootstrapped confidence interval for (5). The number of bootstrap iterations to estimate the confidence intervals was 1,000.

Gentzkow and Shapiro between their ideology measure and that obtained from another source of newspaper slant was 0.40.

For further insight into how well our model generalises, we use data from Gordon and Dahl (2013) to compare our predicted and groundtruth ideologies to responses provided by economists for a survey conducted by the Chicago Booth School of Business through 30 October 2012. The panel sets out to capture a diverse set of views from economists at top-ranked departments in the United States. Each question asks for an economist's opinion on a particular statement. The questions reflect issues of contemporary and/or long-standing importance such as taxation, minimum wages or the debt ceiling. Valid responses are 'Did Not Answer', 'No Opinion', 'Strongly Disagree', 'Disagree', 'Uncertain', 'Agree', 'Strongly Agree'.[20] Of importance here is that Gordon and Dahl (2013) categorised a set of questions where agreement with the statement implies belief in 'Chicago price theory' and disagreement implies concern with market failure. The former of these also implies a rightward lean, while the latter is consistent with left-leaning beliefs. While Gordon and Dahl (2013) found no evidence of a conservative/liberal divide in the survey responses, we find a significant correlation between the responses and our predicted ideologies. We also know the groundtruth ideology of twenty members on the panel and the correlation between groundtruth ideologies and survey responses is also significant. Following recent methods proposed by Cattaneo *et al.* (2022), Figure 1 shows binned scatterplots from a linear probability specification, conditional on question fixed effects for each of our four ideology measures. There is a clear correlation between the predicted ideology scores and the IGM-based measure of partisanship.

In order to examine this more formally, Table 4 further presents results from logit and ordered logit regressions of the form

$$Pr(response_{i,j} = C) = \Lambda(\beta_1 \widehat{\theta_i} + \delta_j),$$

where Λ is the logistic link function. In the logistic version (columns (1)–(3)), $response_{i,j}$ is a binary variable indicating whether the panellist agreed with the conservative viewpoint or not.[21] In the ordered logistic version (columns (4)–(6)), the response variable is coded with the following order: 'Strongly Disagree', 'Disagree', 'Uncertain', 'Agree', 'Strongly Agree'.[22] As seen in Table 4, the coefficients between our predicted ideology variable and the conservative viewpoint

[20] For further details on the data, see Gordon and Dahl (2013) and Sapienza and Zingales (2013). The latter showed that the IGM panel's answers to the questions are different from the answers of a random sample of the public.

[21] 'Uncertain', 'No Opinion' and 'Did Not Answer' responses where dropped for the binary logistic analysis.

[22] 'No Opinion' and 'Did Not Answer' responses were dropped for the ordered logit analysis.

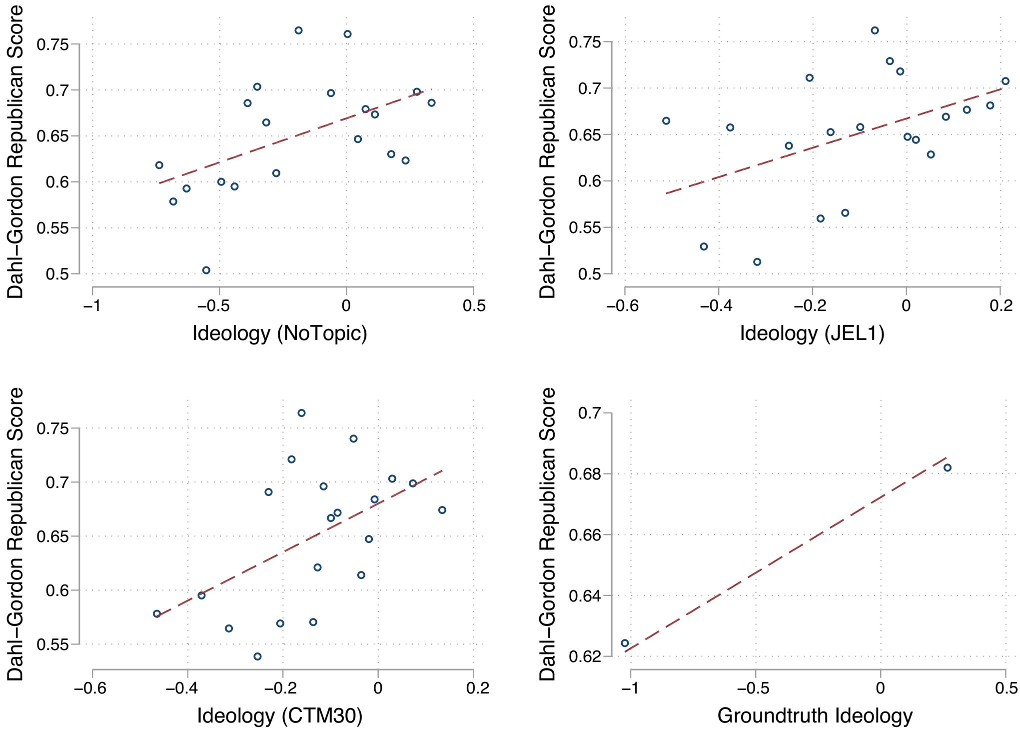

Fig. 1. *Partial Binned Scatterplots of IGM Responses on Predicted Ideology Measures.*
Notes: Plots of the mean IGM conservative answers by ventiles of predicted author ideology, conditional on question fixed effects.

are all in the expected directions and are all significant. Across all the different topic adjustments, the logit and ordered logit results in Table 4 show a significant positive relationship between our predicted ideology variables and the probability of being in an increasingly conservative category. Columns (3) and (6) add the same controls as Gordon and Dahl (2013), which are the years of the awarding of a PhD and the indicator variables for PhD institution, NBER membership, gender and experience in federal government.[23]

Finally, we present evidence that our predicted ideologies are primarily a function of individuals, not journals or time. We rerun our prediction model to produce predicted ideologies for each paper rather than each author.[24] We then decompose the variation across these paper-level predicted ideologies for each author into an author fixed effect, a journal fixed effect and a time fixed effect, following the labor economics literature using matched worker-firm data (see Abowd *et al.* (1999), henceforth AKM). For each article (co-)written by author i, in journal j, published in year t, we model ideology θ as additively separable, estimating

$$\widehat{\theta}_{ijt} = \delta_i + \delta_j + \delta_t + \epsilon_{ijt}.$$

[23] As an additional validation exercise, we run our algorithm on a corpus of editorials written by Israeli and Palestinian authors and show that we can achieve high prediction accuracy in classifying who wrote them. We discuss our performance relative to other political scaling methods more completely in our companion paper (Jelveh *et al.*, 2014).

[24] The paper-level prediction algorithm uses paper-level phrase counts, F^{train}, to predict paper-level ideologies.

Table 4. *Correlation between Predicted Author Ideology and IGM Responses.*

	(1)	(2)	(3)	(4)	(5)	(6)
Groundtruth ideology	0.274***	0.843***	15.61**	0.266***	0.393***	3.186***
	(0.0681)	(0.220)	(4.917)	(0.0640)	(0.0819)	(0.712)
JEL1	0.961*	2.265*	2.387	0.727*	1.214*	1.071*
	(0.376)	(1.081)	(1.302)	(0.333)	(0.506)	(0.442)
JEL2	1.373**	3.178**	4.318**	1.121**	1.907**	3.236***
	(0.456)	(1.230)	(1.674)	(0.407)	(0.622)	(0.605)
CTM30	1.502**	3.270*	2.818	1.145**	1.607**	1.268*
	(0.493)	(1.370)	(1.579)	(0.399)	(0.593)	(0.546)
CTM50	1.781***	3.960**	3.954*	1.430***	2.032***	2.060**
	(0.445)	(1.401)	(1.601)	(0.352)	(0.568)	(0.634)
CTM100	1.916***	4.205**	4.278*	1.497**	2.213**	1.825*
	(0.553)	(1.563)	(1.960)	(0.465)	(0.739)	(0.719)
NoTopic	0.574***	1.393**	1.025*	0.572***	0.824***	0.866***
	(0.206)	(0.573)	(0.578)	(0.173)	(0.263)	(0.227)
Question FEs	No	Yes	Yes	No	Yes	Yes
Controls	No	No	Yes	No	No	Yes
Observations	598	438	414	715	715	673
Individuals	39	39	37	39	39	37

Notes: SEs are clustered by economist. Controls include year of PhD, and binary indicators for gender, PhD university and any Federal government experience. Columns (1)–(3) are logit regressions predicting the author as conservative, as coded by Gordon and Dahl (2013) (which omits neutral answers), while columns (4)–(6) are ordered logit regressions using the five different levels of respondent agreement with statements coded by Gordon and Dahl (2013) as conservative (which includes neutral answers, and hence the larger sample size). * $p < .1$, ** $p < .05$, *** $p < .01$.

