



Munich Personal RePEc Archive

Who Said or What Said? Estimating Ideological Bias in Views Among Economists

Javdani, Moshen and Chang, Ha-Joon

University of British Columbia - Okanagan Campus, University of
Cambridge

4 February 2019

Online at <https://mpra.ub.uni-muenchen.de/91958/>

MPRA Paper No. 91958, posted 12 Feb 2019 09:33 UTC

Who Said or What Said? Estimating Ideological Bias in Views Among Economists*

Mohsen Javdani[†]

Ha-Joon Chang[‡]

February 2019

Abstract

There exists a long-standing debate about the influence of ideology in economics. Surprisingly, however, there is no concrete empirical evidence to examine this critical issue. Using an online randomized controlled experiment involving economists in 19 countries, we examine the effect of ideological bias on views among economists. Participants were asked to evaluate statements from prominent economists on different topics, while source attribution for each statement was randomized without participants' knowledge. For each statement, participants either received a mainstream source, an ideologically different less-/non-mainstream source, or no source. We find that changing source attributions from mainstream to less-/non-mainstream, or removing them, significantly reduces economists' reported agreement with statements. Using a model of Bayesian updating we examine two competing hypotheses as potential explanations for these results: unbiased Bayesian updating versus ideologically-biased Bayesian updating. While we find no evidence in support of unbiased updating, our results are consistent with biased Bayesian updating. More specifically, we find that changing/removing sources (1) has no impact on economists' reported confidence with their evaluations; (2) similarly affects experts/non-experts in relevant areas; and (3) affects those at the far right of the political spectrum much more significantly than those at the far left. Finally, we find significant heterogeneity in our results by gender, country, PhD completion country, research area, and undergraduate major, with patterns consistent with the existence of ideological bias.

Keywords: Ideology, ideological bias, authority bias, Bayesian updating, views among economists.

JEL Codes: A11, A14.

*We thank John List, Syngjoo Choi, Carey Doberstein, Julien Picault, Michele Battisti, Erik Kimborough, Ali Rostamian, Leonora Risse, Behnoush Amery, Sam van Noort, and Tony Bonen for helpful comments and discussions. We would also like to thank Michele Battisti, Julien Picault, Taro Yokokawa, and Allan Vidigal for their help in translating the survey into Italian, French, Japanese, and Brazilian Portuguese, respectively. We are also grateful to our research assistant, Jay Bell, for his assistance. We would also like to thank all the people who took the time to participate in this survey. Funding for this project was provided by Social Sciences and Humanities Research Council of Canada (SSHRC) FAS# F16-00472. This study is registered in the AEA RCT Registry and the unique identifying number is: "AEARCTR-0003246".

[†] Corresponding Author. Department of Economics, University of British Columbia – Okanagan, 3333 University Way, Kelowna, BC V1V 1V7, Canada. Tel: 250-807-9152, E-mail: mohsen.javdani@ubc.ca

[‡] Faculty of Economics, University of Cambridge, Cambridge, UK, E-mail: hjc1001@cam.ac.uk

1. Introduction

One of the dominant views in mainstream (Neoclassical) economics emphasises the positivist conception of the discipline and characterizes economists as objective, unbiased, and non-ideological. Friedman (1953) describes in his famous essay that “positive economics is, or can be, an 'objective' science, in precisely the same sense as any of the physical sciences.” Similarly, Alchian asserts that “[i]n economics, we have a positive science, one completely devoid of ethics or normative propositions or implications. It is as amoral and non-ethical as mathematics, chemistry, or physics.”¹ Boland (1991) suggests that “[p]ositive economics is now so pervasive that every competing view has been virtually eclipsed.” There exists, however, a long-standing debate about the role of ideology in economics, which some argue has resulted in rigidity in the discipline, rejection and isolation of alternative views, and narrow pedagogy in economic training (e.g. Backhouse 2010, Chang, 2014, Colander 2005, Dobb 1973, Fine and Mikonakis 2009, Fullbrook 2008, Frankfurter and McGoun 1999, Galbraith 1989, Harcourt 1969, Hoover 2003, Krugman 2009, Morgan 2015, Robinson 1973, Romer 2015, Rubinstein 2006, Samuels 1992, Stiglitz 2002, Thompson 1997, Wiles 1979, and others).

Unfortunately, there is not much direct and concrete empirical evidence to evaluate the claims made about the existence and the role of ideological bias in economics. This is partly because different sides of the argument have not been engaged in any constructive conversation about this critical topic, which has led to an increasing divide between them. Economists who reject the influence of ideological bias simply ignore and dismiss these arguments without engaging in a conversation with the critics, while some critics have taken the ideological bias in mainstream economics as a given and use it to denigrate neoclassical economics. We argue that the absence of an engaging conversation in this area is clearly not a good sign. We take the words of Alice Rivlin in her 1987 American Economic Association presidential address to heart that “economists need to be more careful to sort out, for ourselves and others, what we really know from our ideological biases.”

In this study, we are not directly investigating the credibility of the different arguments about the influence of ideological bias in economics by checking the validity of their evidence and the consistency of the conclusions drawn. We will instead take an agnostic view on these

¹ Letter from Armen Alchian to Glenn Campbell, January 20, 1969. See Freedman (2016).

discussions and rather take them as alarming signs that invoke important questions which require further investigation. We believe that the answer to whether there is an ideological bias among economists has important intellectual implications, both theoretical and practical. Theoretically, it will help us investigate the extent to which the theoretical arguments behind the positivist methodology of neoclassical economics are consistent with empirical evidence. In terms of practical implications, the answer to this question will inform the debate about the ideological views in economics, as it relates to the evolution of the mainstream economics discourse and economic training.

In order to examine the effect of ideological bias on views among economists, we use an online randomized controlled experiment involving economists in 19 countries.² More specifically, we ask participants in our online survey to evaluate statements from prominent (mainly mainstream) economists on a wide range of topics (e.g. fairness, inequality, role of government, intellectual property, globalization, free market, economic methodology, women in economics, etc.). All participants receive identical statements in the same order. However, source attribution provided for each statement is randomized without participants' knowledge. For each statement, participants randomly receive either a mainstream source (Control Group), a relatively less-/non-mainstream source (Treatment 1), or no source attribution at all (Treatment 2).

We then measure whether economists agree/disagree with identical statements to different degrees when statements are attributed to authors with different views (ideologies), which put them at different distances to mainstream economics, or when no source attributions are provided for the statements. As we discuss in more detail in [Section 4](#), implementing two different treatments could potentially allow us to distinguish between the influences of ideological bias and authority bias, which operate in similar ways. Traditional norms of modern science, such as universalism, organized skepticism, and disinterestedness, are intended to make the identity of the source of an argument irrelevant (Merton 1973). Therefore, evidence showing that changing the identity of sources affects economists' level of agreement could be consistent with the existence of ideological bias.

² By economists we mean those with a graduate degree in economics who are either academics, or work in government agencies, independent research institutions, or think tanks. The majority of economists in our sample (around 92%) are academics with a PhD degree in economics. See the data section and Table A1 in our [online appendix](#) for more details.

We find clear evidence that changing or removing source attributions and ideologies attached to them significantly affects economists' level of agreement with statements. More specifically, we find that changing source attributions from mainstream to less-/non-mainstream on average reduces the agreement level by around one-fourth of a standard deviation. These results hold for 12 out of 15 statements evaluated by participants, across a wide range of topics and ideological distances between sources. Similarly, we find that removing mainstream source attributions on average reduces the agreement level by more than one-third of a standard deviation. These result holds for all 15 statements evaluated by participants.

We use a model of Bayesian learning to examine two competing hypotheses as potential explanations for our results: unbiased Bayesian updating versus ideologically-/authority-biased Bayesian updating. Under both models, economists use source attributions as signals to form beliefs about the validity of statements. Under unbiased Bayesian updating, higher level of agreement with statements that are attributed to mainstream sources is justified by *objective* differences in credibility of mainstream sources relative to less-/non-mainstream sources. In contrast, under ideologically-biased Bayesian updating, economists *interpret* mainstream sources as more credible, not based on objective evaluation but because they are more (less) likely to confirm (disconfirm) their prior views as mainstream economists. Similarly, under authority-biased Bayesian updating, while economists might not have any particular priors or ideological views, they are more likely to agree with statements attributed to mainstream sources since they are considered as authority figures. While these different mechanisms are likely to generate similar treatment effects, we use our model to derive empirically-testable implications to examine their validity.

While we find no evidence in support of implications derived from unbiased Bayesian updating, our empirical results are all consistent with the implications from biased updating among economists. More specifically, and in contrast (consistent) with unbiased (biased) Bayesian updating, we find that changing/removing sources (1) has no impact on economists' precision of their posterior beliefs, proxied for by their reported confidence with their evaluations; (2) similarly affects experts/non-experts in relevant areas; and (3) affects those at the far right of the political spectrum much more significantly than those at the far left. In addition, as it is discussed in more detail in [Section 6](#), differences in our estimated effects of treatment 1 and treatment 2, and their

heterogeneity patterns, highlight the role of both ideological bias and authority bias in influencing views among economists.

In addition to the empirical results that support the existence of ideological/authority bias, participants' own expressed views on how to evaluate a statement lends more credibility to the hypothesis that biased updating is the driving mechanism behind our estimated treatment effects. More specifically, in an accompanying questionnaire at the end of the survey, a majority of participants (82 percent) report that a statement should be evaluated based on its content only, as opposed to its author (0.5 percent), or a combination of both (around 18 percent), which is in sharp contrast with how they actually evaluate statements. This suggests that perhaps part of the ideological/authority bias evident in our results operates through implicit or unconscious modes (Bertrand and Duflo 2017).

We also use background information collected from participants to examine whether our results vary systematically by characteristics such as gender, country, area of research, country where PhD was completed, and undergraduate major. We find that the estimated ideological bias among female economists is around 40 percent less than their male counterparts. Interestingly, on one statement in our survey which examines the issue of gender gap in economics, there is a clear significant disagreement between male and female economists, with women much more strongly agreeing with the existence of a serious and persisting gender gap in the discipline.³ In addition, on this specific statement, while men still exhibit strong ideological bias, women display no signs of ideological bias. This is perhaps due to the fact that when it comes to the important issue of gender gap in economics, which involves female economists at a personal level, women put aside ideology and focus on the content of the statement as opposed to its source.

We also find systematic and significant heterogeneity in ideological bias by country, area of research, country where PhD was completed, and undergraduate major, with some economists exhibiting no ideological bias on average and some others showing very strong bias. In addition, the heterogeneity patterns found in our results remain consistent with the existence of ideological bias (see [Section 6.4](#)).

The remainder of the paper proceeds as follows. [Section 2](#) provides a brief overview of the discussion about economics and ideology. [Section 3](#) describes our experimental design. [Section 4](#)

³ This is despite the fact that on average women's average level of agreement with our statements is lower than men, even after controlling for observed characteristics.

discusses the two competing hypotheses for our results (i.e. unbiased versus biased updating) and uses a model of Bayesian updating to organize them. [Section 5](#) describes our data and empirical methodology. [Section 6](#) presents and discusses our results. [Section 7](#) concludes.

2. Economics and Ideology: a Brief Overview

Our hypothesis regarding the potential influence of ideological bias among economists is rooted in a long-standing debate about the influence of ideology in economics. Therefore, a better understanding of this literature will better inform our analysis and the interpretation of any results associated with ideological bias. Milberg (1998) elegantly summarizes the long-standing debate about the influence of ideology in economics by stating that “the history of economic thought can in fact be read as a series of efforts to distance knowledge claims from the taint of ideology, a continuing struggle to establish the field’s scientific merit.”

About a century ago, Irving Fisher, in his presidential address to the American Economic Association, raised his concern about ideological bias in economics by stating that, “academic economists, from their very open-mindedness, are apt to be carried off, unawares, by the bias of the community in which they live.” (Fisher 1919). Other prominent economists such as Joseph Schumpeter and George Stigler also made substantial contributions to this discussion over the next few decades (see Schumpeter (1949) and Stigler (1959, 1960, 1965) for examples). However, the change in the nature of economic discourse, the increasing use of mathematics and statistics, and the increasing dominance of the positivist methodology, represented by Friedman’s ‘Methodology of Positive Economics’, have reduced the concern with ideological bias in economics, which has gradually given way to a consensus that “economics is, or can be, an objective science.”⁴

Due to this prevailing consensus, the issue of ideological bias has been largely ignored within mainstream economics in the last few decades. Critics, however, argue that the increasing reliance of economics on mathematics and statistics has not freed the discipline from ideological bias; it has simply made it easier to disregard it (e.g. Myrdal 1954, Lawson 2012).

There also exists evidence that could suggest that economics has not successfully rid itself of ideological bias. For example, Hodgson and Jiang (2007) argue that due to ideological bias in economics, the study of corruption has been mainly limited to the public sector, when there is abundant evidence of corruption in the private sector (sometimes in its relation to the public sector

⁴ See Friedman (1953).

but also internally). Jelveh et al. (2018) point to ideological overtones that could be identified in public debates between prominent economists over public policy during the last financial crisis as an example of ideological bias in economics. They also point out that these perceptions of ideological bias among economists have even affected the selection of economists as experts for different government positions.⁵ Other prominent manifestations of this ideological bias include the so-called fresh-water/salt-water divide in macroeconomics (Gordon and Dahl 2013), the conflicts between liberal/conservative camps in economics (especially regarding the possible distribution-efficiency trade-off), the Borjas versus Card debate on immigration, and the ideologically charged debates over the controversial book by Thomas Piketty (2014) or over Paul Romer (2015) and his criticism that “mathiness lets academic politics masquerade as science.”

There also exists a long-standing charge laid mainly by non-neoclassical economists regarding the prevalence of ideological bias among neoclassical economists (e.g. Backhouse 2010, Fine and Milonakis 2009, Fullbrook 2008, Frankfurter and McGoun 1999, Morgan 2015, Samuels 1992, Thompson 1997, Wiles 1979). For example, summarizing the views of the Post-Autistic economics movement in France, Fullbrook (2003) argues that the economic profession is the “opposite of pluralistic” and is “dogmatically tied to value-laden neoclassical orthodoxy.” Samuels (1980) suggests that economics is much more a “system of belief than it is a corpus of verified logical positivist knowledge” and that many uses of economics “may represent only the clothing of normativism with the garments of science”. Rothbarb (1960) criticizes what Hayek calls ‘scientism’ in economics and argues that it is a “profoundly unscientific attempt to transfer uncritically the methodology of the physical sciences to the study of human action.” McCloskey (2017) asserts that economics has “deliberately clad itself in a garb of positivism, even when scholars knew the critical importance of the historical, social, and political embeddedness of their interventions.”

