# The Conflict-of-Interest Discount in the Marketplace of Ideas

John Barrios Filippo Lancieri Joshua Levy Shashank Singh
Tommaso Valletti Luigi Zingales*

May 26, 2025

We examine how conflicts of interest (CoI)—financial, professional, or ideological stakes held by authors - affect the perceived credibility of economics research. A randomized survey of economists and the US public shows that CoI reduces trust in findings by 28%, with variation by conflict type. We model the "CoI Discount," estimating that conflicted papers are worth 39% less than non-conflicted ones. These findings are validated through disclosure and citation analysis, a medical meta-analysis, and large-language model simulations. Our results highlight a credibility gap in economics research that is not eliminated by current disclosure practices.

*Barrios : Yale School of Management; Lancieri : Georgetown University Law Center; Levy : University of Southern California; Singh : University of Chicago Booth School of Business; Valletti : Imperial College Business School; Zingales : University of Chicago Booth School of Business. We thank Henry Cummins and Jonathan Lee for their excellent research assistance. We thank Erzo F. P. Luttmer for sharing the data from the American Economic Review. We also thank Elliott Ash, Jon Baker, Stefan Bechtold, Marianne Bertrand, Adam Chilton, John Ioannidis, Steve Kaplow, Ioannis Lianos, John Matsusaka, Giorgio Monti, Aviv Nevo, Rupprecht Podzun, Eric Posner, Amit Seru, Alex Stremitzer, Thomas Wollmann, Amit Zac, participants in conferences and seminars in Brussels, Chicago, Columbia, ETH Zurich, George Mason, Harvard, Imperial, Mannheim, University of Southern California, FGV Sao Paulo/INSPER, and Yale for their very useful comments. The project received approval from the IRBs of the University of Chicago, Imperial College London, ETH Zurich, and The Washington University in St. Louis (UChicago IRB AURA Number - IRB23-1177). The project was also pre-registered on the AEA RCT Registry under the number AEARCTR-0012017 Barrios et al., 2023

**Discloure Statement:** Following the AEA Disclosure Policy Requirements, the authors disclose the following potential conflicts of interest arising from the past three years: *Barrios:* No potential conflicts of interest to report; *Lancieri:* Received a research grant from CERRE, a Brussels' think tank, to engage in research on data access in European Regulations; *Levy:* No potential conflicts of interest to report; *Singh:* No potential conflicts of interest to report. *Valletti:* Was awarded a Leverhulme Major Research Fellowship to support his overall research projects; *Zingales:* Serves on the advisory board of the Opportunity & Inclusive Growth Institute of the Minneapolis Fed, advises a financial startup called Jurna, produced a report on financial markets competitiveness in Chile commissioned by the Inter-American Development Bank, and was compensated for speeches given at different universities and others institutions. You can find Zingales' full disclosure statement here. The authors would like to thank Imperial College London, ETH Zurich, the Stigler Center at UChicago Booth, and WashU's Olin Center for Research in Economics and Strategy for financially supporting this research.

# 1 Introduction

The currency of academic research is ideas, but not all ideas are treated equally in the intellectual marketplace. Their influence depends not only on their novelty or rigor, but also on their trustworthiness: the degree to which a paper influences individuals' prior beliefs about its conclusions. Trustworthiness, in turn, is shaped by multiple factors, one of the most significant being whether the author of a paper has a conflict of interest (CoI).

CoI can be diverse. They can be financial—for example, a researcher who holds stock options in a manufacturer of a drug they are investigating could stand to materially gain from "good news" about the effectiveness of that drug; professional: companies may be more likely to share access to high-level proprietary, otherwise inaccessible datasets with scholars who are more likely to write papers aligned with the company's interests; ideological or others. Since empirical and theoretical analyses often involve subjective judgments, the greater a researcher's interest in a specific direction of an investigation's results, the greater the potential for the results to be biased in that direction. Just as a government regulator can favor large incumbents without being dishonest or corrupt, a researcher can introduce bias in favor of a sponsor without engaging in dishonesty or corruption (Zingales, 2013). Motivated reasoning alone can shape the findings to align them with the desired outcome (Dawson et al., 2002). However, despite widespread recognition that CoI can impact academic research—almost all major economics journals currently have a CoI disclosure policy—there is only scant evidence on which practices lead to a meaningful conflict and how different conflicts affect the trustworthiness of research findings.