We restrict attention to articles published in journals with at least fifty articles, and include indicators for each co-author for co-authored articles. Under this additive separable assumption on the determinants of article slant, the variance of predicted ideology can be decomposed into the share explained by individual authors, the share explained by journals and the share explained by time, along with covariances across these terms. Figure 2 shows that across measures of θ, the variance is most explained by individual author fixed effects. We also show that, while explained variance is less than 50%, journals only explain 10%–15%, while individual authors explain 20%–25% and the rest is explained by the covariance of authors and journals, which suggests sorting of authors to journals along predicted ideology. Given that the original training data were individual political behaviour, the result that individual authors explain the majority of what can be explained raises our confidence that we are recovering an individual measure of ideology.

We can also use this specification to examine the contribution of journals to predicted article ideology. Figure 3 shows that the resulting estimates of δ_j correspond to existing estimates of political ideology across journals. Davis *et al.* (2011) surveyed economists and asked them their favourite journal along with an assessment of their free-market orientation, and then scored journals by the mean free-market orientation of the economists that rank them as favourite. On the sample of our journals that overlaps with theirs, their measure of 'free-market orientation' largely agrees with our predictions of conservative ideology, with a Spearman correlation of 0.87. For example, our most left-wing journal fixed effect comes from the *Journal of Post-Keynesian Economics*, and our most right-wing journal is *Public Choice*, which are exactly the lowest and highest 'free-market' journals, respectively, coded by Davis *et al.* (2011).[25] The *Journal of Political Economy* is the most conservative out of the 'top five' journals, and the *Journal of*

[25] The *Journal of Feminist Economics* has the lowest free-market score assigned by Davis *et al.* (2011), but it is not in our sample.

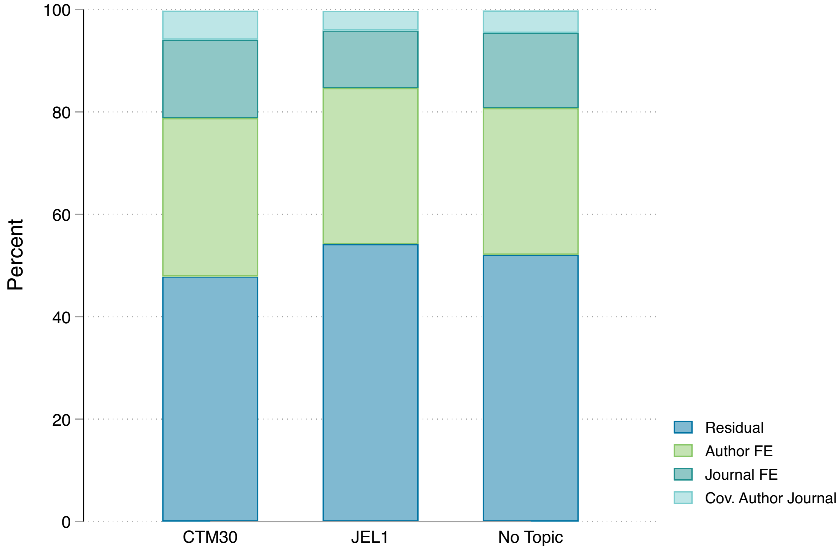

Fig. 2. *Variance Decomposition of Article-Level Predicted Ideologies.*
Notes: This figure shows variance decomposition of article-level predicted CTM30 ideology under various topic adjustments into author, journal and year components, together with covariances and the residual unexplained variation. Covariances between year and author and year and journal are too small to visualise and so are not labelled. Co-authored papers have each author fixed effect included.

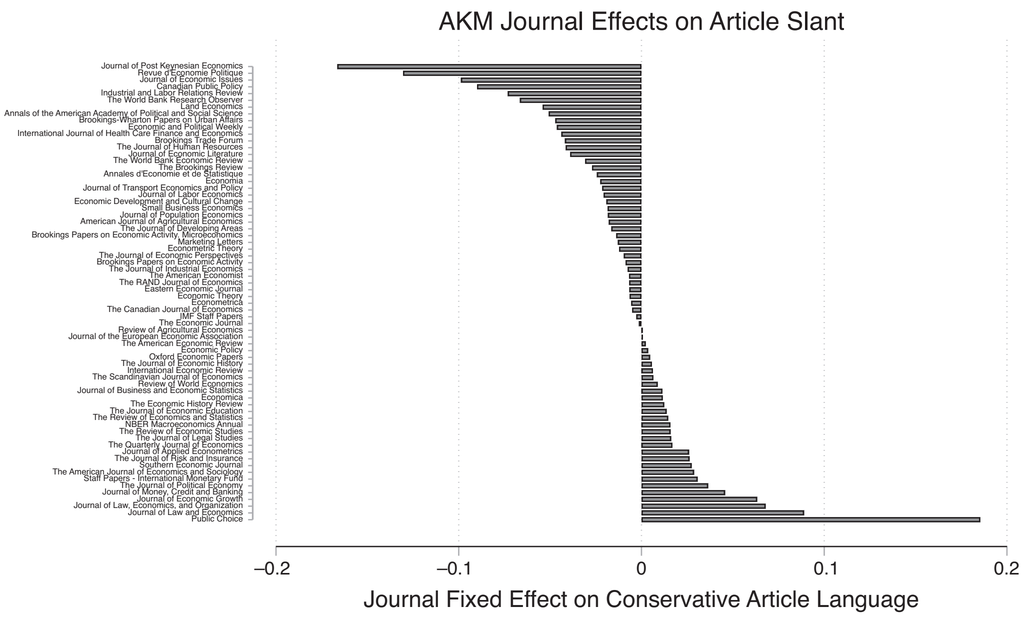

Fig. 3. *Journal Fixed Effects on CTM30 Predicted Slants.*
Notes: This figure plots journal fixed effects from the regression of predicted article ideology using the CTM30 topic adjustment on author, journal and year fixed effects.

Law and Economics (a generally conservative field historically; see Ash *et al.*, 2022) is among the most right-wing journals. Labour and development economics journals, and non-English-language journals, on the other hand, tend to be more left wing. These fixed effects show that specific journals are associated with specific political slants in the economics articles published in them, even within authors. We stress that we cannot interpret the journal effects as causal, as authors might select into publishing particular articles with particular journals, but just note that they are intuitively plausible and accord with prior research, raising confidence in our methodology.

Most relevant for the rest of the paper, individual author fixed effects explain the bulk of the variation in predicted article ideology, not year or journal effects. This finding suggests that our predicted ideology is primarily determined by authors rather than by secular changes over time or particular journals.

3. Sorting by Professional Characteristics

We link CVs of economists to our predictions and document cross-sectional patterns of predicted ideology. We start by first describing these descriptive patterns of ideology across fields of economics as well as school and career characteristics. We collect data from CVs of economists at top twenty-five departments and top ten business schools in spring 2011. We collect year and department of PhD and all subsequent employers, nationality and birthplace where available, and use self-reported fields of specialisation. Looking at self-declared primary fields, we examine labour economics, public economics, financial economics (including corporate finance), international economics and macroeconomics as determinants of political behaviour, as these are among the most policy-relevant fields in economics, but we also examine a number of other fields. We classify each department as saltwater or freshwater or neither following Önder and Terviö (2015). An economist is saltwater or freshwater if either went to grad school, had their first job or had their current job at a saltwater school (i.e., situated on the west or east coast) or freshwater school (i.e., situated in a city by one of the Great Lakes). A saltwater school is likely to be more liberal than a freshwater school.

We are interested to see if there are significant correlations between predicted political ideology and field of research. Note that while our ideology predictions account topics, self-reported fields of individuals vary independently of topic-adjusted predicted paper ideologies. Secondly, we are interested in institutional affiliations. We construct a variable for being at a business school, a top five department (Harvard, MIT, Stanford, Chicago and Princeton), as well as our indicator for 'freshwater' and 'saltwater' schools. Finally, we consider a set of demographic and professional characteristics such as Latin American origin (measured by undergraduate institution), European origin (measured by European undergraduate institution), doctoral degree year, years between undergraduate degree and economics PhD, and the number of different employers per year since obtaining a PhD.