There are also studies that point to the ideological biases in economic training. Based on a survey of graduate students in economics, Colander (2005) raises concerns regarding how graduate training in economics may lead to biases in students’ views. For example, he argues that graduate

⁵ They point out to the following two examples: “The rejection of Peter Diamond, a Nobel laureate in economics, by Senate Republicans, as the nominee to the Federal Reserve Board, with one of the top Republicans on the Banking Committee calling him “an old-fashioned, big government, Keynesian” at the nomination hearing (see [here](#)). And, the withdrawal of Larry Summers from his candidacy for the chairmanship of the Federal Reserve Board due to strong opposition from a coalition group over several issues related to ideology, including his role to “push to deregulate Wall Street”” (see [here](#)).

training in economics induces conservative political beliefs in students. Allgood et al. (2012) also find evidence that suggests that “undergraduate coursework in economics is strongly associated with political party affiliation and with donations to candidates or parties”. Using laboratory experiments, other studies find that compared to various other disciplines, economics students are more likely to be selfish (Frank et al. 1993 and 1996, Frey et al. 1993, Rubinstein 2006), free-riding (Marwell and Ames 1981), greedy (Want et al. 2012), and corrupt (Frank and Schulze 2000)

Frey et al. (1993) attribute these results to the economic training which “neglects topics beyond Pareto efficiency [...] even when trade-offs between efficiency and ethical values are obvious.” Frank et al. (1993) highlight the exposure of students to the self-interest model in economics where “motives other than self-interest are peripheral to the main thrust of human endeavor, and we indulge them at our peril.” Rubinstein (2006) argues that “students who come to us to 'study economics' instead become experts in mathematical manipulation” and that “their views on economic issues are influenced by the way we teach, perhaps without them even realising.” Stiglitz (2002) also argues that “[economics as taught] in America’s graduate schools ... bears testimony to a triumph of ideology over science.”

Surprisingly, however, there is very thin empirical evidence to rule out or establish the existence of ideological views among economists. We are only aware of two studies that examines this issue to some extent. Gordon and Dahl (2013) use data from a series of questions from IGM Economic Expert Panel to examine to what extent prominent economists (41 economists) from the top seven economics departments disagree about key economic issues. Their results suggest that “there is close to full consensus among these panel members when the past economic literature on the question is large. When past evidence is less extensive, differences in opinions do show up.”⁶ They also find that “there is no evidence to support a conservative versus liberal divide among these panel members, at least on the types of questions included so far in the surveys.”⁷

Jelveh et al (2018) use purely inductive methods in natural language processing and machine learning to examine the relationship between political ideology and economic research.

⁶ The variable that measures the size of the economic literature related to a certain question is constructed based on judgment calls by Gordon and Dahl (2013).

⁷ They use two approaches here. First, they use different distance-based clustering methods to examine whether panel members are clustered into “two or even a few roughly equal-sized camps” based on their responses. As their second approach, they identify a subset of questions that are likely to generate disagreement among panel members, and then classify answers to these questions as either consistent with “Chicago price theory” or consistent with concerns regarding distributional implications or market failures. They then test whether participants’ responses are homogenous as a panel or are divided into two groups. They find evidence that supports the former.

More specifically, using the member directory of the AEA, they identify the political ideology (i.e. Republican versus Democrat) of a subset of these economists by (fuzzily) matching their information to publicly disclosed campaign contribution and petition signings (35 petitions). Next, using the set of JSTOR and NBER papers written by these economists with an identified political ideology, they estimate the relationship between ideology and word choice to predict the ideology of other economists.⁸ Finally, examining the *correlation* between authors' predicted ideology and their characteristics, they find that predicted ideology is "robustly correlated with field of specialization as well as various department characteristics." They suggest that results are suggestive of "substantial ideological sorting across fields and departments in economics."

3. Experimental Design

It is well understood that examining issues such as the impact of bias, prejudice, or discrimination on individual views and decisions is very challenging, given the complex nature of these types of behaviour. For example, the issue of discrimination in the labour market has long been an issue of importance and interest to labour economists. However, as Bertrand and Duflo (2017) put it, "it has proven elusive to produce convincing evidence of discrimination using standard regression analysis methods and observational data." This has given rise to a field experimentation literature in economics that has relied on the use of deception, for example through sending out fictitious resumes and applications, to examine the prevalence and consequences of discrimination against different groups in the labour market (see Bertrand and Duflo (2017) and Riach and Rich (2002) for a review. Also see Currie et al. (2014) as another example of experimental audit studies with deception).

Given that answering our question of interest is subjected to the same challenges, we take a similar approach, namely using fictitious source attributions, in order to produce reliable results.⁹ More specifically, we employ a randomized controlled experiment embedded in an online survey. Participants are asked to evaluate a series of statements presented to them by choosing one of the following options: strongly agree, agree, neutral, disagree, and strongly disagree. They are also asked to choose a confidence level on a scale from 1 to 5 for their selected answer. These statements are on a wide range of topics in economics and while they are mainly from prominent

⁸ These estimates are relied on the strong assumption that the relationship between word choice and ideology is the same among economists whose political ideology is identified through their campaign contributions and petition signings and those whose political ideology is unidentified.

⁹ See Section 3 in our [online appendix](#) for a more detailed discussion on the use of deception in economics.

(mainstream) economists, most of them challenge, to different extents, certain aspects of mainstream economics.

Our choice of critical statements, as opposed to neutral or supportive statements, is based on the idea that ideological reactions are more likely to be invoked, especially though changing sources, when one encounters views that are in contrast to his views/ideologies. Changing sources on views that one agrees with are less likely to induce an ideological reaction. Another important issue to highlight is that most of our statements are not clear-cut one-dimensional statements. Given the complex nature of ideological bias, it is more likely to arise, or to be revealed by individuals, in situations where the issues discussed are more dense, complex and multi-dimensional. This is partly due to the fact that ideological arguments are more easily concealed when it comes to more complex multi-dimensional issues.

All participants in our survey receive identical statements in the same order. However, source attribution for each statement is randomized without participants' knowledge.¹⁰ For each statement, participants randomly receive either a mainstream source (Control Group), or a relatively less-/non-mainstream source (Treatment 1), or no source attribution (Treatment 2).¹¹ See Table A8 in our [online appendix](#) for a complete list of statements and sources. Participants who are randomized into treatment 2 for the first statement remain there for the entire survey. However, those who are randomized into control group or treatment 1 are subsequently re-randomized into one of these two groups for each following statement. Moreover, those randomized into treatment 2 were clearly informed, before starting to evaluate the statements, that “All the statements that you are going to evaluate are made by scholars in economics, and do not necessarily reflect the views of the researchers. We have not provided the actual sources of these statements to make sure they are evaluated based on their content only.”

¹⁰ For the most part, the randomization was done across countries. Participants got randomized into different groups upon visiting the online survey. Since in most cases the survey was run concurrently in different countries, this led to randomization of subjects across countries.

¹¹ We are aware that in economics the notions of mainstream and less-/non-mainstream has changed a lot over time. Similar to many other classifications, this is simply an attempt to summarize a complex notion and provide an easy-to-understand relative comparison, even though we believe these differences between mainstream and less-/non-mainstream economics are more appropriately understood as a continuum rather than a dichotomy. For the lack of better classifications, therefore, we categorize our sources into these two groups to clearly summarize and represent the relative ideological distance between them. Of course, it is well-understood that this classification does not readily apply to some sources, such as older ones (e.g. Karl Marx) or sources from other disciplines (e.g. Michael Sandel or Sigmund Freud) in the same way as it applies to others. However, to remain consistent and avoid confusion for the reader, we stick to the same naming convention for all sources.

Three points are worth highlighting here. First, the actual and the altered sources for each statement were carefully paired such that they can be associated to commonly known but different views (such as different schools of thoughts, political leanings, disciplines, attitudes towards mainstream economics, etc.), so that there is an ideological distance between them and switching the source can potentially induce an ideological bias. For each source, we also provide information on their discipline, their affiliation, and the title of one of their publications. This is to further accentuate the ideological differences between the two sources for each statement, especially in cases where sources might be less known by economists.¹² Second, statements were carefully selected so that their attribution to fictitious sources is believable by participants. Third, all selected statements were relatively obscure so the misattribution would not be easily noticed by participants.¹³

4. Biased Versus Unbiased Bayesian Updating: A Model

We consider two competing hypotheses that could explain how changes in the identity of a source attributed to a statement could affect one's agreement level with the statement. We use a Bayesian updating model to organize these two hypotheses.¹⁴

Consider an economist i who is provided with a statement j to evaluate. The economist's level of agreement (A_{ij}) depends on the statement's degree of validity (θ_j^*),

$$A_{ij} = \theta_j^* + v_{ij}. \quad (1)$$

Statement j 's degree of validity is an *unknown* point on a continuum. Imperfect information about the validity of the statement could be due to not having enough knowledge about the subject, lack of conclusive empirical evidence, the statement being open to interpretation, etc.. In addition, $v_{ij} \sim N(0, d_{ij}^{-1})$ represents the component of evaluation that is not systematically

¹² For example, while some economists might not know Richard Wolff or Anwar Shaikh, knowing that they are affiliated with the University of Massachusetts Amherst or the New School for Social Research, the two famous heterodox schools in economics, makes it more likely to induce an ideological reaction. Similarly, titles of selected publication for each source, such as “*Rethinking Marxism*”, “*The Crisis of Vision in Modern Economic Thought*”, or “*What Money Can't Buy: The Moral Limits of Markets*”, serve the same purpose.

¹³ We received less than a dozen emails from people who had recognized the misattribution of a statement to a source. In all but one of these cases, the statement identified as being misattributed was statement 13 (see Table A8 in the online appendix), which is perhaps the least obscure statement used in our survey. This statement is from Adam Smith's Wealth of Nations, which was misattributed to Karl Marx. All the emails we received, however, made it clear that this was perceived as a mistake in our survey and not part of our survey design. Nevertheless, identifying misattributions, which seems to be quite rare based on emails we received, would only lead to an underestimation of the true bias effect.

¹⁴ The notation is adapted from Moretti (2011), where a similar model is used to study peer effects in movie consumption.

related to the statement's degree of validity and captures factors such as political and ideological views that could affect one's level of agreement.

We assume that the economist's belief about θ_j^* is normally distributed: $\theta_j^* \sim N(\theta_j, k_{ij}^{-1})$ where k_{ij}^{-1} is the variance of the prior (i.e. k_{ij} is its precision) which is known to the economist. The variance indicates the degree of confidence the economist attributes to his belief. It is allowed to vary across statements, reflecting the idea that the precision of economist i 's prior belief might be different across statements, for example due to his varying levels of knowledge on different topics. The variance is also allowed to vary across economists, since some economists might have less/more precise beliefs than others.

Suppose that economist i also receives a noisy signal (s_{ij}) regarding the validity of statement j in the form of an attributed source

$$s_{ij} = A_{ij} + \epsilon_{ij}, \quad (2)$$

where $\epsilon_{ij} \sim N(\beta_j, m_{ij}^{-1})$. We have in mind that the signal is a measure of how much credibility economist i gives to the attributed source in determining the validity of statement j . The noise component of the signal ϵ_{ij} has a mean of β_j . If $\beta_j = 0$, then the signal is interpreted to provide unbiased information about the validity of statement j (i.e. $s_{ij} \sim N(\theta_j^*, h_{ij}^{-1})$). In contrast, $\beta_j \neq 0$ implies perceived bias regarding the signal (i.e. $s_{ij} \sim N(\theta_j^* + \beta_j, h_{ij}^{-1})$). The variance of the signal is known and depends on the economist *interpretation* of how credible the statement's source is in determining its validity. The variance is allowed to vary across economists and statements to reflect the possibility that different economists have different interpretations regarding the validity of different signals.

We also allow ϵ_{ij} to be correlated with v_{ij} with $cov(\epsilon_{ij}, v_{ij}) = \tau_{ij}$. This is to reflect the possibility that unobserved individual-level factors, such as ideological views, that could affect one's agreement level could also affect his interpretation regarding the credibility of signals. The normal Bayesian learning model indicates that the representative economist will have posterior belief $\theta_j^* | s_{ij} \sim N(\theta_{1j}, k_{1ij}^{-1})$ where:¹⁵

$$\theta_{1j} = E[A_{ij} | s_{ij}] = (1 - \omega_{ij})\theta_j + \omega_{ij}(s_{ij} + \beta_j), \quad (3)$$

$$k_{1ij}^{-1} = \frac{1}{k_{ij} + h_{ij}}, \quad (4)$$

¹⁵ See DeGroot (1970) for more details.

where $\omega_{ij} = \frac{h_{ij}}{h_{ij}+k_{ij}}$, $h_{ij} = \frac{d_{ij}m_{ij}}{d_{ij}+m_{ij}+d_{ij}m_{ij}\tau_{ij}}$. As Equation (3) suggests, the expected agreement level is a precision-weighted average of the signal and the prior belief.

Now suppose that signals could be either a mainstream source $s_{ij(m)} \sim N(\theta_{j(m)}, h_{ij(m)}^{-1})$ or a less-/non-mainstream source $s_{ij(n)} \sim N(\theta_{j(n)}, h_{ij(n)}^{-1})$. We assume that if there exists a mainstream bias in the form of $\tau_{ij} \neq 0$, then this bias will lead to attributing more credibility to mainstream sources and less credibility to non-mainstream sources compared to the situation where there is no bias. This implies that keeping m_{ij} constant, $\tau_{ij(n)} > 0$ and $\tau_{ij(m)} < 0$.¹⁶ This is however not the only channel through which bias can manifest itself in this model. As we will discuss in more detail below, perceived differences in signal means (i.e. $\theta_{j(m)} > \theta_{j(n)}$) or in signal precision (i.e. $m_{ij(m)} > m_{ij(n)}$) could be also driven by biased evaluations of signals.

A mainstream source will induce a higher level of agreement if $\theta_j^* | s_{ij(m)} = \theta_{1j(m)} > \theta_j^* | s_{ij(n)} = \theta_{1j(n)}$. Therefore, the probability that the economist will agree more with a statement when it is attributed to a mainstream source versus a less-/non-mainstream source is:

$$P_1 = \text{Prob}(\theta_{1j(m)} > \theta_{1j(n)}), \quad (5)$$

$$P_1 = \text{Prob}\{[(1 - \omega_{ij(m)})\theta_j + \omega_{ij(m)} s_{ij(m)}] > [(1 - \omega_{ij(n)})\theta_j + \omega_{ij(n)} s_{ij(n)}]\}.^{17}$$

It therefore follows that

$$P_1 = \Phi\left(\frac{N_1}{\frac{\omega_{ij(m)}^2}{h_{ij(m)}} + \frac{\omega_{ij(n)}^2}{h_{ij(n)}}}\right), \quad (6)$$

Where $N_1 = \theta_j \left(\frac{h_{ij(n)}}{h_{ij(n)+K_{ij}}} - \frac{h_{ij(m)}}{h_{ij(m)+K_{ij}}} \right) + \frac{h_{ij(m)}}{h_{ij(m)+K_{ij}}} \theta_{j(m)} - \frac{h_{ij(n)}}{h_{ij(n)+K_{ij}}} \theta_{j(n)}$ and $\Phi(\cdot)$ is the standard normal cumulative function.