In this paper, we examine how different types of CoI shape perceptions of trustworthiness and the value of research. To this end, we survey both academic economists and a representative sample of the US population. Economists are further split into two samples, those at the top of the academic profession (denoted as selected economists) and others (denoted as ordinary economists). Respondents were initially presented with vignettes describing findings from both economics and the physical sciences and asked to indicate their initial trust in the described results. We then treated respondents with randomized exposure to a disclosure statement that presented different types of CoI, asking them how much this added information affected their beliefs. Each vignette focused on one type of CoI. We begin by examining monetary conflicts, where authors receive consulting

fees or research grants from interested entities (a key requirement in journals' disclosure policies). Next, we explore career conflicts that arise when research outcomes can influence the professional trajectory of an author. We then consider data conflicts, whereby study data are controlled by a public or private organization with a vested interest (or not) in the direction of the results. In addition, we analyze academic conflicts, which occur when an author's reputation is linked to a particular set of findings, and political conflicts, which arise when an author's political ideology aligns with specific results.[1]

The survey responses enable the calculation of the impact of the presence of a CoI on trust. On average, CoI reduce the trust in the results of economics articles by about 30%; we call this variable the *CoI Trust Reduction*. This reduction is higher among non-economists than among economists, but it is positive and large even among the most selected group of economists, suggesting that the erosion in trust is not merely the manifestation of an "ignorance" discount by those less familiar with academic research. The extent of this CoI Trust Reduction, however, varies greatly depending on the nature of the CoI. For example, monetary conflicts reduce trust by 27%. Instead, data-access conflicts result in reductions between 20% and 52%, depending on the level of control of the interested party over the data. As another example, ideological (political) CoI reduce trust by 17% on average, but this average masks significant variation between how much Democratic readers trust Republican authors and vice-versa—evidence of growing polarization that can shape how academics read and trust the work of scholars who do not share their political views.

Since we randomized treatments within vignettes, we can shed light on the causal mechanisms behind these CoI Trust Reductions. For instance, while financial incentives matter, we find that the specific amount of compensation matters less than expected. Trust falls by 44% when an author receives \$1M, compared to 37% for an author receiving \$10,000. This result is hard to reconcile with the idea that the primary reason why research results are biased is that researchers sell their results to sponsors. It is more consistent with bias arising from motivated reasoning.

We find that the control and availability of data play a critical role in shaping trust. We

[1]While we do not study the potential bias arising from authors' fears of publishing controversial results, some of our findings suggest that this reticence may be another important source of influence on research, albeit one distinct in nature from the class of CoI-induced biases analyzed in this paper. We plan to explore this issue in future work.

observe significant differences in trust based on whether the data used in a study are proprietary or public. Regardless of the direction of a study's findings, use of proprietary data leads to a 20% reduction in trust, while the use of public data results in a 4% *increase* in trust. If the study finds results in the interest of the private data provider, the reduction in trust rises to 27%. Moreover, if respondents learn that the company providing private data waived the right to review the paper prior to publication, trust *increases* by 12%. In contrast, if the company retains the right to review results, trust decreases by 52%.

Our three samples (ordinary economists, selected economists, and average Americans) have varying levels of expertise. In addition, we asked economists about their area of specialization. Therefore, we are also able to examine the correlation between trust reductions and the expertise of the respondents. Experts tend to trust the result more to begin with and tend to exhibit a lower reduction in trust when a conflict is revealed. Interestingly, this trend also applies within expertise areas. For example, industrial organization economists exhibit a lower reduction in trust than financial economists when the vignette refers to an antitrust paper written by an expert witness in a similar antitrust case. However, the same industrial organization economists show a higher reduction in trust than financial economists when a financial consultant writes a paper about a financial trading strategy. This latter result cannot be justified by expertise alone (these are all Ph.D. economists) and is more compatible with a form of cognitive dissonance of researchers who are willing to admit the bias of others but not of their closer colleagues.