We then look at the correlation between predicted author ideology and various CV characteristics. The estimating equation is

$$\widehat{\theta_i} = \sum_F \delta_F F_i + \gamma X_i + \delta_{phd(i)\times Year} + \epsilon_i .$$

Here $\widehat{\theta_i}$ denotes predicted ideology, F_i is a set of indicator variables for different fields of economics, X_i is a vector of other economist characteristics. We also include fixed effects

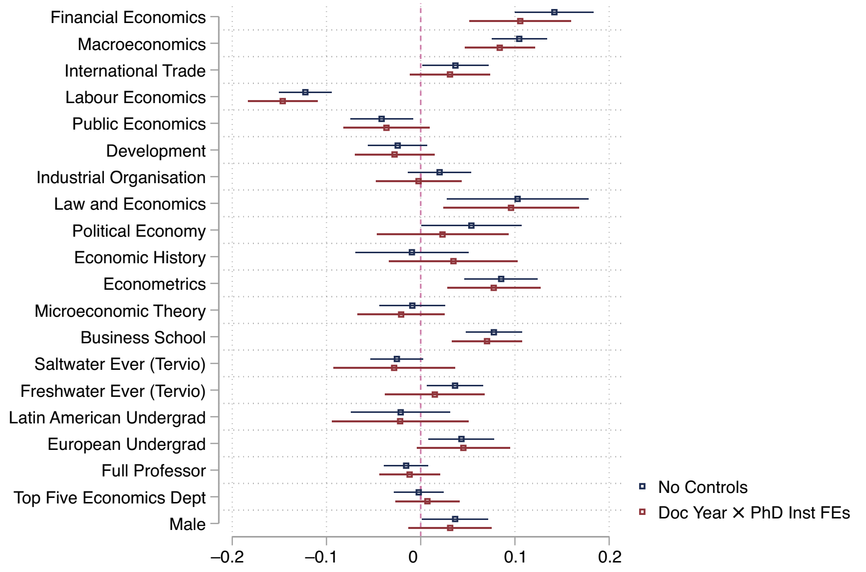

Fig. 4. *Regression Coefficients on Economist Characteristics.*
Notes: This figure plots coefficients and 95% confidence bands for coefficients on covariates from two regressions using the CTM30 ideology scores as the outcome. Coefficients are similar for all other ideology measures. The bottom set of coefficients includes no other controls; the top set of coefficients controls for the five-year interval when the PhD was obtained interacted with PhD institution fixed effects.

for PhD institution of economist i interacted with year, to see if the correlations remain robust within PhD cohorts. SEs are clustered at the department level. We vary this specification with different sets of controls, including department fixed effects and university fixed effects (there are fifteen business schools in the same university as economics departments in our sample).

Figure 4 summarises the results from the baseline specification for the CTM100 measure of ideology, although, as we show in Online Appendix A.7, results are extremely similar for all the other topic adjustments. We see that the fields of finance, macroeconomics and industrial organisation are more conservative, while labour is considerably more liberal than the average. Other fields, such as history and international trade, show less political valence. We further see that faculty at business schools are more conservative, as are professors affiliated with 'freshwater' schools, while 'saltwater' schools have a left-wing bent. Professors of European origin also seem to be somewhat more conservative, and there seems to be no association with Latin American origin, full professor rank or top five department ranking.

The finding that both the finance subfield and business schools tend to attract (or influence) economists with more conservative predicted ideology is interesting in light of the patterns documented by Fourcade *et al.* (2014), who showed that there has been a pronounced increase in economists with business school affiliations as well as in the importance of financial economics as a subfield within economics over the past few decades. These two trends, together with the political preferences documented here, may have contributed to the perception that economics is a 'conservative' field.

The magnitudes of all these coefficients should be interpreted as changes in the predicted probability of an economist being right leaning. For example, a coefficient of 0.2 indicates that the author was 10 percentage points (20 divided by the 2 that we rescale all the ideology scores by) more likely to be classified as a Republican by our prediction algorithm.

We also find that predicted ideology is persistent within individuals. As documented more fully in Online Appendix Table A.9, we split authors' writings chronologically by their first and second 50% of publications. We then predict ideology separately for each set of publications, and find that the correlation between early predicted ideology and late predicted ideology is quite high. We use this below to isolate 'early career' ideology as less likely to be influenced by the results of research.

4. Ideology and Policy Elasticities

Part of economists' influence on policy is arguably its quantitative precision. Economic theory identifies important empirical estimates that in turn imply particular optimal policies. Introductory microeconomics teaches thousands of students every semester about supply and demand elasticities, and how knowing the magnitude of the relevant elasticity tells you about the economic incidence of various policies. Economic literatures have thus developed around key empirical estimates of behavioural responses to policy. These elasticities are then used to argue, either formally or informally, for various policies. For example, the labour demand elasticity for low-wage workers can tell policymakers what are the costs and benefits of the minimum wage and empirical fiscal multipliers gauge the efficacy of government stimulus spending. Various government agencies, such as the Congressional Budget Office, the Federal Reserve and the Federal Trade Commission actively incorporate empirical economic research into policy evaluations.

This marriage of economic theory and data is well articulated, again, by Stigler (1959, p.531):

> In general there is no position, to repeat, which cannot be reached by a competent use of respectable economic theory. The reason this does not happen more often than it does is that there is a general consensus among economists that some relationships are stronger than others and some magnitudes are larger than others. This consensus rests in part, to be sure, on empirical research.

An important question, therefore, is whether predicted author political ideology predicts the magnitude of an elasticity reported in a published paper in these policy-relevant literatures. If it does, it may suggest that economists are selecting into methods or implementations of methods (e.g., *p*-hacking; see Brodeur *et al.*, 2020) that yield elasticities consistent with political beliefs. Of course, there is a possibility of reverse causation, whereby, say, liberal economists who discover elasticities that suggest that market interference is highly costly decide subsequently to contribute to the Republican party or become conservative on other issues as well. It is very difficult to identify causally any effect of predicted political ideology on empirical estimates, as any exogenous shock to partisanship could also influence the decision to be an economist, as well as the selection into what field of economics to work in. Therefore, we limit ourselves to a descriptive analysis, and discuss mechanisms below. In a robustness exercise below, we mitigate endogeneity concerns by using only ideology predicted from the first 50% of an author's papers that are in our corpus.

We select policy-relevant elasticities drawing on Fuchs *et al.* (1998) (henceforth FKP). FKP surveyed labour and public finance economists about their views on politically salient policies and parameters. FKP estimated the correlation between policy preferences and beliefs about

Table 5. *Elasticities, Meta-Analyses and Political Orientations Identified by Fuchs et al. (1998).*

Labour/ public	Type of elasticity	Surveys found	Usable data?	Policy relevant	Political orientation
Labour	Job training	Card *et al.* (2010)	No	Yes	−
Labour	Job training	Heckman *et al.* (1999)	Some	Yes	−
Labour	Labour supply	Bargain and Peichl (2016)	Some	Yes	+
Labour	Labour supply	Chetty *et al.* (2011)	Yes	Yes	+
Labour	Labour supply	McClelland and Mok (2012)	Some	Yes	+
Labour	Labour supply	Whalen and Reichling (2017)	No	Yes	+
Labour	Minimum wage	Neumark and Wascher (2008)	Yes	Yes	−
Labour	Minimum wage	Belman and Wolfson (2014)	Yes	Yes	−
Labour	Union productivity	Doucouliagos and Laroche (2003)	Yes	Yes	−
Labour	Gender wage gap	Stanley and Jarrell (1998)	No	Yes	−
Labour	Gender wage gap	Jarrell and Stanley (2004)	No	Yes	−
Labour	Gender wage gap	Weichselbaumer and Winter-Ebmer (2005)	Some	Yes	−
Labour	Labour demand	Lichter *et al.* (2015)	Yes	No	
Public	Elasticity of gasoline demand	Brons *et al.* (2008)	No	Yes	+
Public	Elasticity of gasoline demand	Espey (1998)	Yes	Yes	+
Public	Elasticity of gasoline demand	Espey (1996)	Yes	Yes	+

Notes: This table shows the set of meta-analyses of elasticities identified by Fuchs *et al.* (1998). *Usable data* indicates that the data were available from the authors. *Policy relevant* denotes whether the elasticity was relevant to a policy identified by FKP. *Political orientation* denotes whether or not the coefficient magnitude is associated with 'conservative' or 'liberal' policy choices (again, as identified by Fuchs *et al.*, 1998).

relevant economic parameter values. For example, estimates of the empirical effect of unions on productivity might influence preferences towards increased unionisation. Similarly, the female labour supply elasticity may influence the desirability of increasing Aid to Families with Dependent Children. The mapping between estimates and policies, as well as the partisan leaning, is provided in Table 5. There is one elasticity, the labour demand elasticity, that FKP did not assign to a clear policy, and consequently we denote it as 'not-policy' relevant.