Case 1: $m_{ij(m)} = m_{ij(n)}$ and $\theta_{j(m)} = \theta_{j(n)} = \tilde{\theta}_j$.

In this case, we will have:

¹⁶ The existence of strong mainstream bias is expected to result in lower agreement with our statements (given their critical nature of mainstream economics) through smaller v_{ij} . At the same time, the same bias will result in giving more credibility (less credibility) to mainstream (less/non-mainstream) sources through larger (smaller) ϵ_{ij} . This suggests that $\tau_{ij(m)} < 0$ and $\tau_{ij(n)} > 0$.

¹⁷ We still allow for the possibility that signals could be interpreted as biased. The biased term, if it exists, will be part of the signals mean (i.e. $\theta_{j(m)}$ and $\theta_{j(n)}$).

$$N_1 = \frac{k_{ij}(h_{ij(m)} - h_{ij(n)})}{(h_{ij(m)} + k_{ij})(h_{ij(n)} + k_{ij})} (\tilde{\theta}_j - \theta_j). \quad (7)$$

In the absence of bias in the form of $\tau_{ij} = 0$, we will have $h_{ij(m)} = h_{ij(n)}$. Therefore, $N_1 = 0$ and $\Phi(0) = \frac{1}{2}$. In other words, if the economist perceives both signals as having the same mean and precision, then as expected in the absence of bias changing the source won't have an impact on the agreement level. However, in the presence of bias with $\tau_{ij} \neq 0$, we will have $h_{ij(m)} > h_{ij(n)}$ since as we discussed before we have $\tau_{ij(n)} > 0$ and $\tau_{ij(m)} < 0$. It therefore follows that $N_1 > 0$ and $P_1 > \frac{1}{2}$ as long as $\tilde{\theta}_j > \theta_j$. In other words, the existence of bias results in higher level of agreement with mainstream sources as long as the economist's evaluation of the validity of a given statement when it is attributed to a source is higher than his prior (i.e. $\tilde{\theta}_j > \theta_j$).

Case 2: $m_{ij(m)} = m_{ij(n)}$ and $\theta_{j(m)} = \tilde{\theta}_j > \theta_{j(n)} = \tilde{\theta}_j - \beta_j$ (where $\beta_j > 0$).

In this case we will have

$$N_1 = \frac{k_{ij}(h_{ij(m)} - h_{ij(n)})}{(h_{ij(m)} + k_{ij})(h_{ij(n)} + k_{ij})} (\tilde{\theta}_j - \theta_j) + \frac{h_{ij(n)}}{h_{ij(n)} + k_{ij}} \beta_j. \quad (8)$$

Similar to Case (1) and under unbiased updating with $\tau_{ij} = 0$, we will have $h_{ij(m)} = h_{ij(n)}$. Therefore, $N_1 = \frac{h_{ij(n)}}{h_{ij(n)} + k_{ij}} \beta_j > 0$ and $P_1 > \frac{1}{2}$. In other words, if the economist perceives the mean of the less-/non-mainstream signal as biased upward (i.e. making inflated claims) then he will correct for the perceived bias before updating his priors (hence $\tilde{\theta}_j - \beta_j$). Therefore, changing the source to a less-/non-mainstream source will reduce the agreement level.

If $\tilde{\theta}_j > \theta_j$, then it is easy to show that the same result holds (i.e. $P_1 > \frac{1}{2}$) when $\tau_{ij} \neq 0$.

Case 3: If $m_{ij(m)} > m_{ij(n)}$ and $\theta_{j(m)} = \theta_{j(n)} = \tilde{\theta}_j$.

In this case, as long as $\tilde{\theta}_j > \theta_j$, we will get $P_1 > \frac{1}{2}$ both when $\tau_{ij} = 0$ and $\tau_{ij} \neq 0$.

Case 4: If $m_{ij(m)} > m_{ij(n)}$ and $\theta_{j(m)} = \tilde{\theta}_j > \theta_{j(n)} = \tilde{\theta}_j - \beta_j$.

Again, similar to Case (3), as long as $\tilde{\theta}_j > \theta_j$, we will have $P_1 > \frac{1}{2}$ both when $\tau_{ij} = 0$ and $\tau_{ij} \neq 0$.

Our Bayesian updating model therefore suggests that in an environment with imperfect information, individuals make judgements using a set of prior beliefs that are updated using Bayes' rule when new signals are received. As we demonstrated above, under both biased and unbiased updating, there are several scenarios that could result in lower agreement level when signals are

less-/non-mainstream. Case (4) is perhaps the most reasonable scenario to consider since it is hard to imagine that the perception of bias does not affect the perception regarding the precision of signals (Case 2) or vice versa (Case 3).

It is important to note however that Bayes' Theorem does not say anything about how one should *interpret* the information received in the process of updating priors. Therefore, Bayes' rule does not preclude the influence of bias in updating priors (Gerber & Green 1999, Bartels 2002, Bullock 2009, MacCoun & Paletz 2009, Fryer et al. 2017). The perceived differences in the mean and the precision of signals is unbiased if mainstream sources are *objectively* more credible than less-/non-mainstream sources. Unbiased Bayesian updating requires the interpretation of signals to be independent from one's prior beliefs (Fischle 2000, Taber and Lodge 2006, Bullock 2009).

Alternatively, the perceived differences in the mean and the precision of signals could be driven by ideological bias. Mainstream sources could be perceived as more credible not based on objective evaluation unrelated to priors, but rather based on the fact that mainstream sources are more likely to confirm a mainstream economist's views. There exists extensive evidence that suggests individuals tend to agree more with findings or views that are more (less) likely to confirm (disconfirm) their beliefs (e.g. McCoun 1998, Gerber and Green 1999, Bartels 2002, Bullock 2009, Hart et al. 2009, MacCoun and Paletz 2009, Fryer et al. 2017, and others). This is broadly referred to as confirmation bias. Beliefs that one seeks to confirm have different natures and could be formed by ingrained, ideological or emotionally charged views. If beliefs an individual is trying to confirm or validate are shaped by ideological views, we are dealing with what is often referred to as ideological bias.¹⁸

For example, MacCoun and Paletz (2009) conduct an experiment to examine how ordinary citizens evaluate hypothetical research findings on controversial topics. They find that when findings challenge people's prior beliefs, they are more skeptical of the findings. Their results also suggest that "citizens, especially those holding conservative beliefs, tended to attribute studies with liberal findings to the liberalism of the researcher, but citizens were less likely to attribute

¹⁸ As Eagleton (1991) suggests, the term "ideology" has been used in different ways by different social scientists. This is partly due to the complex and multi-dimensional nature of the concept, which does not yield itself very easily to a neat definition. We therefore see little advantage in providing a narrow definition by singling out one trait among a complex of traits. It is the complex itself that we are interested in, and in this paper, we examine a clearly-defined manifestation of this complex notion.

conservative findings to the conservatism of the researcher.” They interpret this as effects of “partisanship and ideology”.

In the context of our experiment, the pair of sources that are randomly assigned to each statement (i.e. treatment 1) are carefully selected so that there is an ideological distance between them. This ideological distance is of different lengths, and takes different forms, including mainstream versus non-mainstream, left versus right, liberal versus conservative, critical versus non-critical of neoclassical economics, economist vs philosopher/sociologist, etc.¹⁹ Therefore, in the presence of ideological bias, this ideological distance could induce an ideological reaction. Given the fact that the majority of our participants are mainstream economists, this ideological reaction is likely to manifest itself through *interpreting* less-/non-mainstream sources as being less credible signals. The distinction with unbiased updating therefore is that this *interpretation* is not made objectively and independent from prior beliefs, but rather based on the fact that mainstream (less-/non-mainstream) sources are more likely to confirm (disconfirm) mainstream views (Bartels 2002, Taber and Lodge 2006, Gentzkow and Shapiro 2006).

Altogether, our simple framework suggests that both unbiased and biased updating could lead to a higher agreement level when the source is mainstream. However, we use this framework to derive three empirically-testable implications that we can then assess using our data to examine the validity of these two different explanations.

4.1. Implication 1: Differences in Confidence Level

As our Bayesian model suggests, the reduction in agreement level associated with our treatments is the result of perceived differences in the mean and the precision of different types of signals. Therefore, as Equation (4) suggests, in the absence of bias in the form of $\tau_{ij} = 0$, higher signal precision attributed to mainstream sources (i.e. $h_{ij(m)} > h_{ij(n)}$) increases the precision of posterior beliefs when the source is mainstream (i.e. $k_{1ij(m)} > k_{1ij(n)}$).

In contrast, ideological/authority bias could affect the level of agreement with little or no impact on confidence level. For example, an ideologically-biased individual is likely to put the same level of confidence in accepting a story if it comes from the Fox News and in rejecting the same story if it comes from the New York Times. This is the nature of ideological bias that could

¹⁹ We should emphasize that this dichotomization of different ideologies is intended to give the reader a better idea about what we refer to as ideological distance. Otherwise, we consider ideological distance as differences along a multi-dimensional continuum rather than a one-dimensional dichotomy which is what is examined in other studies such as Gordon and Dahl (2013) and Jelveh et al. (2018).

affects one's judgement often without casting doubt on his confidence with his judgement. Consistent with this idea, and in the context of our Bayesian model, the presence of ideological bias in the form of $\tau_{ij} \neq 0$ implies that:

$$k_{1ij(m)} > k_{1ij(n)} \text{ IFF } \frac{m_{ij(m)} - m_{ij(n)}}{m_{ij(m)} m_{ij(n)}} - (\tau_{ij(n)} - \tau_{ij(m)}) > 0. \quad (9)$$

On the one hand, higher precision attributed to mainstream sources means $m_{ij(m)} - m_{ij(n)} > 0$. On the other hand, since $\tau_{ij(n)} > 0$ and $\tau_{ij(m)} < 0$, we will have $\tau_{ij(n)} - \tau_{ij(m)} > 0$. This suggests that while in the absence of bias we should get $k_{1ij(m)} > k_{1ij(n)}$, in the presence of bias there are scenarios under which $k_{1ij(m)} = k_{1ij(n)}$.

In [Section 6.2.1](#), we use participants' reported confidence level with their evaluations as a proxy for the precision of their posterior beliefs. We then test whether the effect of attributed sources on the precision of beliefs is consistent with the implication of biased versus unbiased updating derived above.

4.2. Implication 2: Differences Between Experts and Non-Experts

It is reasonable to assume that experts have more precise priors in their area of expertise due to their higher level of knowledge in that area (i.e. $k_{ij}^{expert} > k_{ij}^{non-expert}$). Therefore, as Equation (3) suggests, in the absence of bias in the form of $\tau_{ij} = 0$, experts will put more weight on own prior beliefs and less weight on the signal compared to less-/non-experts. Therefore, under unbiased updating, changing the signal from mainstream to less-/non-mainstream should have a smaller effect on experts relative to non-experts.

In contrast, under biased updating, treatment 1 could produce a similar effect among experts and non-experts. Experts' stronger views on a subject could create a stronger bias in favour of mainstream sources. In other words, given that $\tau_{ij(m)} < 0$, we will have $\tau_{ij(m)}^{expert} < \tau_{ij(m)}^{non-expert} < 0$, which in turn implies that $h_{ij(m)}^{experts} > h_{ij(m)}^{non-experts}$.²⁰ Therefore, in the presence of bias, the combination of more precise priors ($k_{ij}^{expert} > k_{ij}^{non-expert}$) and larger mainstream biases among experts ($h_{ij(m)}^{experts} > h_{ij(m)}^{non-experts}$) could produce a similar treatment effect for experts and non-experts. We test these two implications in [Section 6.2.2](#), by identifying experts

²⁰ Assuming $m_{ij}^{experts} = m_{ij}^{non-experts}$.

and non-experts for each statement based on their area of research and comparing their estimated treatment effects.

4.3. Implication 3: Differences by Political Orientation

Political views are one of the important driving forces behind ideological bias. In the context of our model, these political views will be a part of v_{ij} in Equation (1). Under unbiased updating, these political views will be uncorrelated with one's interpretation regarding the validity of different signals, which means we will have $cov(\epsilon_{ij}, v_{ij}) = \tau_{ij} = 0$. As a result, treatment 1 should have a similar effect among groups with different political orientations.

In contrast, under biased updating, political views will be correlated with how one interprets the credibility of a source. The closer the political orientation of a source to that of the economist, the higher will be the credibility attributed to the source. Our less-/non-mainstream sources often represent views or ideologies that are politically to the left of mainstream sources. Therefore, conditional on the mainstream source, the effect of treatment 1 is expected to increase (i.e. the reduction in agreement level will be larger) as participants' political orientation gets closer to the far right of the distribution and further away from the less-/non-mainstream source.

In our model, this translates into an increase in $cov(\epsilon_{ij}, v_{ij} | s_{ij(n)}) = \tau_{ij(n)} > 0$ as we move to the far right of the distribution of the political orientation. This is because as we move to the far right, the level of agreement with our statements decreases (given their critical and "leftist" nature) through decrease in v_{ij} . At the same time, moving to the far right will also result in lower credibility attributed to less-/non-mainstream sources through smaller ϵ_{ij} . As a result, moving to the far right and increase in $\tau_{ij(n)}$ will increase $h_{ij(m)} - h_{ij(n)}$. It is therefore easy to see from both Equations (7) and (8) that this will lead to an increase in the value of N_1 in all the four cases discussed. This suggests that under biased updating, moving to the far right is expected to increase the probability that the economist will agree more with a statement when it is attributed to a mainstream source versus a less-/non-mainstream source. In other words, as we move to the far right the reduction in agreement level resulting from treatment 1 is expected to increase.²¹

²¹ Assuming that moving to the far right will also increase the credibility attributed to mainstream sources (i.e. increase $\tau_{ij(m)}$ in absolute value) will induce an even larger decrease in agreement resulting from treatment 1.