Up to this point, we have focused on how CoI erode trust in research findings by influencing respondents' perceptions of credibility. This reduction in trust may imply a deeper concern: that CoI not only affect how research is perceived but also diminish its overall value. To quantify this broader impact, we introduce a formal model that builds on the work of Ioannidis (2005) and Maniadis et al. (2014). This framework is centered on the concept of the post-study probability—the likelihood that a research finding is true after considering the statistical power of the empirical test and any potential biases introduced by CoI. This enables us to infer the reduction in the value of a conflicted paper from the survey responses, where the value of a paper is measured by its ability to change a prior belief. We call this the *CoI Discount*, that is, the percentage reduction in the value of a paper due to the presence of a conflict of interest. The CoI Discount equals the CoI Trust

Reduction only if respondents did not anticipate any potential presence of a CoI when providing their initial answers on how much they trusted the headline results of a given article.

To measure the expected frequency of CoI, we estimate the frequency of such conflicts by analyzing the disclosures of all articles published in the *American Economic Review* and the *RAND Journal of Economics* during the 2019-2023 quinquennium. Our analysis shows that, on average, 39% of the AER papers and 29% of RAND papers are potentially conflicted. Theory papers are less conflicted than empirical papers. The most common form of CoI is the reliance on discretionary, gated data and formal employment affiliations with entities that have an interest in the research project. Political conflicts are less common.[2]

We can then use these estimates to infer the reduction in value of an article given the frequency of the underlying conflicts. We find that, on average, conflicted papers are worth 39% less than non-conflicted ones. The CoI Discount is the most pronounced for papers relying on gated data with a right to review (CoI Discount of 58%), and lowest for academic conflicts, that is, scholars with a history of publications in a given direction (13%).

One potential criticism of our findings is that they simply reflect a decrease in trust due to *perceived* conflicts rather than *actual* ones. Even if our results merely capture a perception, they would be important. If economists and the public perceive certain results as being less valuable and are less convinced of them, these results are inferior quality products in the marketplace of ideas, less likely to influence public policy and individual or corporate behavior. This happens regardless of the correctness of the perception of bias. Nevertheless, we conduct four "sanity-check" exercises to show that our results are not mere perceptions but are in line with the evidence on the actual biases of conflicted studies.

First, we conduct a comprehensive review of the existing literature on CoI in economics and medicine. In economics, there are only two observational studies (Asatryan et al., 2020; Fabo et al., 2021), which both find that CoI bias research findings toward outcomes favorable to the interested parties. Experimental evidence from Borjas and Breznau (2024) further illustrates this effect, show-

[2]We measure US political appointments by creating a dataset of economists with federal government appointments. Similarly, we used Federal Election Commission (FEC) contributions data to map AER/RAND authors in the 2019-23 quinquennium to their donations (cumulatively greater than $10,000) to federal election candidates during 2000-24.

ing that researchers' pre-existing attitudes significantly affect effect sizes estimated from identical data. In medicine, aggregating meta-studies published in leading journals from 1998 to 2023, we find industry-sponsored research is 1.4 times more likely to produce statistically significant results favoring the sponsor. This average masks substantial methodological heterogeneity: observational studies sponsored by industry have odds of reporting favorable results 22.4 times greater than non-sponsored ones, whereas randomized controlled trials exhibit a smaller discrepancy.

Second, we calibrate our theoretical model using empirical findings from Fabo et al. (2021), demonstrating that optimal belief updating closely mirrors the trust reductions observed in our survey experiments involving similar CoI scenarios, specifically those involving public-entity employees. We further validate and extend our analysis by simulating an additional 1,000 survey responses per randomization using GPT-4 Omni, a large language model trained to emulate human decision making. The simulated responses align closely with our empirical survey results, confirming a generalizable pattern in which CoIs consistently lead to reductions in trust, closely mirroring attitudes observed in the general population.

Finally, we compare our results with those of Leuz et al. (2023). Leuz et al. (2023) estimate the reduced value of conflicted medical studies by analyzing their inclusion or exclusion from a database of up-to-date and important research studies for family physicians (the University of Chicago's Priority Updates from the Research Literature, PURL). This method is likely to underestimate the reduction in the value of a conflicted paper since colleagues are leery of "punishing" a paper of a powerful but conflicted author by not including it in PURL. Indeed, Leuz et al. (2023) find that disclosing an author's CoI decreases the likelihood that their paper is included in PURL by 7% to 16.5%, which is lower than our average 39% CoI Discount for the value of a conflicted paper.