We focus on estimated rather than calibrated or simulated parameters, which are mostly from the labour economics literature, as these are more comparable and studied in meta-analyses. We then looked through the literature for meta-analyses of these parameters, obtained the data from the authors where available and then merged each estimate's authors with our predicted slant measures. The list of meta-analyses is also provided in Table 5. In addition, we obtained a number of other meta-analyses from the meta-analysis archive maintained at Deakin University by Chris Doucougliasis, enabling a placebo exercise where we check the correlation between predicted author ideology and non-policy-relevant parameters.[26] We expect the correlation between predicted ideology of the authors and policy-irrelevant parameters to be insignificant.

Meta-analyses necessarily rely on the judgements of the authors about what to include and what to exclude.[27] With such diverse literatures, we take the datasets of estimates as they are, and do not process them extensively. One exception is the female gender gap, where the literature reports both the total gender gap as well as the unexplained gender gap. We construct the measure of gender wage discrimination to be the ratio of the unexplained gender wage gap to the total gender wage gap, to better account for idiosyncrasies in choices of control variables.

There are often many estimates from a single paper. When SEs are provided, we weight estimates by the inverse of the SE; otherwise, we take the simple average of estimates. These give a single estimate from each paper. We show robustness to unweighted estimates below. We

[26] See http://www.deakin.edu.au/buslaw/aef/meta-analysis/ (last accessed: 6 March 2016).
[27] Andrews and Kasy (2019) examined the econometrics of meta-analyses rigorously, and developed tests for publication bias, finding that publication bias in the minimum wage literature cannot be rejected.

adjust the sign of each estimate so that higher is more conservative, following FKP, and present these adjustments in Table 5.

Meta-analyses may have distributions of estimates that are skewed, multi-modal or truncated (as shown in Andrews and Kasy, 2019). Consequently, our primary measure is the rank of the coefficient within the category. Category refers to the policy-relevant literature (e.g., the effect of changing the minimum wage on employment). In additional specifications, we consider alternative outcome measures. We also look at a binary indicator for a coefficient being greater than the median in its category. Finally, in order to give quantitative interpretations to our point estimates, we also normalise each paper-level estimate within the survey paper, taking the z-score of its value using the mean and the SD of the elasticities reported in the survey paper.

As many estimates have multiple co-authors, we average the predicted author ideology, only for the authors for whom we are able to predict ideology (i.e., they have enough papers in our JSTOR and NBER corpus), to construct an estimated average author ideology for each paper. Unfortunately, this means that, for some papers, we only have predicted ideology for a subset of the authors, but Online Appendix Table A.14 shows that this missing predicted ideology does not seem correlated with either average predicted ideology of the co-authors we do have in our sample or with the magnitude of the FKP parameter estimates. Let β_{ps} denote the elasticity measure (rank, greater than median or standardised) from paper p in survey paper s. Our baseline regression equation is given by

$$\beta_{sp} = \gamma \overline{\theta_p} + \delta_s + \epsilon_{sp}, \tag{1}$$

where $\overline{\theta_p} = (1/|N_p|) \sum_{i \in N_p} \widehat{\theta_i}$ is the mean predicted ideology of the N_p authors of paper p from our methodology above, δ_s is a meta-analysis fixed effect, which will be included in all specifications, and ϵ_{sp} is an error term. We illustrate the basic variation using binned scatterplots in Figure 5, which shows that there is a strong correlation between our ideology measures and the coefficient rank, adjusting for meta-analysis fixed effects. This is true across our different topic adjustments, and, in fact, there is a positive correlation between groundtruth ideology and coefficient estimates.

One piece of evidence showing that our topic adjustments are indeed picking up fields of research is that the correlation between the topic-adjusted slants (JEL1 or CTM30) and the coefficients is larger than without topic adjustments; if our results were driven solely by sorting across subfields then the coefficient would shrink.

An issue arises from the generated nature of our independent variable, which, at a minimum, could bias our SEs downwards (Murphy and Topel, 2002) and could also attenuate the coefficient towards 0. As is common in high-dimensional prediction, our algorithm does not yield a straightforward SE on the prediction. While a standard solution would be to bootstrap the whole procedure, the computationally costly prediction algorithm makes the bootstrap impractical. We instead examine robustness to a split-sample instrumental variables (IV) procedure discussed below that will account for biases due to prediction error in both the coefficient as well as the SE. The bias of OLS versus IV in this case depends on an untestable assumption on whether the prediction error is uncorrelated with the truth (classical error requiring IV) or uncorrelated with the mismeasured variable (in which case OLS is unbiased, but the SEs are too small). That our estimates are qualitatively similar in both cases is thus reassuring.

Table 6 shows estimates of γ, the coefficient on mean predicted author ideology, from (1). Panel A shows results for the CTM30 adjustment, panel B for the CTM100 adjustment and panel C for the JEL1 adjustment. Column (1) shows results with coefficient values as outcome

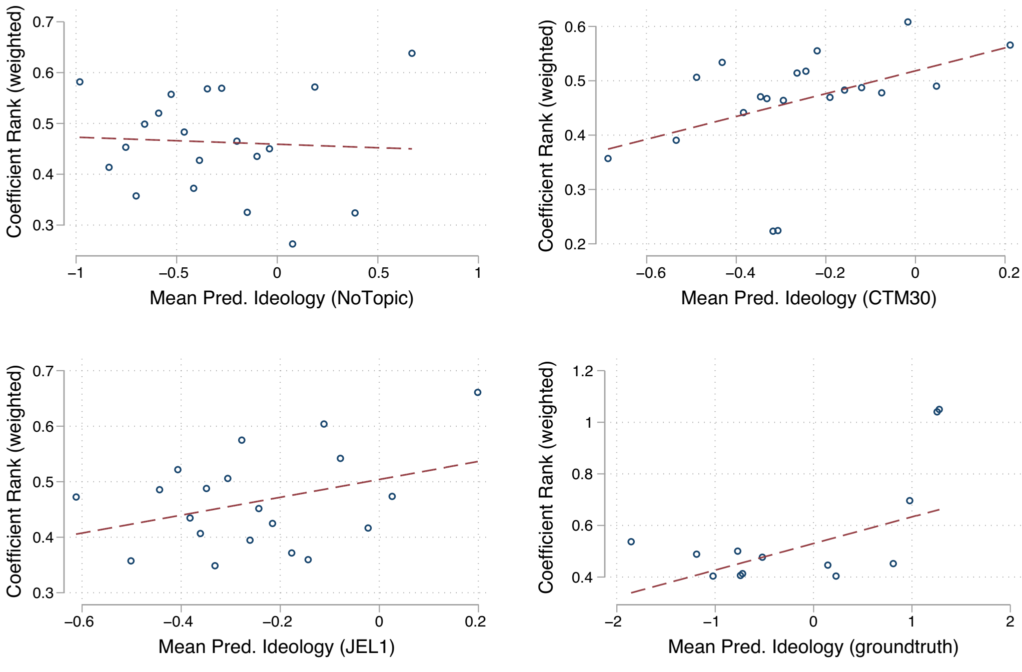

Fig. 5. *Binned Scatterplots of Coefficient Rank Against Predicted Ideology (FKP Elasticities).*
Notes: Plots of the mean elasticity rank (within category) by vintiles of predicted author ideology, conditional on meta-analysis fixed effects.

variables, with signs adjusted as described above. Column (2) shows results with the coefficient rank as the outcome variable, while column (3) shows γ when the outcome is the binary indicator variable for a high coefficient. Column (4) shows the standardised coefficient as the outcome (standardised to be a unit normal within the meta-analysis). All estimates are positive and significant.