4.4. Authority Bias

It could be argued that statements with less-/non-mainstream sources could receive lower agreement not necessarily due to ideological bias but (partly) due to the influence of authority bias. Authority bias is the tendency to assign more credibility to views that are attributed to an authority figure (Milgram 1963). Under authority-biased Bayesian updating, an economist might consider mainstream sources as authority figures whose statements are more credible, which is a view that is not necessarily ideologically driven.²² Our implementation of a second treatment where we remove the source attribution altogether allows us to examine whether one or both biases are present. In the context of our model, and assuming that removing the source will remove the signal altogether, this translates into

$$P_2 = \Phi \left(\frac{\omega_{ij(m)} [\theta_{j(m)} - \theta_j]}{\frac{\omega_{ij(m)}^2}{h_{ij(m)}}} \right). \quad (10)$$

Therefore, as long as $\theta_{j(m)} > \theta_j$, we will have $P_2 > \frac{1}{2}$. In other words, as long as the mainstream signal lends more validity to the statement relative to the economist's prior, removing the source will reduce the level of agreement. In addition, comparing Equations (10) and (6) suggests that if $\theta_{j(m)} \geq \theta_j$, then $P_1 < P_2$. In other words, if the economist's interpretation of the validity of the statement when it is attributed to a less-/non-mainstream source is also higher than his prior, then changing the source to less-/non-mainstream will result in a smaller reduction in agreement level compared to removing the source.²³

Altogether, as our simple framework suggests, both unbiased and biased updating could potentially lead to a higher agreement level when the source is mainstream. However, in [section 6.2.](#), we use the three empirically-testable implications we derived from our model to examine the

²² From a theoretical perspective, the main distinction we draw between ideological bias and authority bias is that under ideological bias individuals are more likely to agree with views that confirm their own ideological views. However, under authority bias individuals are more likely to agree with views that are attributed to an authority figure, while they might not have any particular ideological views on the subject. For example, authority bias will result in higher admiration for a poem if it is attributed to a famous poet, but lower admiration if it is attributed to a school teacher, with neither of the two assessments influenced by ideology.

²³ It is also plausible to imagine a scenario where removing the source attribution still contains a signal that could affect agreement level. As mentioned before, most of our statements challenge certain aspects of mainstream views. Therefore, even in the absence of a source attribution, a participant is likely to attribute the statement to someone who is considered critical of mainstream economics. This could induce an ideological bias effect similar to the one discussed above, which will result in reduction in agreement level as we demonstrated before.

validity of biased versus unbiased updating. If we find evidence of biased updating, then our implementation of a second treatment, where we remove the source attribution altogether, could allow us to examine whether one or both biases are present. For example, finding an effect *only* for treatment 2 would be consistent with the existence of authority bias but not ideological bias. Similarly, finding an effect *only* for treatment 1 would be in line with the existence of ideological bias but not authority bias. Finding an effect for both treatments would be consistent with the existence of both ideological bias and authority bias.

5. Data

The target population for this study were economists from 19 different countries.²⁴ We used Economics Departments, Institutes and Research Centers in the World (EDIRC) website, which is provided by the Research Division of the Federal Reserve Bank of St. Louis, to identify economic institutions (economics departments, government agencies, independent research institutions, and think tanks) in each target country. We then used the website of each institution (provided on the EDIRC website) to manually extract the email addresses of economists in each institution. The extracted email addresses were then used to send out invitations and reminders to ask economists to participate in the survey. The survey was conducted between October 2017 and April 2018. While the survey's exact opening and closing dates were different for some countries, the survey was open in each country for approximately two months.

In many cases during email extraction, especially in the case of multidisciplinary departments (e.g. school of business, management, and economics), research institutions, and government agencies, it was not clear from the institution's website which listed faculty members or researchers were economists and which ones held a degree from other disciplines. In these cases, we asked our team of research assistants to extract all listed email addresses. Our rationale was that sending email invitations to some non-economists was clearly better than risking to exclude some economists, especially since this exclusion could be systemically related to the type of institution and lead to sample selection. We made sure however that non-economists who received

²⁴ These countries include Australia, Austria, Brazil, Canada, Denmark, Finland, France, Germany, Ireland, Italy, Japan, Netherlands, New Zealand, Norway, South Africa, Sweden, Switzerland, the UK, and the US. The entire (English) survey was translated into French, Italian, Japanese and Brazilian Portuguese to allow participants from corresponding countries to complete the survey in their own native language if they choose to.

the survey invitation were self-filtered out by making it clear in our email invitation as well as on the first page of the survey that the target population of the survey are economists.²⁵

As a result, we are not able to provide a reasonable and reliable estimate of the participation rate in our survey since that would require the total number of economists in the target population, which is considerably smaller than the total number of email addresses we extracted online, for the reason discussed above. In addition, this calculation is further complicated by the fact that upon sending email invitations we received a considerable number of auto-replies from people who had left their institution, were on sabbatical, parental, or sick leave, or temporarily had no access to their email. With these in mind, a very rough estimate of the participation rate in our survey is around 15%. Although we cannot measure a reliable participation rate for our survey for the reasons discussed above, our summary statistics (Table A1 in our [online appendix](#)) suggest that we have a very diverse group of economists in our final sample. We have also reported the distribution of responses by institution of affiliation in the US, Canada, and the UK in figures A3 to A5 in our [online appendix](#) as examples to show that participants in our survey come from a very diverse group of institutions in each country and are not limited to certain types of institutions.

Participants in our survey were required to complete each page in order to proceed to the next page. As a result, they could not skip evaluating some statements. However, participation in the survey was entirely voluntary and participants could choose to withdraw at any point during the survey, without providing any reason, by simply closing the window or quitting the browser. Participants were assured that any responses collected up until the point of withdrawal will not be included in the study. For this reason, we are not allowed, by the terms of our ethics approval, to use data collected from people who did not complete the entire survey. As a result, we have restricted our sample to participants who completed the entire survey.²⁶ Our final sample includes 2,425 economists from 19 different countries. We run several tests to ensure that our focus on participants who completed the entire survey does not introduce sample selection bias in our results and we find no evidence of such a bias. See Section 1 in our [online appendix](#) for more details.

²⁵ As expected, we received many emails from faculty members who were not economists (historians, statisticians, sociologists, political scientists, engineers, etc.) asking us to remove them from the email list.

²⁶ A total of 3,288 economists participated in our survey. There were 454 participants who quit the survey at the very beginning (in the questionnaire section where they were asked to provide background information). Another 409 people withdrew from the survey at some point after they started evaluating the statements. See Table A3 in our [online appendix](#) for more details.

The primary dependent variable in our analysis is the reported agreement level with each statement. In our baseline analysis, we estimate linear regression models in which the agreement variable is coded as 1 for “strongly disagree”, 2 for “disagree”, 3 for “neutral”, 4 for “agree”, and 5 for “strongly agree”. We also estimate ordered logit models for robustness check. The agreement level of participant i with statement j is represented by the variable y_{ij} and is modeled as:

$$y_{ij} = \gamma_1 S1_{ij} + \gamma_2 S2_{ij} + X_i \beta + \epsilon_{ij}, \quad (11)$$

where $S1_{ij}$ and $S2_{ij}$ are indicators that are equal to one if for statement j participant i received a less-/none-mainstream source, or no source, respectively. The estimated coefficients of interest are γ_1 and γ_2 and measure average difference in agreement level between those who randomly received a less-/non-mainstream source or no source, respectively, compared to those who received a mainstream source. We also include several individual-level control variables (X_i) in some of our specifications.²⁷ However, if our randomization is carried out properly, including these control variables should not affect our results (and as reported later on we find that they don't).

6. Results

6.1. Main Findings

Table 1 displays the results from estimating linear models. Column (1) uses a simplified model with no additional control variables, while columns (2) to (4) add personal and job characteristics as well as individual fixed effects.²⁸ We find clear evidence that changing source attributions from mainstream to less-/non-mainstream significantly reduces the agreement level by 0.26 points. This is around one-fourth of a standard deviation or a 7.3 percent reduction in an average agreement level of 3.6 in our control group. Our results also suggest that removing mainstream sources (i.e. providing no source) also significantly reduces the agreement level by 0.41 points (an 11.3 percent reduction which is equal to 35% of a standard deviation).

As estimates reported in Columns (2), (3), and (4) suggest, controlling for different individual characteristics and individual fixed effects does not change our results, which provides

²⁷ Our primary control variables include: gender, PhD completion cohort (15 categories), Current Status (8 categories), Country (19 categories), Research Area (18 categories). Additional control variables used in some specifications include age cohort (13 categories), country/region of birth (17 categories), English proficiency (5 categories), department of affiliation (8 categories), country/region where PhD was completed (16 categories). See Table A1 in the [online appendix](#) for more detail on different categories.

²⁸ Refer to Table A2 in our [online appendix](#) for estimated coefficients of our control variables.

further support that our randomization protocol was implemented properly. In addition, results from our specification with individual fixed effects suggests that our estimate of treatment 1 is unlikely to suffer from sample selection bias due to non-random attrition across treatment groups. Finally, as the results reported in Table A11 in our [online appendix](#) suggest, estimating the same specifications while clustering the standard errors at the individual level does not have any appreciable impact on our results. More specifically, clustering has virtually no impact on standard errors for treatment 1, while slightly increasing standard errors for treatment 2.²⁹ However, the t-statistics for estimates of treatment 2 are so large (around 27 before clustering) that this slight increase has virtually no impact on the outcomes of our hypothesis testing.

While OLS estimates are perhaps easier to summarize and report, given the discrete ordered nature of our dependent variable, a more appropriate model to use in this context is an ordered logit model. Another advantage of using ordered logit is that it allows us to examine whether our treatments have heterogeneous effects. Changing a source from mainstream to less-/non-mainstream might have opposite effects on those who (strongly) agree/(strongly) disagree with a statement. For example, on the one hand, those who strongly agree with a statement that is critical of mainstream economics are more likely to be less-/non-mainstream, and therefore treatment 1 might induce higher agreement among them. On the other hand, those who strongly disagree with the same statement are more likely to be mainstream, and therefore treatment 1 might induce higher disagreement among them.³⁰ These potential heterogeneous effects will not be captured by our OLS model.

Table 2 reports the estimates from our ordered logit model. Overall, we find results similar to those reported in Table 1 using OLS.³¹ We find that changing sources from mainstream to less-/non-mainstream, or to no source, significantly increases (decreases) the probability of disagreement (agreement) with statements. More specifically, we find that providing a less-/non-mainstream source on average increases the probability of “strong disagreement” by 2.2

²⁹ These changes are consistent with Abadie et al (2017) who suggest clustering is relevant “when clusters of units, rather than units, are assigned to a treatment.”

³⁰ This potential heterogeneous effect might also be responsible for the larger effect estimated for treatment 2. For example, providing a less-/non-mainstream source might generate biases in two directions, with the positive effect on agreement partially canceling out the negative effect. Removing the source however might only introduce a negative effect. As a result, our treatment 2 will have a larger effect than treatment 1. Ordered logit model will allow us to examine this possibility as well.

³¹ We also estimate multinomial logit models for robustness check. Results from these models are very similar to those from ordered logit models.

percentage points or 44 percent, increases the probability of “disagreement” by 5 percentage points or 30 percent, increases the probability of reporting “neutral” by 2.1 percentage points or 12.6 percent, reduces the probability of “agreement” by 3.6 percentage points or 9 percent, and reduces the probability of “strong agreement” by 5.7 percentage points or 27 percent. This suggests that regardless of the extent to which participants agree or disagree with a statement, changing the source to less-/non-mainstream significantly decreases (increases) their agreement (disagreement) level. In addition, while the effect of treatment 1 on increasing the probability of (strong) disagreement is larger relative to its impact on reducing the probability of (strong) agreement, we find no evidence of opposite treatment effects.

Moreover, similar to our linear estimates, we also find larger effects in the same direction when no sources are provided for the statements. Also, in line with our linear estimates, we get almost identical results when we include control variables in our specification. Several other robustness checks performed (see Section 2 in our [online appendix](#)) also fully confirm the robustness of our results.

6.2. Ideological/Authority Bias or Unbiased Bayesian Updating?

The important issue that we will address at this point is determining the underlying sources behind these substantially large and robust effects. We organize the remainder of this section to examine the potential role of unbiased updating versus ideologically/authority-biased updating in driving our main results.

As we discussed in [Section 4](#), the validity of unbiased updating hinges on the assumption that the assessment of sources’ credibility is independent from prior beliefs and that mainstream sources systematically provide *objectively* more credible signals regarding the validity of the statements compared to less-/non-mainstream sources. It is important to note however that our (fictitious) less-/non-mainstream sources are not random people but rather prominent scholars in the field with views that put them at different distances, sometimes relatively close and sometimes rather far, to mainstream economics. Therefore, one main problem with unbiased updating as a potential explanation is that there are no *objective* measures that could be used to assess the credibility of these sources. Any claims of systematic differences between these sources in terms of credibility is inevitably based on subjective metrics that correlate with where one stands relative to mainstream views and its academic norms. It is exactly for this reason that traditional norms of

modern science suggest that any serious evaluation of an argument should be based on the content of the argument as opposed to the source attributed to it (Merton 1973, McCoun 1998).

In fact, economists in our sample strongly agree with this view. More specifically, as part of the questionnaire that appears *at the end* of the survey, we ask participants to express their own views regarding several issues including how they believe “a claim or argument should be rejected?” A strong majority of participants (around 82%) report that “a claim or argument should be rejected only on the basis of the substance of the argument itself.” Around 18% of participants report that “a claim or argument should be rejected based on what we know about the views of the author or the person presenting the argument as well as the substance of the argument.” There exists only a tiny minority (around 0.5%) who report “a claim or argument should be rejected based on what we know about the views of the author or the person presenting the argument.”³² Nevertheless, to provide more rigorous empirical evidence, we use the three empirically-testable implications we derived from our model in [Section 4](#) to examine the validity of biased versus unbiased updating.

6.2.1. Test 1: Differences in Confidence Level

As we discussed in [Section 4.1](#), our model suggests that under unbiased updating the precision of the posterior beliefs should be higher when the source is mainstream. However, this does not necessarily hold under biased updating as we showed in [Section 4.1](#). To test the validity of biased versus unbiased updating which underlie these two predictions, we use participants’ reported confidence level with their evaluations (on a scale from 1 to 5, with 1 being least confident and 5 being most confident) as a proxy for the precision of their posterior beliefs.³³ We then estimate models similar to those reported in Table 1 but instead use as our dependent variable the level of confidence reported for one’s evaluation of a statement.

As the results reported in Table 3 suggest, the estimated coefficients are quantitatively very small and statistically insignificant.³⁴ This suggests that in contrast (consistent) with unbiased

³² We also examine differences in our estimated treatment effects between those who claim statements should be evaluated only based on their content versus those who claim both the content and the views of the author matter. We find almost identical estimates for treatment 1 (-0.22 versus -0.20) and treatment 2 (-0.34 versus -0.35) for both groups. This suggests that our results are not driven by one of these two group.

³³ See Figure A6 in our [online appendix](#) for probability of different confidence levels by statement.