Our findings contribute to many inter-disciplinary streams of scholarship. The medical field pioneered the study of conflicts of interest in academic research. Dozens of studies and meta-analyses, such as Ioannidis (2005), have shown that studies with strong financial backing are more prone to bias due to explicit monetary CoI, raising concerns about the integrity of scientific findings in medicine. Our framework complements this work by providing a formal structure to quantify the trust reductions caused by such conflicts in economics.

In economics, the literature on publication bias dates back to at least De Long and Lang

(1992), who highlighted the tendency of statistically insignificant results to go unpublished (for a comprehensive survey, see Christensen & Miguel, 2018). Ioannidis et al. (2017) find that almost 80% of the effects reported in the empirical economics literature are exaggerated. While this bias leads to the publication of too many unimportant results, it does not necessarily lead to systematically biased results. Earlier research on other forms of CoI primarily examined how political biases influence the choice of key model parameters (Jelveh et al., 2024; Saint-Paul, 2012). As mentioned above, only two economic studies focus directly on CoI. Asatryan et al. (2020) find that government-funded studies tend to report larger fiscal multipliers, suggesting a bias toward favorable outcomes that support government policies. Fabo et al. (2021) demonstrate that central bank-affiliated papers report significantly larger effects of quantitative easing (QE) on both output and inflation, indicating potential institutional biases.

Our survey approach to identifying the effect of CoI is novel in economics but not in medicine, which hosts a vast literature on the effects of CoI on academic independence. Kesselheim et al. (2012) survey 500 board-certified internists and finds that physicians are half as willing to prescribe drugs studied in industry-funded trials as they were to prescribe drugs studied in NIH-funded trials. But they look only at one type of conflict, while we compare a variety of potential conflicts and distinguish between the reduction in trust in a result and the reduction in the value of a paper.[3]

To our knowledge, our work is the first study in economics to quantify the *CoI Discount* applied to academic papers generated by disclosing to readers different author-held conflicts of interest. Our research therefore extends the existing literature by providing a comprehensive examination of how various types of CoI (monetary, career, data access, academic, and political conflicts) impact the perceived trustworthiness of economics research. Our data also allows us to understand how individuals' political and economic priors—their personal biases and beliefs—influence both the interpretation of economics research and the potential CoI associated therewith.

The implications of our findings go beyond the standard appeal for enhanced disclosure. Scholars and others have long documented how conflicted research can be used to influence public

[3]Surveys have also been used to assess patients' willingness to participate in industry-sponsored trials. Snyder et al. (2009) find that low trust in Alzheimer's research represents a major obstacle to recruiting elderly patients for studies, with industry-sponsored drug trials being the least trusted of all. In contrast, others find that industry links do not affect cancer patients' willingness to enroll in experimental drug trials (e.g., Hampson et al., 2006).

policy (McGarity & Wagner, 2008; McIntire & Kantor, 2024; Michaels, 2008). Policymakers and academics are increasingly concerned about the potential dissolution of the lines that separate independent academia from business (or other) advocacy (Kanter, 2024; Lianos, 2024; Nevo, 2024; Valletti, 2024). The public trust in science and universities is also declining (Gallup Inc, 2023), and commentators have identified academics who appear to be "self-interested rather than disinterested" as a reason for this decline (Mills, 2023; Rosenbluth, 2024). At the same time, modern academic research is becoming more expensive and requires access to data increasingly controlled by private parties (Edelson et al., 2023; Wagner, 2023). High-quality research can directly benefit from interaction with industry participants, government agents, think-tanks, and others. This project does not study the *benefits* to science that arise from the relationships that are considered potential CoI. Therefore, our paper should ***not*** be interpreted as recommending a ban on such relations. We do not know the socially optimal amount of industry-academia and other relationships.

That notwithstanding, our data fill a significant gap and point to many problems in current academic practices. For example, our analysis demonstrates that the current disclosure system lacks credibility—only 2% of our economists surveyed believe that academic economists appropriately disclose their conflicts of interest. This lack of trust creates a *negative externality* born by independent, unconflicted researchers. In addition, the relationships that led to the largest CoI Trust Reduction involve private data access, a growing practice that deserves much more attention.