One way to make sense of these magnitudes is to consider tax policy and the taxable income elasticity as a particular example. Building on Saez (2001), Diamond and Saez (2011) suggested top tax rates of $\tau^* = 1/(1 + 1.5 \times \epsilon)$, where ϵ is the taxable income elasticity of top income earners. The mean of the Chetty *et al*. survey on the labor supply elasticity is 0.31, suggesting a top tax rate of 68%. However, the mean ideology among people who estimate taxable income elasticities in this sample is more left than average (e.g. −0.26 in CTM30 adjusted ideology), but researchers in this area also exhibit a considerable range of ideology, from −0.55 to 0.09. Using our estimates of the impact of ideology on the elasticity (0.367), moving from the most left wing to the most right wing within this sample would change the elasticity by 0.24 points, changing the optimal top tax rate from 77% to 60%. Extrapolating to the most liberal ideology of −1 to the most conservative ideology of 1, we end up with optimal tax rates from 95% to 46%. While 46% is still a high tax rate (resulting from the small elasticities uniformly found in the literature, even by conservatives), this result shows that same standard optimal taxation formula may yield quite different prescriptions depending on the estimate, and so if partisanship is correlated with estimates, the implied policy prescription will depend on the researcher producing the elasticity.

Table 6. *Correlation between Predicted Ideology and Policy-Relevant Elasticity Coefficient Rank.*

	(1)	(2)	(3)	(4)
Panel A: CTM30 adjustment				
Mean predicted ideology strong	0.279***	0.210**	0.363**	0.832***
	(0.102)	(0.086)	(0.161)	(0.292)
Meta-analysis FEs	Yes	Yes	Yes	Yes
R^2	0.85	0.09	0.05	0.04
Observations	237	237	237	237
Ideology range	1.22	1.22	1.22	1.22
Panel B: CTM100 adjustment				
Mean predicted ideology strong	0.381***	0.230**	0.416**	1.118***
	(0.144)	(0.107)	(0.200)	(0.409)
Meta-analysis FEs	Yes	Yes	Yes	Yes
R^2	0.85	0.09	0.05	0.04
Observations	237	237	237	237
Ideology range	0.94	0.94	0.94	0.94
Panel C: JEL1 adjustment				
Mean predicted ideology strong	0.220*	0.161	0.224	0.735**
	(0.131)	(0.102)	(0.188)	(0.369)
Meta-analysis FEs	Yes	Yes	Yes	Yes
R^2	0.84	0.08	0.04	0.03
Observations	237	237	237	237
Ideology range	1.30	1.30	1.30	1.30
Panel D: groundtruth measure				
Mean predicted ideology strong	0.101	0.103	0.132	0.182
	(0.099)	(0.070)	(0.092)	(0.217)
Meta-analysis FEs	Yes	Yes	Yes	Yes
R^2	0.91	0.43	0.46	0.23
Observations	46	46	46	46
Ideology range	2.00	2.00	2.00	2.00

Notes: Robust SEs, clustered by author combination, are reported in parentheses. Ideology is calculated as the mean ideology of the authors, using ideology predicted from papers written prior to the published estimate. Coefficient rank is the rank of the average elasticity reported in the paper in the set of elasticities of the same category. High coefficient is an indicator variable for the paper elasticity being higher than the median elasticity within the same category. Standardised coefficient value is the paper's elasticity normalised by the mean and SD within category. * $p < .1$, ** $p < .05$, *** $p < .01$.

For comparison, panel D shows results with the groundtruth measure of ideology. While all the coefficients are positive and comparable in magnitude to the results in panel A, the sample of elasticities is, at $N = 31$, quite small, and the resulting SEs make the estimates insignificant at conventional levels. This shows the utility of our text-based measure: with only the groundtruth measure constructed from campaign contributions and petition signings we would not be able to predict the ideology of very many economists, but the groundtruth measure together with academic text allows us to predict ideology for many more economists, and thus expand the sample used in this regression considerably.

We examine robustness to a variety of specifications, shown in Table 7 for the coefficient rank and the CTM-adjusted ideology prediction. Column (1) in Table 7 shows coefficients from a specification that includes fixed effects for each category of estimate (e.g., labour supply elasticity) interacted with five-year bin indicators for the publication date, in order to capture

Table 7. *Correlation between Author Ideology and Policy-Relevant Elasticity Coefficient Rank Robustness (CTM30).*

	(1) Cat. × five-year FEs	(2) Unweighted	(3) US control	(4) Early pred.	(5) No GT	(6) IV	(7) Placebo
Mean predicted ideology (CTM30)	0.276***	0.211**	0.178**		0.193**		−0.012
	(0.088)	(0.085)	(0.084)		(0.096)		(0.097)
US estimate			0.083*				
			(0.046)				
Mean predicted ideology (CTM30)—early				0.167**			
				(0.075)			
Mean predicted ideology (CTM30)—IV						0.569**	
						(0.241)	
Meta-analysis FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.23	0.09	0.11	0.09	0.13	0.05	0.00
Observations	237	237	210	231	191	194	262
F-statistic						12.70	

Notes: This table presents the robustness specifications for one outcome (rank) and one topic adjustment (CTM30). Robust SEs are clustered by author combination. Outcome variable is coefficient rank within category. Ideology is calculated as the mean ideology of the authors for whom we are able to predict ideology. Column (1) includes the category of estimate × five-year period fixed effects. Column (2) uses the raw average of estimates reported in a paper, not weighting by the precision of the estimates. Column (3) controls separately for estimates on the US data. Column (4) omits any estimate where any author has a groundtruth (GT) observation. Column (5) uses ideology estimated from the first 50% of an author's written text (measuring 'Early ideology'). Column (6) presents an IV estimate using a random split of the words for each author to calculate two measures of predicted ideology and uses the first to instrument for the second. Column (7) presents a placebo estimate using non-policy-relevant elasticities from Deakin University, as described in the text. * $p < .1$, ** $p < .05$, *** $p < .01$.

observed heterogeneity in methods, data or simple improvements in estimates over time. Column (2) uses a measure that ignores the SEs attached to estimates, and instead uses the simple unweighted average of estimates within a paper. Column (3) adds an indicator variable for whether the estimate was obtained on the US data. While US estimates seem to be in a more conservative direction, the effect of predicted author ideology remains statistically significant with all three measures (albeit sometimes at only 10% significance).

In column (4), we restrict attention to predictions made using the first 50% of papers written by authors to minimise reverse causality running from empirical results to predicted ideology. These predictions are necessarily going to have more error, as they use less of the available text for each economist. Indeed, five papers (out of 197) in our sample are lost as none of the authors have enough text in the first 50% of their writings to estimate ideology. Nonetheless, the results remain positive and statistically significant despite the attenuation we would expect from the additional prediction error. In column (5), we omit any papers that have an author that is in the groundtruth sample, and the similarity of this coefficient to the rest of the table indicates that our results are not driven by the groundtruth subsample.

In column (6), we adapt split-sample instrumental variables to deal with possible prediction error in our main estimates. As discussed above, while this instrumental variables strategy does not handle endogeneity, it can address prediction error that is important to the generated nature of our independent regressor. Because our independent variable is a prediction of ideology, it has an error, akin to measurement error that attenuates the true regression slope towards zero. We split each author's writings into two random samples and predict ideology in both. Under the assumption that prediction error is orthogonal to the true ideology, then using the ideology in one sample to instrument for the ideology in the other sample will eliminate the resulting attenuation bias. Formally, if the true second-stage equation is (1), but we have prediction error in the main independent variable, we have

$$\overline{\theta_p} = \overline{\theta_p^{True}} + \eta_{sp},$$

where η_{sp} is the mean prediction error, $\eta_{sp} = (1/|N_p|)\sum_{i \in N_p} \eta_i$, akin to measurement error. Furthermore, even if η_{si} is uncorrelated with either the true value of the independent variable or any omitted variable, the estimated coefficient $\widehat{\gamma}$ will be attenuated by the well-known factor $\mathrm{var}(\overline{\theta^{True}})/[\mathrm{var}(\overline{\theta^{True}}) + \mathrm{var}(\eta)] < 1$.[28] Thus, our coefficients will be too small, relative to the true value.