³⁴ Results from ordered logit (not reported here) also suggest that our treatments have no impact on confidence level. Our estimated differences in predicted probabilities for all five categories are small and statistically insignificant.

(biased) updating, altering or removing source attributions does not influence participants' confidence in their evaluations (i.e. precision of their posterior beliefs).

6.2.2. Test 2: Differences Between Experts and Non-Experts

The second implication, derived from our model in [Section 4.2.](#), suggests that, under unbiased updating, changing the source from mainstream to less-/non-mainstream is expected to have a smaller effect on experts relative to non-experts. In contrast, under biased updating, treatment 1 could produce a similar effect among experts and non-experts. We create an indicator that is equal to 1 if a participant's reported area of research is more likely to be relevant to the area of an evaluated statement and zero otherwise (see Table A10 in our [online appendix](#) for more details). We then estimate linear models similar to those reported in Table 1 where we allow our treatment effects to vary by expertise.

Results from this model are reported in Table 4. First, we find no difference in average agreement level between experts and non-experts in the control group. This suggests that one's expertise in the subject area does not affect one's evaluation of a statement when sources are mainstream. Second, we find that both estimated treatments effects are similar among those categorized as experts and non-experts. These findings are therefore inconsistent with unbiased updating which suggests higher level of expertise is expected to result in smaller treatment effects. However, these results could be consistent with biased updating which as we showed in [Section 4.2.](#) could result in similar treatment effects between experts and non-experts.

6.2.3. Test 3: Differences by Political Orientation

As our third and perhaps most important test, we examine whether our estimated treatment effects vary across different groups with different political orientations. If the reduction in agreement level associated with changes in sources is based on *objective* differences in credibility of the sources, by definition this objective difference should not depend on one's political views, especially after controlling for other observed characteristics. As a result, as we showed in [Section 4.3.](#), our estimates should not vary systematically by political orientation under unbiased updating. In fact, if anything, those on the right should be less affected by changing the source attributions since they are significantly more likely to report that a statement should be evaluated based on its content only. More specifically, among those at the far right, 86.7 percent of participants report that in evaluating a statement only its content matters, while 13.3 percent report that both content

and author matter. In contrast, among those at the far left, these numbers are 73.8 percent and 25.1 percent, respectively.

As we showed in [Section 4.3.](#), evidence suggesting that the effect of treatment 1 varies systematically by political orientation is consistent with ideological bias. More specifically, our less-/non-mainstream sources often represent views or ideologies that are (politically) to the left of mainstream sources. Therefore, if our results are driven by ideological bias, reduction in agreement level should be larger among those more to the right of the political spectrum since altering the sources creates a larger contrast with their prior beliefs which will in turn induce a larger ideological reaction among this group.

We estimate linear models similar to Equation (11) above where we allow the effect of each treatment to vary by political orientation. Political orientation is reported by participants on a scale from -10 (far left) to 10 (far right). We use the reported values to group people into 5 categories.³⁵ Results from this model are reported in the first three columns of Table 5 and suggest that, for those on the far left, altering the source only reduces the average agreement level by 0.05 points of a standard deviation, which is less than one-fourth of the overall effect we reported in Table 1 (0.22 points). Moving from the far left to the far right of the political orientation distribution, however, consistently increases this effect, with the effect of altering the source being 6 times larger at the far right compared to the far left (-0.31 versus -0.050, respectively). We reject, at 0.1% confidence level, both the null hypothesis that the effect at the far left is equal to the effect at the far right, and that the effects are equal across all five groups (F-statistics are 39.7 and 14.57, respectively).

As estimates reported in columns (4) to (6) suggest, our results remain the same if we categorize people based on quintiles of the political orientation distribution. Moreover, including additional control variables do not affect our results.³⁶ In line with our previous tests, these findings are consistent (inconsistent) with the existence of ideological bias (unbiased updating).

It is reasonable to argue, however, that the self-reported measure of political orientation used to categorize people depends on political environments and contexts that could vary significantly from one country to another. For example, someone who is considered a centrist or centre-right in the UK could perhaps be categorized as left in the US. This could complicate the

³⁵ Far left = [-10 -7], Left = [-6 -2], Centre = [-1 1], Right = [2 6], Far Right = [7 10].

³⁶ These results are not reported here, but are available upon request.

interpretation of our results. To address this issue, we use participants' answers to a series of questions at the end of our survey that are designed to identify their political typology.³⁷ More specifically, we regress our self-reported political orientation measure on a series of indicators created based on answers to these questions. We then use predicted values from this regression to categorize people into five groups based on its distribution quintiles. Results reported in Columns (7) to (9) of Table 4 are based on this alternative categorization and remain similar to those reported in other columns confirming the same pattern.

Finally, another important pattern worth highlighting, reported in Table 5, is that, contrary to our estimated effects of treatment 1, our estimated effects of treatment 2 do not follow any meaningful pattern of change by political orientation. In fact, as reported in Column (3), we fail to reject the equality of estimated effects for treatment 2 across different categories of political orientation. These differences are consistent with the theoretical distinction we drew between ideological bias and authority bias in [Section 4.4.](#). To reiterate, under ideological bias, individuals' interpretation of signals and their level of agreement is influenced by their ideological views. However, under authority bias, it is the presence/absence of an authority figure that affects individuals' interpretation of signals and agreement level, while they might not have any particular ideological views on the subject.

As we discussed in [Section 4.4.](#), given that we find a robust and significant estimated effect for both treatment 1 and treatment 2, up until this point, we could not rule out that both ideological bias and authority bias contribute to each effect. However, our finding that our estimated effect of treatment 2 does not follow the same meaningful pattern by political orientation as treatment 1 is consistent with the theoretical distinction between ideological and authority bias. This therefore suggests that there are important differences in underlying forces driving our estimated effects of treatment 1 and treatment 2, with the former (latter) being more likely to be driven by ideological (authority) bias.

6.2.4. Addressing Two Potential Caveats

In this section, we address two potential caveats regarding the validity of ideological/authority bias as an explanation for our results. First, it can be argued that, given the low stake nature of our survey, economists did not have the incentive to exert much effort and read

³⁷ Participants were asked to read a series of binary statements and for each pair pick the one that comes closest to their view. See Table A9 in our [online appendix](#) for a list of these statements.

each statement carefully. Therefore, when the attributed source for a statement was a prominent mainstream economist who they recognized and trusted as a scholar, they glossed over the statement and relied on the source for their evaluation. As a result, statements attributed to mainstream sources received a higher level of agreement.

If this hypothesis is valid, then one of its implications is that participants in our control group should have spent less time completing the survey compared to those in the two treatment groups. However, our estimates reported in Table A5 in our [online appendix](#) suggest that there are no differences in average survey completion time between control group and treatment 1. We find, however, that those randomized into treatment group 2 on average took less time to complete the survey, but the estimated difference is very small (less than a minute) and is to be expected since people in this group have less text to read, given that there are no source attributions provided.

Consistent with these results, our estimates in Table A6 in our [online appendix](#) suggest that restricting our sample to individuals with different survey completion times (a potential proxy for different levels of effort exerted to read the statements) has also no impact on our results. Finally, this hypothesis fails to explain the findings from our three tests, especially why our estimated treatment effects vary significantly by political orientation ([Test 3](#)).

The second potential caveat we would like to address is that we have not established whether our participants have in fact been able to identify the ideological distance between sources, and therefore ideological bias might not be a valid explanation. If this is the case, then one needs to put forward an alternative explanation, other than authority bias, that explains why changing sources systematically reduces the agreement level. We cannot think of a plausible alternative hypothesis that is able to explain all or even most of our results. More importantly, it is hard to imagine that our participants could not identify the ideological distance between sources such as Marx versus Smith, Mill versus Engels, Summers versus Varoufakis, and Keynes versus Arrow. In addition, as we mentioned in [Section 3](#), we also provided information on the sources' discipline, their affiliation, and the title of one of their publications, to further accentuate the ideological differences between the two sources. Finally, significant differences in estimated treatment effects by political orientation does not support the idea that participants were not able to identify the ideological distance between different pairs of sources.

Altogether, our results presented and discussed above are consistent with the existence of ideological and authority biases in views among economists and are in contrast with the

implications of unbiased Bayesian updating. We would like to emphasize that, while some of our results might be considered (at least by some of our readers) as more compelling than others, it is the entirety of the body of evidence we have provided that should be considered when one ponders the validity of these two competing hypotheses (i.e. biased versus unbiased updating), or other potential hypotheses. This is of course consistent with the methodology of modern economics, that the relative superiority of a hypothesis, relative to other alternative hypotheses, is determined by its relative degree of success in explaining the observed patterns. Therefore, to suggest that ideological/authority bias is not driving our results, one needs to posit an alternative hypothesis that is on average more successful in explaining all the observed patterns, and not only a cherry-picked few.

6.3. Statistical Power and Reproducibility of Findings

A growing body of literature is raising concern regarding the dominant practice in empirical studies to restrict attention to type-I-error and highlights the importance of statistical power as a critical parameter in evaluating the scientific value of empirical findings (see Maniadis and Tufano (2017) for a review). This is especially important since there is growing evidence that suggests empirical findings in economics, and several other disciplines, are significantly underpowered. For example, assessing more than 6,700 empirical studies in economics, Ioannidis et al. (2017) find that “half of the areas of economics research assessed have nearly 90% of their results under-powered. The median statistical power is 18%, or less.” As Maniadis and Tufano (2017) suggest, “an important implication of the overall inadequate power of empirical research in economics is that a sizable majority of its studies have less than 50% probability of detecting the phenomenon under investigation.”

Following List et al. (2011) and Maniadis et al. (2017), we calculate the optimal sample size for our control and treatment groups that allows us to calculate a minimum economically relevant treatment effect at what is considered a reasonable significance level ($\alpha=0.05$) and statistical power ($1-\beta=0.80$) in the literature (e.g. List et al. 2011, Ioannidis et al. 2017). We find that *for a given statement*, in order to find a treatment effect equal to 15% of a standard deviation, which we consider an economically relevant treatment effect in our study, we need to have approximately 800 participants in each group, which is the number of participants we actually have. This suggests that, given our sample size, the treatment effects we have estimated for each individual statement meet a high standard in terms of statistical power and degree of

reproducibility. Doing the same exercise for our analysis of all statements combined, which is most of our analysis, suggests that our sample size of 36,375 individual-statements allows us to find a treatment effect as small as 6% of a standard deviation at $\alpha=0.01$ and $1-\beta=0.99$.

Another important and related question raised by Maniadis et al. (2017) is that, “given publication of a newly discovered finding, how much confidence should we have that it is true?” Maniadis et al. (2014) propose a measure of post-study probability (PSP) that allows one to measure “the probability that a declaration of research finding, made upon reaching statistical significance, is true.” As Maniadis et al. (2014) show, the value of PSP depends on significance level α , statistical power β , and priors about the true probability of association between two phenomena, π . In the context of our study, π translates into priors regarding the probability that the existence of ideological bias among economists is a true phenomenon. We consider a low, a medium, and a high prior to measure how much these priors should change in light of our results. We use $\beta=0.99$ and $\alpha=0.01$ which we showed above are reasonable numbers given our sample size and estimated effects when all statements are combined.

We find that, if one holds a prior that there is only a 10% probability that the existence of ideological bias among economists is a true phenomenon ($\pi=0.1$), then the probability that our estimated treatment effects are true is $\text{PSP}=91.6\%$. This probability is 99% and 99.8%, for $\pi=0.5$ and $\pi=0.9$, respectively. This suggests significant updating in those prior beliefs in light of our results.³⁸

6.4. Heterogeneity Analysis

In this section, we examine how our estimated treatments vary by statement and also by different characteristics including gender, country of residence, country where PhD was completed, undergraduate major, and main research area. It is interesting and important to understand how the biases we have found in our analysis vary across different groups. This could help to shed more light on some of the factors underlying ideological/authority bias. As it is discussed in more detail in the next section, we find evidence of significant and systematic heterogeneity in our estimated treatment effects. Consistent with our previous results, we cannot square these findings with unbiased Bayesian updating since we cannot think of compelling reasons to explain why unbiased updating will lead to systematically different treatment effects

³⁸ Even for a prior of 1%, the $\text{PSP} = 50\%$.

for people from a certain gender or from certain countries. However, as we discuss in more detail in the following section, these differences by personal characteristics remain consistent with the existence of ideological/authority bias among economists.

We continue to use OLS to estimate treatment effects since its estimates are easier to summarize and present, and since they are similar to estimates from ordered logit models as discussed before.

6.4.1. Heterogeneity by Statement

First, we investigate the effect of our two treatments on agreement level separately for each statement. These results are summarized in Figure 1. Consistent with our overall findings, we find that for all but three statements, changing source attributions to a less-/non-mainstream source significantly reduces the agreement level. The estimated reductions range from around one-tenth of a standard deviation to around half of a standard deviation. Interestingly, we find that the largest reduction in agreement level for treatment 1 occurs for Statement 6, which is arguably the statement that is most critical of mainstream economics and its methods, and also brings up the issue of ideological bias in mainstream economics.³⁹ This again is consistent with ideological bias where views that are more likely to disconfirm previously held beliefs are more strongly discounted.

Regarding the three statements with no reduction in agreement level (i.e. statements 1, 3, and 7), one potential explanation is that the ideological distance between the sources for each statement (i.e. the degree to which they confirm/disconfirm prior beliefs) is not large enough to induce ideological bias. Taking a closer look at the sources for each statement seems to suggest that this is indeed a plausible explanation. The sources for these statements are Dani Rodrick vs. Paul Krugman, Hayek vs. Freud, and Irving Fisher vs. Kenneth Galbraith. Interestingly and consistent with authority bias, we find that, for the same three statements, removing the source

³⁹ The statement reads: “Economic discourse of any sort - verbal, mathematical, econometric-is rhetoric; that is, an effort to persuade. None of these discursive forms should necessarily be privileged over the others unless it is agreed by the community of scholars to be more compelling. Only when economists move away from the pursuit of universal knowledge of 'the economy' and towards an acceptance of the necessity of vision and the historical and spatial contingency of knowledge will the concern over ideological 'bias' begin to fade. Such a turn would have important implications for economic method as well, as knowledge claims would increasingly find support, not in models of constrained optimization, but with such techniques as case studies and historical analyses of social institutions and politics. Increasing reliance of economics on mathematics and statistics has not freed the discipline from ideological bias, it has simply made it easier to disregard.”

attribution (i.e. treatment 2) significantly reduces the agreement level, highlighting again the difference in driving forces behind the estimated effects of treatment 1 and treatment 2.

Results displayed in the right panel of Figure 1 suggest that removing the source attributions significantly reduces the agreement level for all 15 statements. Similar to our results reported in Table 1, the estimated effects of treatment 2 are larger than treatment 1 in almost all of the statements. Results reported in Figure 1 also suggest that it is not just extreme differences in views (e.g. Smith vs. Marx) that invoke ideological bias among economists. Even smaller ideological differences (e.g. Deaton vs. Piketty or Sen vs. Sandel) seem to invoke strong ideological reactions by economists.

6.4.2. Heterogeneity by Gender

Next, we examine gender differences in our estimated treatment effects. These results are reported in Table 6 and suggest that on average female economists who are randomized into control group agree less with our statements compared to their male counterparts in the control group. The estimated difference in agreement level is around 6 percent of a standard deviation. In addition, we find that the estimated ideological bias is 42% larger among male economists as compared to their female counterparts (24% of a standard deviation reduction in agreement level versus 14%, respectively), a difference that is statistically significant at the 0.1%. These results are consistent with studies from psychology which suggest that women exhibit less confirmation bias than men (Meyers-Levy 1986, Bar-Tal and Jarymowicz 2010). Gordon and Dahl (2013) also find evidence that suggests that male economists are less cautious in expressing an opinion. This seems to be consistent with stronger ideological bias among male economists found in our results, since ideological bias and assigning higher levels of certainty to our own views usually work hand in hand.

We find, however, that the gender difference in authority bias is much smaller (33% of a standard deviation reduction for males versus 36% for females) and statistically insignificant. In other words, removing mainstream sources seems to affect men and women in similar ways. We find similar results when we estimate gender differences in treatment effects separately for each statement. In 9 out of 15 statements, the estimated ideological bias is larger for men than for women, while the results are more mixed for our estimates of authority bias (see Figure A2 in our [online appendix](#)).

We would like, however, to highlight the estimated gender difference in ideological bias for Statement 5, which involves the issue of gender gap in economics.⁴⁰ Overall, and without considering group assignment, there exists a very large difference in the level of agreement with this statement between male and female economists. More specifically, conditional on observed characteristics, the average agreement level among male economists is 0.78 points lower than female economists, a very large difference that is around 2/3rd of a standard deviation and statistically significant at the 0.1% confidence level. Taking group assignment into account, female economists who randomly receive *Carmen Reinhart* as the statement source (i.e. control group) report an agreement level that is on average 0.73 points higher compared to their male counterparts in the control group. Moreover, while switching the source from *Carmen Reinhart* to *Diane Elson* does not affect the agreement level among female economists (estimated effect is 0.006 points), it significantly decreases the agreement level among male economists by 0.175 points (around 15% of a standard deviation). It seems that when it comes to the important issue of the gender gap in economics, which involves female economists at a personal level, women put aside ideology and merely focus on the content of the statement as opposed to its source.

These results highlight a large divide between male and female economists in their perception and concerns regarding the gender gap in economics.⁴¹ This is of critical importance since the discussion around the gender problem in economics has recently taken the centre stage. During the recent 2019 AEA meeting, and in one of the main panel discussions titled “how can economics solve its gender problem?”, several top female economists talked about their own struggles with the gender problem in economics. In another panel discussion, Ben Bernanke, the current president of the AEA, suggested that the discipline has “unfortunately, a reputation for hostility toward women....”⁴² This is following the appointment of an Ad Hoc Committee by the

⁴⁰ The statement reads: “Unlike most other science and social science disciplines, economics has made little progress in closing its gender gap over the last several decades. Given the field’s prominence in determining public policy, this is a serious issue. Whether explicit or more subtle, intentional or not, the hurdles that women face in economics are very real.” The actual (mainstream) source of the statement is *Carmen Reinhart, Professor of the International Financial System at Harvard Kennedy School and the author of This Time is Different: Eight Centuries of Financial Folly (2011)*. The altered (less-/non-mainstream) source of the statement is *Diane Elson, British Economist and Sociologist, Professor Emerita at the University of Essex, and the author of Male bias in the development process (1995)*.

⁴¹ Gender differences in perception of gender discrimination is not limited to economics and has been documented in other studies and in various contexts (e.g. Fisman and O’neill 2008, Miller and Katz 2018, Raggins et al. 1998).

⁴² Reported in a New York Times article titled “Female Economists Push Their Field Toward a #MeToo Reckoning”, published on January 10, 2019.

Executive Committee of the AEA in April 2018 to explore “issues faced by women [...] to improve the professional climate for women and members of underrepresented groups.”⁴³ AEA also conducted a climate survey recently to “provide more comprehensive information on the extent and nature of these [gender] issues.” It is well-understood that approaching and solving the gender problem in economics first requires a similar understanding of the problem by both men and women. However, our results suggest that there exists a very significant divide between male and female economists in their recognition of the problem.

6.4.3. Heterogeneity by Country of Residence/PhD Completion

Next, we examine how our estimated effects vary by country of residence. These results are reported in Table 7. For both treatments, we find that the estimated effects vary significantly across countries, ranging from around half of a standard deviation to zero. We also reject the null that the estimated effects of treatment 1/treatment 2 are the same across countries at the 0.1% confidence level. More specifically, we find that economists in Austria, Brazil, and Italy exhibit the smallest ideological bias (for Brazil and Austria, the estimated effects are also statistically insignificant). On the other side of the spectrum we find that economists in Ireland, Japan, Australia, and Scandinavia exhibit the largest ideological bias. Economists in countries such as Canada, the UK, France, and the US stand in the middle in terms of the magnitude of the estimated ideological bias. In addition, when we examine the effect of authority bias, these countries maintain their positions in the distribution, although the estimated effects of authority bias remain larger than the estimated effects of ideological bias for most of the countries.

Table 8 reports results that examine heterogeneity by country/region where PhD was completed. We find that economists who completed their PhD in Asia, Canada, Scandinavia, and the US exhibit the strongest ideological bias, ranging from 43% to 24% of a standard deviation. On the opposite end we find that economists with PhD degrees from South America, Africa, Italy, Spain, and Portugal exhibit the smallest ideological bias (statistically insignificant for South America and Africa). These results are somewhat consistent with those reported in Table 7 and suggest that some of the countries where economists exhibit the largest/smallest ideological bias are also those that induce strongest/weakest ideological bias in their PhD students (e.g. Brazil,

⁴³ American Economic Association, Ad Hoc Committee on the Professional Climate in Economics, Interim Report, April 6, 2018.

Italy, Scandinavia). In addition, we find that our estimated effects of authority bias, while larger in size, largely follow the same patterns as our estimates of ideological bias.

6.4.4. Heterogeneity by Area of Research

In Table 9 we take up the issue of heterogeneity by the main area of research. We find that economists whose main area of research is history of thought, methodology, heterodox approaches; cultural economics, economic sociology, economic anthropology; or business administration, marketing, accounting exhibit the smallest ideological and authority bias.⁴⁴ We find however that economists whose main area of research is macroeconomics, public economics, international economics, and financial economics are among those with the largest ideological bias.

Another interesting point to highlight is that, while for economists in all research areas the estimated effect of ideological bias is (significantly) smaller than the estimated effect of authority bias, macroeconomists are the only group for whom the estimated ideological bias is significantly larger than the estimated authority bias ($1/3^{\text{rd}}$ versus $1/5^{\text{th}}$ of a standard deviation). This is potentially driven by the fact that our less-/non-mainstream sources induce a stronger ideological reaction in macroeconomists than when we remove the sources altogether.

6.3.4. Heterogeneity by Undergraduate Major

Lastly, we examine heterogeneity by undergraduate major. As we discussed before, there exists growing evidence that suggests economic training, either directly or indirectly, could induce ideological views in students (e.g. Allgood et al. 2012, Colander and Klamer 1987, Colander 2005, Rubinstein 2006). Consistent with these studies, we find that economists whose undergraduate major was economics or business/management exhibit the strongest ideological bias (around $1/4^{\text{th}}$ of a standard deviation). However, we find that economists with an undergraduate major in law; history, language, and literature; or anthropology, sociology, psychology, exhibit the smallest ideological bias (statistically insignificant in all three cases).⁴⁵

7. Conclusion

We use an online randomized controlled experiment involving economists in 19 countries to examine the influence of ideological and authority bias on views among economists. Economists

⁴⁴ For the latter group, this could be driven by lack of familiarity with where different sources stand in relation to mainstream economics and their ideology.

⁴⁵ Of course, this systematic difference could be driven by self-selection of individuals into different undergraduate majors and is not necessarily causal.

who participated in our survey were asked to evaluate statements from prominent economists on different topics. However, source attribution for each statement was randomized without participants' knowledge. For each statement, participants either received a mainstream source, a less-/non-mainstream source, or no source. We find that economists' reported level of agreement with statements is significantly reduced when statements are randomly attributed to less-/non-mainstream source with views or ideologies that put them at different distances to mainstream economics, even when this distance is relatively small. In addition, we find that removing the source attribution also significantly reduces the agreement level with statements.

We use a Bayesian updating framework to organize and test the validity of two competing hypotheses as potential explanations for our results: unbiased Bayesian updating versus ideologically-/authority-biased Bayesian updating. While we find no evidence in support of unbiased Bayesian updating, our results are all consistent with the existence of biased updating. More specifically, and in contrast (consistent) with implications derived from our model under unbiased (biased) Bayesian updating, we find that changing/removing sources (1) has no impact on economists' precision of their posterior beliefs (proxied for by economists' reported level of confidence with their evaluations); (2) similarly affects experts/non-experts in relevant areas; and (3) affects those at the far right of the political spectrum much more significantly than those at the far left. We also find systematic and significant heterogeneity in our estimates of ideological/authority bias by gender, country, country/region where PhD was completed, area of research, and undergraduate major with patterns consistent with ideological/authority bias.

Scholars hold different views on whether economics can be a 'science' in the strict sense and free from ideological underpinnings. However, perhaps one point of consensus is that the type of ideological bias that could result in endorsing or denouncing an argument on the basis of its author's views rather than its substance, is unhealthy and in conflict with scientific tenor and the subject's scientific aspiration, especially when the knowledge regarding rejected views is limited. It is hard to imagine that these biased reactions only emerge in a low-stake environment, such as our experiment, without spilling over to other areas of academic life. After all, a strong majority of experimental studies in economics and psychology are based on low-stake experiments, but we rarely discount the importance of their findings and their implications based on the low-stake nature of the experiments. Moreover, there already exists some evidence that suggested political ideology influences economic research (Jelveh et al. 2018, Saint-Paul 2018).

It is well understood that ideological bias could impede the engagement with alternative views, narrow the pedagogy, and bias and delineate research parameters. We believe that recognising their own biases, especially when there exists evidence that suggests they could operate through implicit or unconscious modes, is the first step for economists who strive to be objective and ideology-free. This is also consistent with the standard to which most economists in our study hold themselves. To echo again the words of Alice Rivlin in her 1987 American Economic Association presidential address, “economists need to be more careful to sort out, for ourselves and others, what we really know from our ideological biases.”

Another important step to minimize the influence of our ideological biases is to understand their roots. As argued by prominent social scientists (e.g. Althusser 1976, Foucault 1969, Popper 1955, Thompson 1997), the main source of ideological bias is knowledge-based, influenced by the institutions that produce discourse. For example, Colander and Klamer (1987) and Colander (2005) survey graduate students at top-ranking graduate economic programs in the US and find that, according to these students, techniques are the key to success in graduate school, while understanding the economy and knowledge about economic literature only help a little. Similarly, Rubinstein (2006) suggests that “students who come to us to 'study economics' instead become experts in mathematical manipulations.” This lack of depth in knowledge acquired, not only in economics but in any discipline or among any group of people, make individuals to lean more easily on ideology.

This highlights the importance of economic training as perhaps the most influential factor in shaping ideological views among economists. It affects the way they process information, identify problems, and approach these problems in their research. In addition, not surprisingly, this training may also affect the policies they favour, and the ideologies they adhere to. As Colander (2005) points out, over time these influences are passed on from one cohort of graduate students to the next: “[i]n many ways, the replicator dynamics of graduate school play a larger role in determining economists’ methodology and approach than all the myriad papers written about methodology.”

Reference:

- [1] Abadie, A., Athey, S., Imbens, G. W., & Wooldridge, J. (2017). When should you adjust standard errors for clustering? (No. w24003). National Bureau of Economic Research.
- [2] Allgood, S., Bosshardt, W., van der Klaauw, W., & Watts, M. (2012). Is economics coursework, or majoring in economics, associated with different civic behaviors? *The Journal of Economic Education*, 43(3), 248-268.
- [3] Althusser, L. (1976). *Essays on ideology*. Verso Books.
- [4] Backhouse, R. E. (2010). *The puzzle of modern economics: science or ideology?* Cambridge University Press.
- [5] Bartels, L. M. (2002). Beyond the running tally: Partisan bias in political perceptions. *Political behavior*, 24(2), 117-150.
- [6] Bar-Tal, Y., & Jarymowicz, M. (2010). The effect of gender on cognitive structuring: who are more biased, men or women?. *Psychology*, 1(02), 80.
- [7] Bertrand M., & E. Duflo (2017). "Field experiments on discrimination" *In Handbook of Economic Field Experiments*, Vol. 1, pp. 309-393, North-Holland.
- [8] Boland, L. A. (1991). Current views on economic positivism. *Companion to contemporary economic thought*, 88-104.
- [9] Carter, J. R., & Irons, M. D. (1991). Are economists different, and if so, why?. *Journal of Economic Perspectives*, 5(2), 171-177.
- [10] Chang, H. (2014). *Economics: The User's Guide*.
- [11] Colander, D. (2005). The making of an economist redux. *The Journal of Economic Perspectives*, 19(1), 175-198.
- [12] Colander, D., & Klammer, A. (1987). The making of an economist. *The Journal of Economic Perspectives*, 1(2), 95-111.
- [13] Currie, J., Lin, W., & Meng, J. (2014). Addressing antibiotic abuse in China: An experimental audit study. *Journal of development economics*, 110, 39-51.
- [14] Dobb, M. (1973). *Theories of value and distribution: ideology and economic theory*. Cambridge University Press, Cambridge.
- [15] DeGroot, M. H. (1970). *Optimal Statistical Decisions*, New York: McGraw-Hill.
- [16] Eagleton T. (1991). *Ideology: an introduction*. Cambridge.
- [17] Fine, B., & Milonakis, D. (2009). *From economics imperialism to freakonomics: The shifting boundaries between economics and other social sciences*. Routledge.
- [18] Fischle, M. (2000). Mass Response to the Lewinsky Scandal: Motivated Reasoning or Bayesian Updating?. *Political Psychology*, 21(1), 135-159.
- [19] Fisher, I. (1919). Economists in Public Service: Annual Address of the President. *The American Economic Review*, 9(1), 5-21.
- [20] Fisman, R., & O'Neill, M. (2009). Gender differences in beliefs on the returns to effort evidence from the world values survey. *Journal of Human Resources*, 44(4), 858-870.