Our IV strategy mitigates this problem. We split the words used by each author into two equally sized random samples, and estimate two separate, independent predictions of ideology, $\overline{\theta_p^0}$ and $\overline{\theta_p^1}$, where 0 and 1 refer to the two random samples. Unsurprisingly, these measures are highly correlated with each other. To show that the IV eliminates the influence of prediction error, we write the relationship between the predictions from the subsamples and the true value as

$$\overline{\theta_p^g} = \overline{\theta_p^{True}} + \eta_{sp}^g, \qquad g = 0, 1,$$

where η_{sp}^1 is independent of η_{sp}^0. We then use the $g = 1$ prediction as an instrument for the $g = 0$ prediction. Keeping the covariates δ_s implicit, this results in an IV coefficient given by

$$\gamma^{IV} = \frac{\mathrm{cov}(\beta_{sp}, \overline{\theta^0})}{\mathrm{cov}(\overline{\theta^0}, \overline{\theta^1})} = \frac{\mathrm{cov}(\gamma(\overline{\theta^{True}}) + \epsilon, \overline{\theta^0})}{\mathrm{cov}(\overline{\theta^0}, \overline{\theta^1})} = \gamma \frac{\mathrm{var}(\overline{\theta^{True}})}{\mathrm{var}(\overline{\theta^{True}})} = \gamma,$$

since ϵ is independent of η^1 and η^0 (which are also independent of each other). We can see the gain from the IV strategy by focusing on just the results for the CTM100-adjusted models in Table 7. As we hoped to achieve with the IV, the first-stage F-statistic is unsurprisingly extremely strong, and the coefficients are generally 20% larger than the OLS estimates, with slightly larger SEs. This serves as additional confirmation that the error in our prediction is random rather than systematically correlated with observable or unobservable variables.

Finally, in column (7) we conduct an identical exercise using ‘non-policy-relevant’ elasticities, described above. These elasticities are beta convergence in cross-country growth regressions, the value of alternative fuels, the effect of institutions on growth, the value of a recreational area and the labour demand elasticity. We again calculate rank within each category of elasticity and estimate the correlation with mean author ideology. We find no significant correlation between predicted ideology and these elasticities, and the point estimates are an order of magnitude smaller than the same specification estimated on the ‘policy-relevant’ elasticities.

While these robustness results are reassuring, they by no means exhaust the space of specifications and measures. Rather than show tables for every specification and every variant of our dependent and independent variables, we show the specification curve (Simonsohn *et al.*, 2020), a procedure to explore the sensitivity of results to modelling choices, in Figure 6. For each of the nine specifications, we estimate the specification using six different measures of ideology, three different outcomes (binary, rank and standardised coefficient) as well as two weighting schemes (coefficients within a paper averaged with the inverse of the SE where available or not). The nine specifications include three sets of covariates (controlling for category × five-year fixed effects, an indicator for the US estimate and no covariates except meta-analysis fixed effects), crossed with three identification strategies: OLS, split-sample IV and the early measure of ideology only.

[28] Even though our groundtruth measure is a binary measure, our prediction is continuous, so the measurement error can still be classical, which would not be the case if our prediction was binary.

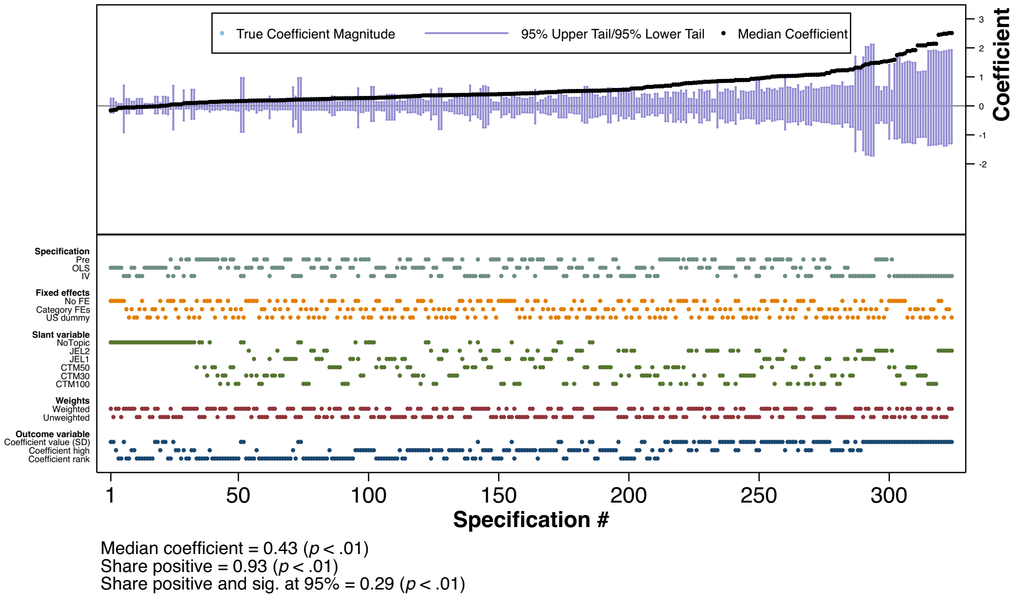

Fig. 6. *Specification Curve.*

Notes: Coefficients from 324 different specifications shown, ordered by size. Bottom left corner shows statistics testing (1) the probability that the median coefficient from a randomly shuffled sample is greater than the true median coefficient, (2) the probability that a randomly shuffled sample has at least the same share of positive coefficients as the true sample and (3) the probability that a randomly shuffled sample has at least the same share of positive and significant coefficients as the true sample.

The solid plot shows the coefficient on γ from all 324 specifications generated by the above five specifications, excluding the placebo and including the main specification from Table 7, ordered by magnitude.

For performing inference, we shuffle the independent variable randomly across observations one hundred times to create one hundred different datasets. For each dataset, we estimate each of the 324 specifications. This procedure gives us the distribution of specification curves under the null hypothesis. The bars in Figure 6 show the 5% confidence intervals for each specification. We test across all specifications jointly by counting the fraction of the hundred samples for which the estimated coefficient is greater than the median coefficient estimated in the true sample. We also measure the fraction of randomized samples that yield more specifications with positive coefficients than the true sample, as well as the fraction with more positive and significant coefficients. Across all of these statistics, less than 1% of randomised samples show more positive coefficients than the true sample. While there are some specifications that do not exceed the 95% percentile across the shuffled datasets, these are sufficiently rare across all the 324 specifications that the tests of joint significance can rule out misspecification at 99% confidence.

Finally, as another check on the general validity of our estimates, Figure 7 shows the results from dropping each category of elasticities one at a time in order to confirm that no one set of elasticities is driving our result. Across our different ideology predictions, the correlation between mean author predicted ideology and average reported elasticity generally remains significant (or nearly so) at 5%, regardless of which category is dropped.

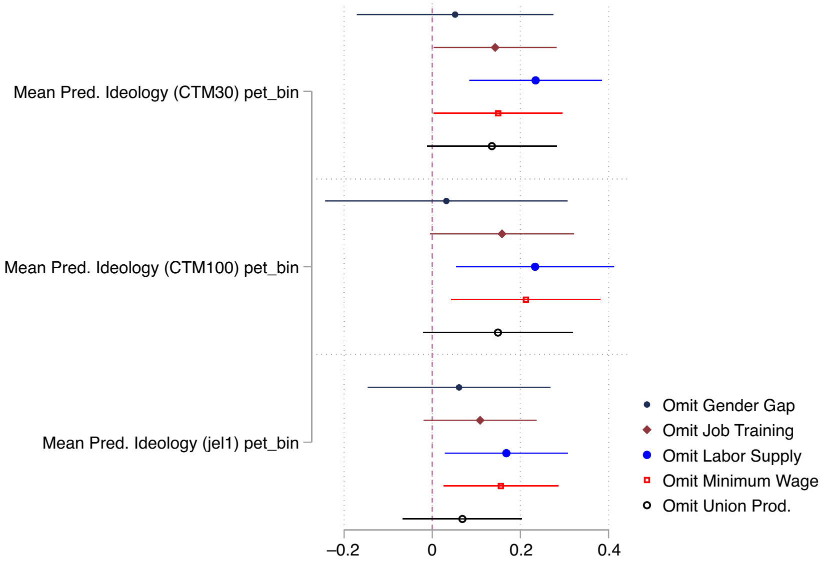

Fig. 7. *Correlation of Coefficient Rank and Ideology, Omitting Each Category of Elasticity.*
Notes: Each estimate shows correlation between the predicted ideology measure and coefficient rank, omitting a category of elasticity. We show 95% confidence windows, together with a vertical line at 0.

4.1. *Assessing Mechanisms and Threats to Validity*

The evidence supplied in this section, as we have stressed, is correlational. In this subsection, we assess the evidence for various interpretations of the correlation between elasticities and predicted partisanship.

4.1.1. *Reverse causality*

A natural concern is that our estimates are driven by reverse causality, in that economists' personal political views are influenced by the results of their research. Our results above on gender point against this, although they may be contaminated, as discussed by omitted variable bias. Furthermore, our results show the same correlation when ideology is predicted only from early papers. More tellingly, as Online Appendix Figure A.11 shows, predicted partisanship is remarkably consistent across estimates using the first 50% of an authors' papers and those using the second 50%. Consistent with a wide variety of evidence from political behaviour (e.g., Sears and Funk, 1999; Kaplan and Mukland, 2011; although see Peltzman, 2019 for evidence that people become more Republican with age), this result suggests that very few economists have their predicted partisanship change over their life cycle, suggesting that there is little evidence of researchers being 'surprised' and changing their views.