- [21] Foucault, M. (1972). *The archeology of knowledge*, trans. AM Sheridan Smith. London: Tavistock.
- [22] Frank, B., & Schulze, G. G. (2000). Does economics make citizens corrupt?. *Journal of economic behavior & organization*, 43(1), 101-113.
- [23] Frank, R. H., Gilovich, T., & Regan, D. T. (1993). Does studying economics inhibit cooperation?. *Journal of economic perspectives*, 7(2), 159-171.
- [24] Frank, R. H., Gilovich, T. D., & Regan, D. T. (1996). Do economists make bad citizens?. *Journal of Economic Perspectives*, 10(1), 187-192.
- [25] Frankfurter, G. & Mcgoun, E. (1999). Ideology and the theory of financial economics. *Journal of Economic Behavior & Organization*. 39. 159-177.
- [26] Freedman C. (2016). *In search of the two-handed economist: ideology, methodology and marketing in economics*. Springer.
- [27] Frey, B. S., Pommerehne, W. W., & Gygi, B. (1993). Economics indoctrination or selection? Some empirical results. *The Journal of Economic Education*, 24(3), 271-281.
- [28] Friedman M. (1953). *Essays in positive economics*. University of Chicago Press.
- [29] Fryer, R. G., Harms, P., & Jackson, M. O. (2018). Updating beliefs when evidence is open to interpretation: Implications for bias and polarization.
- [30] Fullbrook, E. (2008). *Pluralist economics*. Palgrave Macmillan.
- [31] Fullbrook, E. (2003). *The crisis in economics*. London: Routledge.
- [32] Galbraith, J. K. (1989). Ideology and economic reality. *Challenge*, 32(6), 4-9.
- [33] Gentzkow, M., & Shapiro, J. M. (2006). Media bias and reputation. *Journal of political Economy*, 114(2), 280-316.
- [34] Gerber, A., & Green, D. (1999). Misperceptions about perceptual bias. *Annual review of political science*, 2(1), 189-210.
- [35] Gordon, R., & Dahl G. B. (2013). Views among economists: professional consensus or point-counterpoint? *American Economic Review*, 103(3), 629–635.
- [36] Halligan L. (2013). Time to stop this pretence – economics is not science, *The Telegraph*, October 19.
- [37] Harcourt, G. C. (1969). Some Cambridge controversies in the theory of capital. *Journal of Economic Literature*, 7(2), 369-405.
- [38] Hart, W., Albarracín, D., Eagly, A. H., Brechan, I., Lindberg, M. J., & Merrill, L. (2009). Feeling validated versus being correct: a meta-analysis of selective exposure to information. *Psychological bulletin*, 135(4), 555.
- [39] Hodgson, G. M., & Jiang, S. (2007). The economics of corruption and the corruption of economics: an institutionalist perspective. *Journal of Economic Issues*, 41(4), 1043–1061.
- [40] Hoover, K. R. (2003). *Economics as ideology: Keynes, Laski, Hayek, and the creation of contemporary politics*. Rowman & Littlefield Publishers.


- [41] Ioannidis, J. P., Stanley, T. D., & Doucouliagos, H. (2017). The power of bias in economics research. *The Economic Journal*, 127(605), F236-F265.
- [42] Jelveh, Z., Bruce K., & Naidu S. (2018). Political language in economics. Columbia Business School Research Paper No. 14-57.
- [43] Krugman, P. (2009). *The conscience of a liberal*. WW Norton & Company.
- [44] Lawson, T. (2012). Mathematical modelling and ideology in the economics Academy: competing explanations of the failings of the modern discipline? *Economic Thought*, 1(1, 2012).
- [45] List, J. A., Sadoff, S., & Wagner, M. (2011). So you want to run an experiment, now what? Some simple rules of thumb for optimal experimental design. *Experimental Economics*, 14(4), 439.
- [46] MacCoun R. J., & Paletz S. (2009) Citizens' perceptions of ideological bias in research on public policy controversies. *Political Psychology* 30(1),43-65.
- [47] MacCoun, R. (1998). Biases in the interpretation and use of research results. *Annual Review of Psychology*, 49, 259–287.
- [48] Maniadis, Z., & Tufano, F. (2017). The research reproducibility crisis and economics of science. *The Economic Journal*, 127(605), F200-F208.
- [49] Maniadis, Z., Tufano, F., & List, J. A. (2017). To replicate or not to replicate? Exploring reproducibility in economics through the lens of a model and a pilot study. *The Economic Journal*, 127(605), F209-F235.
- [50] Maniadis, Z., Tufano, F., & List, J. A. (2014). One swallow doesn't make a summer: New evidence on anchoring effects. *American Economic Review*, 104(1), 277-90.
- [51] Marwell, G., & Ames, R. E. (1981). Economists free ride, does anyone else?: Experiments on the provision of public goods, IV. *Journal of public economics*, 15(3), 295-310.
- [52] Merton R. K. (1973). *The sociology of science*. University of Chicago Press, Chicago.
- [53] Meyers-Levy, J. (1986). *Gender differences in information processing: A selectivity interpretation* (Doctoral dissertation, Northwestern University).
- [54] McCloskey, D. (2017). *The Many Transgressions of Deirdre McCloskey*, Institute for New Economic Thinking, June 28, 2017, accessed 30 January 2019 at <https://www.ineteconomics.org/perspectives/blog/the-many-transgressions-of-deirdre-mccloskey>
- [55] Milberg W. S. (1998). “Ideology” in *The Handbook of Economic Methodology*. Cheltenham: Edward Elgar, pp. 243-246.
- [56] Miller, J., & Katz, D. (2018). Gender differences in perception of workplace experience among anesthesiology residents. *The journal of education in perioperative medicine*, 20(1).
- [57] Moretti, E. (2011). Social learning and peer effects in consumption: Evidence from movie sales. *The Review of Economic Studies*, 78(1), 356-393.

- [58] Morgan, J. (Ed.). (2015). What is neoclassical economics?: debating the origins, meaning and significance. Routledge.
- [59] Myrdal, G. (1954). The political element in the development of economic thought. Harvard University Press.
- [60] Piketty, T. (2014). Capital in the twenty-first century. Harvard University Press.
- [61] Popper, K. R. (1962). On the sources of knowledge and of ignorance.
- [62] Rabin, M., & Schrag, J. L. (1999). First impressions matter: a model of confirmatory bias. *The Quarterly Journal of Economics*, 114(1), 37-82.
- [63] Ragins, B. R., Townsend, B., & Mattis, M. (1998). Gender gap in the executive suite: CEOs and female executives report on breaking the glass ceiling. *Academy of Management Perspectives*, 12(1), 28-42.
- [64] Riach, P., & Rich J. (2002). Field experiments of discrimination in the market place. *The economic journal* 112, no. 483.
- [65] Rivlin, A. M. (1987). Economics and the political process. *The American Economic Review*, 1-10.
- [66] Robinson, J. (1973). Ideology and analysis. *Collected Economic Papers*, 5, 254-61.
- [67] Romer, P. M. (2015). Mathiness in the theory of economic growth. *The American Economic Review*, 105(5), 89-93.
- [68] Rothbarb M. (1960). "The mantle *Issues*, of science" In *Scientism and Values*. Princeton NJ.
- [69] Rubinstein, A. (2006). A sceptic's comment on the study of economics. *The Economic Journal*, 116(510), C1-C9.
- [70] Saint-Paul, G. (2018). The possibility of ideological bias in structural macroeconomic models. *American Economic Journal: Macroeconomics* 10(1), 216-241.
- [71] Samuels W. J. (1992). "Ideology in economics" In *Essays on the Methodology and Discourse of Economics*, pp. 233-248. Palgrave Macmillan UK.
- [72] Schumpeter, J. A. (1949). Science and ideology. *The American Economic Review*, 39(2), 346-359.
- [73] Stigler, G. J. (1965). The economist and the state. *The American Economic Review*, 1-18.
- [74] Stigler, G. J. (1960). The influence of events and policies on economic theory. *The American economic review*, 36-45.
- [75] Stigler, G. J. (1959). The politics of political economists. *The Quarterly Journal of Economics*, 522-532.
- [76] Stiglitz J. E. (2002). There is no invisible hand. The Guardian, 20 December, accessed 14 August 2018 at <https://www.theguardian.com/education/2002/dec/20/highereducation.uk1>.
- [77] Taber, C. S., & Lodge, M. (2006). Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science*, 50(3), 755-769.
- [78] Thompson, H. (1997). Ignorance and ideological hegemony: A critique of neoclassical economics. *Journal of Interdisciplinary Economics*, 8(4), 291-305.

- [79] Wang, L., Malhotra, D., & Murnighan, J. K. (2011). Economics education and greed. *Academy of Management Learning & Education*, 10(4), 643-660.
- [80] Wiles, P. (1979). Ideology, methodology, and neoclassical economics. *Journal of Post Keynesian Economics*, 2(2), 155-180.

Tables and Figures

Figure 1: OLS estimates of differences in agreement level between control and treatment groups – By statement



Note: Agreement levels is z-normalized for each statement. Control variables include: gender, PhD completion cohort, current status, country, research area. Both 90% and 95% confidence intervals are displayed for each estimate. The two horizontal lines on each confidence interval band represent where the 90% confidence interval ends.

First (second) listed source for each statement is the actual (altered) source. Bold source for each pair refers to the less-/non-mainstream source. See Table A8 in our online appendix for more details.

Table 1: OLS Estimates of differences in agreement level between control and treatment groups

<i>A: In Units of Agreement Level</i>	(1)	(2)	(3)	(4)
Treatment 1 (none-/less-mainstream source)	-0.264*** (0.014)	-0.261*** (0.014)	-0.262*** (0.014)	-0.268*** (0.014)
Treatment 2 (no source)	-0.415*** (0.015)	-0.404*** (0.015)	-0.406*** (0.015)	†
<i>B: In Units of Standard Deviation</i>				
Treatment 1 (none-/less-mainstream source)	-0.223*** (0.012)	-0.220*** (0.012)	-0.221*** (0.012)	-0.226*** (0.012)
Treatment 2 (no source)	-0.350*** (0.012)	-0.341*** (0.012)	-0.343*** (0.012)	†
P-value: Treatment 1 = Treatment 2	0.000	0.000	0.000	NA
Controls	No	Yes	No	No
More Control	No	No	Yes	No
Fixed Person Effects	No	No	No	Yes
Number of observations	36375	36375	36375	25185

Note: Omitted category is receiving a mainstream source. Heteroskedasticity-robust standard errors are reported in parentheses. The dependent variable is agreement level on a scale from 1 (strongly disagree) to 5 (strongly agree). For panel (B), the dependent variable is standardized to have mean zero and standard deviation of one. The average agreement level in our sample is 3.35 with standard deviation of 1.185. Significance levels: *** < 1%, ** < 5%, * < 10%.

Controls include: gender, PhD completion cohort, current status, country, and research area. More Controls include all the previously listed variables as well as age cohort, country/region of birth, English proficiency, department of affiliation, and country where PhD was completed.

† We cannot identify the effect of treatment 2 in models with individual fixed effects since those who are sorted into treatment 2 receive all statements without a source and therefore there is no variation in treatment within a person and across statements. We therefore exclude these participants from the fixed effects model.

Table 2: Ordered logit estimates of differences in agreement level between control and treatment groups

	Outcome:				
	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
Panel A: Without Controls					
Predicted probability of outcome	0.050***	0.168***	0.166***	0.403***	0.212***
<i>Control Group (mainstream source)</i>	(0.001)	(0.002)	(0.002)	(0.002)	(0.003)
Difference in predicted probability	0.022***	0.050***	0.021***	-0.036***	-0.057***
<i>mainstream Vs. less-/non-mainstream</i>	(0.001)	(0.003)	(0.001)	(0.002)	(0.003)
Difference in predicted probability	0.039***	0.083***	0.029***	-0.067***	-0.085***
<i>mainstream Vs. no source</i>	(0.001)	(0.003)	(0.001)	(0.002)	(0.003)
Panel B: With Controls					
Predicted probability of outcome	0.048***	0.166***	0.169***	0.411***	0.206***
<i>Control Group (mainstream source)</i>	(0.001)	(0.002)	(0.002)	(0.002)	(0.002)
Difference in predicted probability	0.021***	0.051***	0.022***	-0.038***	-0.056***
<i>mainstream Vs. less-/non-mainstream</i>	(0.001)	(0.003)	(0.001)	(0.002)	(0.003)
Difference in predicted probability	0.037***	0.083***	0.030***	-0.068***	-0.082***
<i>mainstream Vs. no source</i>	(0.001)	(0.003)	(0.001)	(0.002)	(0.003)
Number of observations	36375	36375	36375	36375	36375

Note: Robust standard errors are reported in parentheses. Significance levels: *** < 1%, ** < 5%, * < 10%.

The dependent variable is agreement level on a scale from 1 (strongly disagree) to 5 (strongly agree).

Controls include: gender, PhD completion cohort, current status, country, research area.

Table 3: OLS estimates of differences in confidence level

<i>A: In Units of Confidence Level</i>	(1)	(2)
Treatment 1 (none-/less-mainstream source)	0.005 (0.011)	0.008 (0.010)
Treatment 2 (no source)	-0.019 (0.012)	†
<hr/>		
<i>B: In Units of Standard Deviation</i>		
Treatment 1 (none-/less-mainstream source)	0.006 (0.012)	0.009 (0.011)
Treatment 2 (no source)	-0.020 (0.013)	†
P-value: treatment 1 = treatment 2	0.037	NA
Controls	Yes	No
Fixed Person Effects	No	Yes
Number of observations	36088	24984

Note: Omitted category is Control Group (i.e. mainstream source). Heteroskedasticity-robust standard errors are reported in parentheses. The dependent variable is confidence level with evaluation on a scale from 1 (least confident) to 5 (most confident). For panel (B), the dependent variable is standardized to have mean zero and standard deviation of one. The average confidence level in our sample is 3.93 with standard deviation of 0.928. Since confidence level was voluntary to report in our survey, compared to agreement level regressions we lose a small number of observations where confidence level is not reported. Significance levels: *** < 1%, ** < 5%, * < 10%.

Controls include: gender, PhD completion cohort, current status, country, research area.

† We cannot identify the effect of treatment 2 in fixed effects model since those who are sorted into this group receive all statements without a source and therefore there is no variation in treatment within a person and across statements. We therefore exclude these participants from the fixed effects model.

Table 4: OLS Estimates of differences in agreement level between control and treatment groups – By expertise

	(1)	(2)	(3)
	Control group	Treatment 1	Treatment 2
Expert	-0.0239 (0.0293)	-0.224*** (0.0182)	-0.341*** (0.0191)
Non-Expert		-0.216*** (0.0162)	-0.341*** (0.0169)
P-value: equality of coefficients		0.736	0.993
F-statistic: equality of coefficients		0.11	0.00
Number of observations	36375		

Note: Control group refers to receiving a mainstream source. Treatment 1 refers to receiving a less-/non-mainstream source. Treatment 2 refers to receiving no source. Omitted category is expert & control group. Expert is an indicator that is equal to 1 if a participant's reported area of research is related to the area of an evaluated statement and zero otherwise. See Table A10 in the Online Appendix for more details. Heteroskedasticity-robust standard errors are reported in parentheses. The dependent variable is agreement level on a scale from 1 (strongly disagree) to 5 (strongly agree) and is z-normalized for each group. Significance levels: *** < 1%, ** < 5%, * < 10%.

Controls include: PhD completion cohort, current status, country, research area.

Table 5: OLS Estimates of differences in agreement level between control and treatment groups – By political orientation

	Main Results			Robustness 1			Robustness 2		
	Author-created categories			Categories by quintiles of political orientation			Categories by quintiles of <i>adjusted</i> political orientation		
	(1) Control group	(2) Treatment 1	(3) Treatment 2	(4) Control group	(5) Treatment 1	(6) Treatment 2	(7) Control group	(8) Treatment 1	(9) Treatment 2
Far Left		-0.050*	-0.355***		-0.073***	-0.357***		-0.080***	-0.261***
		(0.027)	(0.030)		(0.023)	(0.025)		(0.025)	(0.027)
Left	0.109***	-0.241***	-0.346***	0.044*	-0.244***	-0.282***	0.098***	-0.237***	-0.427***
	(0.023)	(0.019)	(0.019)	(0.026)	(0.030)	(0.030)	(0.024)	(0.025)	(0.027)
Center	0.130***	-0.286***	-0.411***	0.095***	-0.291***	-0.413***	0.113***	-0.262***	-0.389***
	(0.026)	(0.026)	(0.029)	(0.024)	(0.026)	(0.027)	(0.028)	(0.032)	(0.031)
Right	0.136***	-0.315***	-0.333***	0.092***	-0.283***	-0.426***	0.098***	-0.274***	-0.381***
	(0.029)	(0.031)	(0.032)	(0.024)	(0.027)	(0.029)	(0.025)	(0.026)	(0.029)
Far Right	0.138***	-0.319***	-0.349***	0.074***	-0.310***	-0.333***	0.113***	-0.312***	-0.362***
	(0.043)	(0.054)	(0.057)	(0.026)	(0.030)	(0.031)	(0.026)	(0.027)	(0.028)
P-value of equality	0.57	0.000	0.372	0.238	0.000	0.003	0.907	0.000	0.000
F-statistic of equality	0.637	14.57	1.06	1.41	14.76	3.96	0.18	11.49	5.10
# observations		36315			36315			36315	

Note: Control group refers to receiving a mainstream source. Treatment 1 refers to receiving a less-/non-mainstream source. Treatment 2 refers to receiving no source. Omitted category is Far Left & control group. Heteroskedasticity-robust standard errors are reported in parentheses. The dependent variable is agreement level on a scale from 1 (strongly disagree) to 5 (strongly agree) and is z-normalized for each group. Political orientation is self-reported by participants on a scale from -10 (far left) to 10 (far right). Significance levels: *** < 1%, ** < 5%, * < 10%. Controls include: gender, PhD completion cohort, current status, country, research area.

For Columns (1) to (3), we use self-reported political orientation to group participants into 5 categories: Far left = [-10 -7], Left = [-6 -2], Centre = [-1 1], Right = [2 6], Far Right = [7 10]. Results reported in Columns (4) to (9) are for robustness check. For results reported in columns (4) to (6), we create the five political groups using the quintiles of political orientation distribution. For results reported in columns (7) to (9), we create the five political groups using the quintiles of the adjusted political orientation distribution. Adjusted political orientation measure is created by running a regression of self-reported political orientation on a series of indicators based on questions asked from participants to identify their political typology. See Table A9 in our online appendix for more details.

Table 6: OLS Estimates of gender differences in agreement level between control and treatment groups

	(1)	(2)	(3)
	Control group	Treatment 1	Treatment 2
Male		-0.242*** (0.013)	-0.337*** (0.014)
Female	-0.063*** (0.020)	-0.140*** (0.025)	-0.363*** (0.027)
P-value: equality of coefficients		0.000	0.395
F-statistic: equality of coefficients		12.32	0.72
Number of observations	36375		

Note: Control group refers to receiving a mainstream source. Treatment 1 refers to receiving a less-/non-mainstream source. Treatment 2 refers to receiving no source. Omitted category is male & control group. Heteroskedasticity-robust standard errors are reported in parentheses. The dependent variable is agreement level on a scale from 1 (strongly disagree) to 5 (strongly agree) and is z-normalized for each gender. Significance levels: *** < 1%, ** < 5%, * < 10%.

Controls include: PhD completion cohort, current status, country, research area.

Table 7: OLS Estimates of differences in agreement level between control and treatment groups – By country

	(1)	(2)	(3)
	Control Group	Treatment 1	Treatment 2
Australia		-0.327*** (0.056)	-0.539*** (0.059)
Austria	-0.266*** (0.082)	0.022 (0.104)	-0.077 (0.102)
Brazil	-0.333*** (0.072)	0.016 (0.084)	0.015 (0.105)
Canada	-0.026 (0.045)	-0.284*** (0.035)	-0.399*** (0.037)
France	-0.099** (0.047)	-0.223*** (0.041)	-0.366*** (0.042)
Germany	-0.130** (0.055)	-0.181*** (0.061)	-0.240*** (0.065)
Ireland	0.009 (0.116)	-0.458*** (0.154)	-0.445*** (0.157)
Italy	-0.195*** (0.048)	-0.124*** (0.042)	-0.261*** (0.045)
Japan	-0.012 (0.065)	-0.397*** (0.081)	-0.402*** (0.081)
Netherlands	-0.074 (0.068)	-0.264*** (0.081)	-0.134* (0.079)
New Zealand	-0.054 (0.071)	-0.237*** (0.082)	-0.355*** (0.087)
Scandinavia	-0.051 (0.050)	-0.321*** (0.047)	-0.427*** (0.053)
South Africa	-0.049 (0.087)	-0.127 (0.116)	-0.359*** (0.106)
Switzerland	0.042 (0.078)	-0.298*** (0.101)	-0.464*** (0.098)
UK	-0.032 (0.052)	-0.229*** (0.051)	-0.392*** (0.052)
US	-0.076* (0.040)	-0.206*** (0.020)	-0.337*** (0.020)
P-value: equality of coefficients	0.000	0.000	0.000
F-statistic: equality of coefficients	3.80	2.63	3.34
Number of observations	36375		

Note: Control group refers to receiving a mainstream source. Treatment 1 refers to receiving a less-/non-mainstream source. Treatment 2 refers to receiving no source. Omitted category is Australia & control group. Heteroskedasticity-robust standard errors are reported in parentheses. The dependent variable is agreement level on a scale from 1 (strongly disagree) to 5 (strongly agree) and is z-normalized for each country. Significance levels: *** < 1%, ** < 5%, * < 10%.

Controls include: gender, PhD completion cohort, current status, research area.

Table 8: OLS Estimates of differences in agreement level between control and treatment groups – By country/region where PhD was completed

	(1)	(2)	(3)
	Control Group	Treatment 1	Treatment 2
Africa		-0.103 (0.129)	-0.305** (0.128)
Asia	0.297** (0.125)	-0.433*** (0.108)	-0.392*** (0.101)
Canada	0.333*** (0.107)	-0.321*** (0.046)	-0.471*** (0.052)
Europe 1 (France, Belgium)	0.118 (0.109)	-0.166*** (0.040)	-0.265*** (0.042)
Europe 2 (Germany, Austria, Netherlands, Switzerland, Luxembourg)	0.243** (0.108)	-0.208*** (0.044)	-0.278*** (0.042)
Europe 3 (Italy, Spain, Portugal)	0.012 (0.108)	-0.116** (0.048)	-0.233*** (0.055)
Europe 4 (Denmark, Finland, Norway, Sweden)	0.338*** (0.110)	-0.314*** (0.058)	-0.465*** (0.061)
Europe 5 (UK, Ireland)	0.219** (0.106)	-0.181*** (0.045)	-0.338*** (0.046)
Not Applicable	0.020 (0.115)	-0.188*** (0.044)	-0.384*** (0.046)
Oceania	0.255** (0.116)	-0.188** (0.081)	-0.331*** (0.079)
Other	0.286* (0.150)	-0.090 (0.182)	-0.884*** (0.210)
South America	0.134 (0.146)	0.013 (0.112)	-0.040 (0.126)
United States	0.275*** (0.102)	-0.245*** (0.018)	-0.362*** (0.019)
P-value: equality of coefficients	0.000	0.008	0.000
F-statistic: equality of coefficients	2.83	2.23	6.41

Number of observations

36375

Note: Control group refers to receiving a mainstream source. Treatment 1 refers to receiving a less-/non-mainstream source. Treatment 2 refers to receiving no source. Omitted category is Africa & control group. Heteroskedasticity-robust standard errors are reported in parentheses. The dependent variable is agreement level on a scale from 1 (strongly disagree) to 5 (strongly agree) and is z-normalized for each country/region. Significance levels: *** < 1%, ** < 5%, * < 10%. Controls include: gender, PhD completion cohort, current status, country, research area. "Other" category includes Central America, Eastern Europe, Rest of Europe, Middle East, The Caribbean. Due to very small cell size for these countries/regions (135 observations in total), we have put them all in one category.

Table 9: OLS Estimates of differences in agreement level between control and treatment groups – By research area

	(1) Control Group	(2) Treatment 1	(3) Treatment 2
Teaching		-0.140** (0.060)	-0.404*** (0.060)
History of Thought, Methodology, Heterodox Approaches	-0.115** (0.053)	-0.111** (0.049)	-0.236*** (0.053)
Mathematical and Quantitative Methods	-0.011 (0.053)	-0.265*** (0.047)	-0.308*** (0.049)
Microeconomics	0.005 (0.050)	-0.223*** (0.042)	-0.368*** (0.042)
Macroeconomics and Monetary Economics	-0.034 (0.048)	-0.334*** (0.036)	-0.198*** (0.037)
International Economics	0.032 (0.051)	-0.269*** (0.045)	-0.495*** (0.049)
Financial Economics	-0.037 (0.059)	-0.265*** (0.063)	-0.274*** (0.061)
Public Economics	0.005 (0.052)	-0.301*** (0.045)	-0.315*** (0.048)
Health, Education, and Welfare	0.010 (0.053)	-0.233*** (0.049)	-0.496*** (0.056)
Labor and Demographic Economics	-0.022 (0.048)	-0.212*** (0.037)	-0.366*** (0.040)
Law and Economics	-0.063 (0.078)	-0.218** (0.102)	-0.367*** (0.111)
Industrial Organization	-0.012 (0.056)	-0.255*** (0.055)	-0.338*** (0.060)
Economic Development, Innovation, Technological Change	-0.010 (0.050)	-0.149*** (0.042)	-0.495*** (0.043)
Agricultural and Natural Resource Economics	-0.053 (0.051)	-0.171*** (0.043)	-0.372*** (0.045)
Urban, Rural, Regional, Real Estate, and Transportation Economics	-0.054 (0.069)	-0.126 (0.079)	-0.335*** (0.080)
Cultural Economics, Economic Sociology, Economic Anthropology	-0.297*** (0.113)	-0.0736 (0.164)	0.0885 (0.173)
Business Administration, Marketing, Accounting	-0.024 (0.079)	-0.0736 (0.164)	0.0885 (0.173)
Other	-0.288*** (0.065)	-0.237** (0.103)	-0.496*** (0.119)
P-value: equality of coefficients	0.000	0.006	0.000
F-statistic: equality of coefficients	2.81	2.04	4.64
Number of observations	36375		

Note: Control group refers to receiving a mainstream source. Treatment 1 refers to receiving a less-/non-mainstream source. Treatment 2 refers to receiving no source. Omitted category is Teaching & control group. Heteroskedasticity-robust standard errors are reported in parentheses. The dependent variable is agreement level on a scale from 1 (strongly disagree) to 5 (strongly agree) and is z-normalized for each research area. Significance levels: *** < 1%, ** < 5%, * < 10%. Controls include: gender, PhD completion cohort, current status, country.

Table 10: OLS Estimates of differences in agreement level between control and treatment groups – By undergraduate major

	(1) Control Group	(2) Treatment 1	(3) Treatment 2
Other Social Sciences (Anthropology, Sociology, Psychology)		-0.062 (0.104)	-0.236* (0.123)
Business, Management	0.147* (0.084)	-0.223*** (0.050)	-0.262*** (0.054)
Biology, Chemistry, Physics	0.206** (0.094)	-0.137* (0.078)	-0.328*** (0.084)
Computer Science, Engineering	0.246*** (0.093)	-0.133* (0.074)	-0.346*** (0.077)
Earth and space sciences, Geography	0.301*** (0.097)	-0.195** (0.092)	-0.438*** (0.087)
Economics	0.224*** (0.077)	-0.254*** (0.015)	-0.348*** (0.016)
History, Language and literature	-0.018 (0.102)	0.001 (0.092)	-0.058 (0.093)
Law	-0.106 (0.143)	0.079 (0.186)	-0.078 (0.310)
Mathematics, Statistics	0.216*** (0.082)	-0.154*** (0.043)	-0.272*** (0.044)
Philosophy, Political Science, International Affairs	0.180** (0.086)	-0.195*** (0.059)	-0.407*** (0.066)
Agricultural/Environmental Sciences	0.192* (0.104)	-0.118 (0.097)	-0.567*** (0.103)
Other	-0.107 (0.161)	0.179 (0.192)	-0.044 (0.175)
Not Reported	0.234*** (0.083)	-0.226*** (0.048)	-0.429*** (0.047)
P-value: equality of coefficients	0.001	0.023	0.021
F-statistic: equality of coefficients	2.97	2.07	2.09
Number of observations	36375		

Note: Control group refers to receiving a mainstream source. Treatment 1 refers to receiving a less-/non-mainstream source. Treatment 2 refers to receiving no source. Omitted category is Other Social Sciences & control group. Heteroskedasticity-robust standard errors are reported in parentheses. The dependent variable is agreement level on a scale from 1 (strongly disagree) to 5 (strongly agree) and is z-normalized for each undergraduate major. Significance levels: *** < 1%, ** < 5%, * < 10%. Controls include: gender, PhD completion cohort, current status, country, research area.

[Click Here for the Online Appendix](#)