4.1.2. *Selection via methodology*

Another interpretation is that the choice of methodology is rationally a result of researcher priors. Imagine a researcher knows that journals only publish significant results, and must choose a methodology that is most likely to generate a publication. The researcher will choose a methodology most likely to generate a significant estimate, and this estimate will be correlated with their prior. While we cannot rule this out definitely, the robustness of our estimates to adjusting for topics suggests that this is not an immediate control. We examined results from the one elasticity in our sample that is estimated using a variety of methodologies, the labour supply elasticity. However, the sample is still too small to yield conclusive results.

4.1.3. *Spurious prediction*

One concern is that the small number of economists who give contributions or sign petitions themselves sort into fields, topics and methodologies that support those views, and that this induces predicted partisanship of the language used by all other economists. We have three pieces of evidence that our predictions are informative even for economists that do not contribute or sign petitions. First, we have a strong correlation between predicted author partisanship and the IGM conservativeness scores from Gordon and Dahl. Second, we have a strong correlation between the journal fixed-effects and existing rankings of journal ideology. Third, our CV regressions in Section 3 are robust to controls for observed political behavior. A variant of this concern is that the language economists use in discussing results that may support left- or right-wing policies mirror the language of advocates for those policies, generating correlations between academic writing and career trajectories or empirical results that are spurious. However, we present all of our results using only our groundtruth data: the contributions or petition signings. The coefficient on predicted partisanship remains significant in these specifications despite a much smaller set of observations, suggesting that there is a correlation between political behaviour and empirical estimates, independent of any text-based prediction.

4.1.4. *Editor political preferences*

Our results may be an outcome of the publishing process driven by the political preferences of editors rather than authors. In a previous version of this paper, we showed that even though predicted journal ideology and predicted editor ideology are highly correlated, once journal fixed effects are included, there is little correlation between predicted editor partisanship and predicted journal partisanship. This finding indicates that either authors themselves adapt their language, or a strong selection process that induces sorting jointly across editors, journals and authors, as the source of the correlation between research findings and predicted partisanship.

4.1.5. *Author political preferences*

Our tentatively favoured explanation is that many economists have political preferences that, consciously or not, may lead to particular empirical findings. At the end of the day, the lack of any clear exogenous variation in a highly persistent variable (e.g., predicted author ideology) makes it difficult to find a clean test. For example, any exogenous shock to lifetime partisanship is unlikely to be excludable, as it would likely affect many different professional choices, including the decision to become an economist in the first place. Furthermore, we do not know to what extent the preferences are driven by personal preferences versus preferences amplified by sorting into fields and methodologies. We leave the disentangling of these issues to future work.

5. Conclusion

There is a robust correlation between patterns of academic writing and political behaviour. If, in fact, partisan political behaviour was completely irrelevant to academic economic writing, then academic writing would be a very poor predictor of political ideology. However, our within-topic ideological phrases are not only intuitive, they also predict political behaviour well out of sample, and even predict the partisanship calculated from completely unrelated IGM survey data by Gordon and Dahl. The patterns of individual ideology we document are also of interest, as they suggest that there are in fact professional patterns of ideology in economics, across universities and subfields. Finally, we show that predicted ideology is correlated with empirical results on policy-relevant elasticities. We cannot claim causal identification; however, we believe that our methodology for measuring ideology and the correlations we show between predicted ideology and academic outcomes are informative.

Our paper suggests that empirical results, particularly without credible and transparent research designs, cannot be assumed to resolve questions of economic interest if results are politically contestable and economists differ too in their politics. As in the literature on self-censorship and political correctness (Loury, 1994; Morris, 2001), policy-relevant academic writing does not just reveal the results of research, but also implicit loyalties and beliefs. As academic economic articles may have potentially multiple audiences, from specialists to general interest economists to policymakers and journalists, modelling the resulting trade-offs in choosing what to say and how to explain ideas, methods and results could be a fruitful area of research (Andrews and Shapiro, 2021).

We have illustrated above how 'ideological adjustments' can, as a first pass, be flagged by considering the sensitivity of implied elasticities to ideological preferences. More ambitiously, one potential route for combining theory with the empirical approach in this paper is to develop methods for 'ideological adjustments' that incorporate the effects of sorting into summaries of parameter estimates, such as weighting results counter to a field's average ideology more highly. One simple observation is that Bayesian updating of parameters will be slower if there is known ideologically driven reporting of estimates.

Owing to a lack of data, we have not restricted analyses to empirical estimates of policy-relevant parameters with credible designs and/or pre-analysis plans. Such tools will likely reduce the influence of ideology on specific parameter estimates, but may also increase ideological sorting across research communities, as scholars search for areas that are ideologically sympathetic. Thus, we are sceptical that any purely technical solution to this fundamentally political problem can be found. Debates in economics about the extent of intervention in the market or the merits of various policies will not be resolved by better methodologies alone. A simpler alternative is to understand partisanship in economic arguments as part of the democratic process of policymaking, and acknowledge that economics itself is not outside of politics.

University of Maryland, USA
Columbia University, USA
Columbia University, USA

Additional Supporting Information may be found in the online version of this article:

Online Appendix
Replication Package

References

Abowd, J.M., Kramarz, F. and Margolis, D.N. (1999). 'High wage workers and high wage firms', vol. 67(2), pp. 251–333.

Andrews, I. and Kasy, M. (2019). 'Identification of and correction for publication bias', *American Economic Review*, vol. 109(8), pp. 2766–94.

Andrews, I. and Shapiro, J.M. (2021). 'A model of scientific communication', *Econometrica*, vol. 89(5), pp. 2117–42.

Ash, E., Chen, D.L. and Naidu, S. (2022). 'Ideas have consequences: The impact of law and economics on American justice', Working paper, National Bureau of Economic Research.

Bargain, O. and Peichl, A. (2016). 'Steady-state labor supply elasticities: A survey', *IZA Journal of Labor Economics*, vol. 5(10), pp. 1–31.

Belman, D. and Wolfson, P.J. (2014). *What Does the Minimum Wage Do?*, Kalamazoo, MI: WE Upjohn Institute.

Berg, T., Burg, V., Gombović, A. and Puri, M. (2020). 'On the rise of fintechs: Credit scoring using digital footprints', *The Review of Financial Studies*, vol. 33(7), pp. 2845–97.

Bishop, C.M. (2006). *Pattern Recognition and Machine Learning*, New York: Springer.

Brodeur, A., Cook, N. and Heyes, A. (2020). 'Methods matter: p-hacking and publication bias in causal analysis in economics', *American Economic Review*, vol. 110(11), pp. 3634–60.

Brons, M., Nijkamp, P., Pels, E. and Rietveld, P. (2008). 'A meta-analysis of the price elasticity of gasoline demand. A SUR approach', *Energy Economics*, vol. 30(5), pp. 2105–22.

Card, D. and DellaVigna, S. (2020). 'What do editors maximize? Evidence from four economics journals', *Review of Economics and Statistics*, vol. 102(1), pp. 195–217.

Card, D., Kluve, J. and Weber, A. (2010). 'Active labour market policy evaluations: A meta-analysis', Economic Journal, vol. 120(548), pp. F452–77.

Cardiff, C.F. and Klein, D.B. (2005). 'Faculty partisan affiliations in all disciplines: A voter-registration study', *Critical Review*, vol. 17(3), pp. 237–55.

Cattaneo, M.D., Crump, R.K., Farrell, M.H. and Feng, Y. (2022). 'On binscatter', *American Economic Review*, vol. 114(5), pp. 1488–514.

Chetty, R. (2013). 'Yes, economics is a science', *The New York Times*, 20 October.

Chetty, R., Guren, A., Manoli, D. and Weber, A. (2011). 'Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins', *American Economic Review*, vol. 101(3), pp. 471–5.

Davis, W.L., Figgins, B., Hedengren, D. and Klein, D.B. (2011). 'Economics professors' favourite economic thinkers, journals, and blogs (along with party and policy views)', *Econ Journal Watch*, vol. 8(2), pp. 126–46.

Diamond, P. and Saez, E. (2011). 'The case for a progressive tax: From basic research to policy recommendations', *Journal of Economic Perspectives*, vol. 25, pp. 165–90.

Doucouliagos, C. and Laroche, P. (2003). 'What do unions do to productivity? A meta-analysis', *Industrial Relations: A Journal of Economy and Society*, vol. 42(4), pp. 650–91.

Ellison, G. (2011). 'Is peer review in decline?', *Economic Inquiry*, vol. 49(3), pp. 635–57.

Ellison, G. (2013). 'How does the market use citation data? The Hirsch index in economics', *American Economic Journal: Applied Economics*, vol. 5(3), pp. 63–90.

Espey, M. (1996). 'Explaining the variation in elasticity estimates of gasoline demand in the United States: A meta-analysis', *The Energy Journal*, vol. 17(3), pp. 49–60.

Espey, M. (1998). 'Gasoline demand revisited: An international meta-analysis of elasticities', *Energy Economics*, vol. 20(3), pp. 273–95.

Fawcett, T. (2006). 'An introduction to ROC analysis', *Pattern Recognition Letters*, vol. 27(8), pp. 861–74.

Fourcade, M., Ollion, E. and Algan, Y. (2014). 'The superiority of economists', *Journal of Economic Perspectives*, vol. 28(4), pp. 1–29.

Friedman, J.H. (2002). 'Stochastic gradient boosting', *Computational Statistics & Data Analysis*, vol. 38(4), pp. 367–78.

Fuchs, V., Krueger, A. and Poterba, J. (1998). 'Economists' views about parameters, values, and policies: Survey results in labor and public economics', *Journal of Economic Literature*, vol. 36(3), pp. 1387–425.

Gentzkow, M., Kelly, B.T. and Taddy, M. (2018). 'Text as data', *Journal of Economic Literature*, vol. 57(3), pp. 535–74.

Gentzkow, M. and Shapiro, J.M. (2010). 'What drives media slant? Evidence from U.S. daily newspapers', *Econometrica*, vol. 78(1), pp. 35–71.

Gordon, R. and Dahl, G.B. (2013). 'Views among economists: Professional consensus or point-counterpoint?', *American Economic Review*, vol. 103(3), pp. 629–35.

Heckman, J.J., LaLonde, R.J. and Smith, J.A. (1999). 'The economics and econometrics of active labor market programs', in (O.C. Ashenfelter and D. Card, eds.), *Handbook of Labor Economics*. vol. 3, pp. 1865–2097, Amsterdam: Elsevier.

Hedengren, D., Klein, D.B. and Milton, C. (2010). 'Economist petitions: Ideology revealed', *Econ Journal Watch*, vol. 7(3), pp. 288–319.

Jarrell, S.B. and Stanley, T.D. (2004). 'Declining bias and gender wage discrimination? A meta-regression analysis', *Journal of Human Resources*, vol. 39(3), pp. 828–38.

Jelveh, Z., Kogut, B. and Naidu, S. (2014). 'Detecting latent ideology in expert text: Evidence from academic papers in economics', in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1804–9, Kerrville, TX: Association for Computational Linguistics.

Kaplan, E and Mukand, S.W. (2011). 'Persistence of political partisanship: Evidence from 9/11', Working paper, Department of Economics, University of Warwick.
Lafferty, J.D. and Blei, D.M. (2006). 'Correlated topic models', in *Advances in Neural Information Processing Systems*, pp. 147–54, Cambridge, MA: MIT Press.
Langbert, M. (2020). 'Republicans need not apply: An investigation of the American Economic Association using voter registration and political contributions', *Econ Journal Watch*, vol. 17(2), pp. 392–404.
Lichter, A., Peichl, A. and Siegloch, S. (2015). 'The own-wage elasticity of labor demand: A meta-regression analysis', *European Economic Review*, vol. 80, pp. 94–119.
Loury, G.C. (1994). 'Self-censorship in public discourse a theory of "political correctness" and related phenomena', *Rationality and Society*, vol. 6(4), pp. 428–61.
Maher, T.V., Seguin, C., Zhang, Y. and Davis, A.P. (2020). 'Social scientists' testimony before Congress in the United States between 1946-2016, trends from a new dataset', *PlosOne*, vol. 15(3), e0230104.
McClelland, R. and Mok, S. (2012). 'A review of recent research on labor supply elasticities', Working Paper 2012-12, Congressional Budget Office.
Merton, R.K. (1942). 'Science and technology in a democratic order', *Journal of Legal and Political Sociology*, vol. 1(1–2), pp. 115–26.
Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S. and Dean, J. (2013). 'Distributed representations of words and phrases and their compositionality', in *Advances in Neural Information Processing Systems*, pp. 3111–19, Red Hook, NY: Curran Associates.
Morris, S. (2001). 'Political correctness', *Journal of Political Economy*, vol. 109(2), pp. 231–65.
Murphy, K.M. and Topel, R.H. (2002). 'Estimation and inference in two-step econometric models', *Journal of Business & Economic Statistics*, vol. 20(1), pp. 88–97.
Navarro, G. (2001). 'A guided tour to approximate string matching', *ACM Computing Surveys*, vol. 33(1), pp. 31–88.
Neumark, D. and Wascher, W.L. (2008). *Minimum Wages*, Cambridge, MA: MIT Press.
Önder, A.S. and Terviö, M. (2015). 'Is economics a house divided? Analysis of citation networks', *Economic Inquiry*, vol. 53(3), pp. 1491–505.
Peltzman, S. (2019). 'Political ideology over the life course', Available at SSRN 3501174.
Porter, M.F. (1980). 'An algorithm for suffix stripping', *Program: Electronic Library and Information Systems*, vol. 14(3), pp. 130–7.
Rodrik, D. (2014). 'When ideas trump interests: Preferences, worldviews, and policy innovations', *Journal of Economic Perspectives*, vol. 28(1), pp. 189–208.
Saez, E. (2001). 'Using elasticities to derive optimal income tax rates', *The Review of Economic Studies*, vol. 68(1), pp. 205–29.
Sapienza, P. and Zingales, L. (2013). 'Economic experts versus average Americans', *The American Economic Review*, vol. 103(3), pp. 636–42.
Sears, D.O. and Funk, C. (1999). 'Evidence of the long-term persistence of adults' political predispositions', *The Journal of Politics*, vol. 61(1), pp. 1–28.
Simonsohn, U., Simmons, J.P. and Nelson, L.D. (2020). 'Specification curve analysis', *Nature Human Behaviour*, vol. 4(11), pp. 1208–14.
Stanley, T.D. and Jarrell, S.B. (1998). 'Gender wage discrimination bias? A meta-regression analysis', *Journal of Human Resources*, vol. 33(4), pp. 947–73.
Stigler, G.J. (1959). 'The politics of political economists', *The Quarterly Journal of Economics*, vol. 73(4), pp. 522–32.
Stigler, G.J. (1960). 'The influence of events and policies on economic theory', *The American Economic Review*, vol. 50, pp. 36–45.
Tahamont, S., Jelveh, Z., Chalfin, A., Yan, S. and Hansen, B. (2021). 'Dude, where's my treatment effect? Errors in administrative data linking and the destruction of statistical power in randomized experiments', *Journal of Quantitative Criminology*, vol. 37(3), pp. 715–49.
Vafa, K., Naidu, S. and Blei, D.M. (2020). 'Text-based ideal points', in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 5345–57, Online, Association for Computational Linguistics.
Weichselbaumer, D. and Winter-Ebmer, R. (2005). 'A meta-analysis of the international gender wage gap', *Journal of Economic Surveys*, vol. 19(3), pp. 479–511.
Whalen, C. and Reichling, F. (2017). 'Estimates of the Frisch elasticity of labor supply: A review', *Eastern Economic Journal*, vol. 43, pp. 37–42.
Wolfers, J. (2015). 'How economists came to dominate the conversation', *New York Times*, 23 January.
Zeng, J., Ustun, B. and Rudin, C. (2017). 'Interpretable classification models for recidivism prediction', *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, vol. 180(3), pp. 689–722.
Zingales, L. (2014). 'Preventing economists' capture', in (D. Carpenter and D. Moss, eds.), *Preventing Regulatory Capture: Special Interest Influence and How to Limit It*, pp. 124–51, Cambridge: Cambridge University Press.

The Economic Journal, **134** (*August*), 2439–2469 https://doi.org/10.1093/ej/ueae026 © The Author(s) 2024. Published by Oxford University Press on behalf of Royal Economic Society. All rights reserved. For permissions, please e-mail: journals.permissions@oup.com.
Advance Access Publication Date: 29 April 2024