Above all, we hope our results will help open a conversation about the benefits and costs of CoI in economics and academia more broadly. Without a structural change in how conflicted research and researchers are evaluated—including better measurements of what benefits potential conflicts bring to science and whether they outweigh the associated costs—the misalignment of private incentives and public trust will continue and will further erode the long-term credibility of academia and expertise.

The rest of the paper proceeds as follows: Section 2 describes the survey. Section 3 presents the main empirical findings on CoI Trust Reduction. Section 4 introduces a simple framework to interpret the data in light of the model. Section 5 derives the CoI Discount of a paper. Section 6 performs several sanity checks on the results. Section 7 discusses the implications of our findings and proposes several avenues for reform. Section 8 concludes.

# 2 Survey Design and Implementation

To assess how CoI affect trust in research findings, we conducted a survey built around a series of randomized vignettes. Each vignette presented a stylized but realistic research scenario involving a potential CoI, such as financial ties, career incentives, restricted data access, academic reputation, or political alignment. Our survey design follows best practices for experimental survey research as outlined by Stantcheva (2023), including the use of randomized treatments to generate identifying variation, structured belief elicitation before and after information exposure, and attention to priming and selection concerns to ensure internal validity.

## 2.1 Survey Structure

In the survey's preamble, we informed participants that they were part of a study developed to understand their perceptions about potential CoI. Participants were told they would evaluate a series of potential CoI in relation to published articles. We informed participants that all articles were of high quality, having been published in a highly reputable peer-reviewed journal and conducted using accepted professional standards.

The participants were presented with nine vignettes in random order. Each vignette began with a statement describing the key finding of the article. The respondents were then asked to rate how much they believed the finding on a 5-point Likert scale ranging from "Not at all" (which we code as zero) to "Completely" (which we code as 4). For example, a vignette read: 'You have just finished reading an article about the role of abortion access in women's career outcomes. In particular, the paper finds that increased access to abortion has *no effect/positive effects* on women's lifetime earnings. How much do you believe the paper's results?". The text in italics indicates the random variation within the vignette. Each participant saw each vignette only once. Subsequently, respondents were randomly treated with information about a potential CoI held by the study's author(s). They were then asked how this disclosure changed their trust in the paper's results using a 7-point Likert scale ranging from "It completely makes me distrust the results" to "It completely makes me trust the results." For example, the treatment in the abortion vignette presented a potential political conflict of interest and read: "Later, you read the disclosure statement associated with the

paper and learn that the author discloses that she is formally a member of the *Democratic/Republican Party*. She additionally discloses that in recent election cycles, she has made contributions in excess of $10,000 to *Democratic/Republican Party* presidential and congressional candidates. To what extent does this disclosure change your belief in the paper's results?"

We specified that the CoI was clearly disclosed in the article, ensuring that respondents did not infer any deception by the author.[4] Appendix A contains a QR code to the full text of the survey, including all random variations.

### 2.1.1 Vignettes

Table 1 summarizes the nine vignettes. In addition to seven vignettes that focus on economic findings, we also presented two based on scientific results to examine potential differences in responses to findings in economics and in the physical sciences (a drug test and a chemical test; see Panel A). Of the seven economics vignettes, five of them addressed economic CoI, specifically related to data access, funding, career opportunities, and consulting engagements (Panel B). The last two vignettes focused on ideological conflicts (Panel C): academic conflicts generated by the desire to have consistent findings across papers (tax vignette) and political conflicts (abortion vignette). The vignettes were partially inspired by real-world situations (such as the controversial approval of a new Alzheimer's drug and the hazardous chemical spill of a Norfolk Southern train in East Palestine, Ohio) but were sufficiently simplified to be understandable without additional context. This simplification was also important because we presented the exact same vignettes to academic economists and to a representative sample of the US population that does not necessarily have any formal training in economics.

At the end of the survey, we asked respondents about their prior experience with relationships that may qualify as potential CoI (e.g., consulting practices, data usage, political appointments, and affiliations), other sources of pressure on academic research, and what they believe are the best CoI disclosure practices. We also collected data on socio-political and the economic-political

[4]There are two exceptions to this rule: the science test vignette, where there is no article, and no possibility of disclosure, and when the CoI involves the author's desire for a future corporate board position, where a standard disclosure is not very credible. See below for further details.

alignment of the respondents, trust in large corporations, voting behavior, PhD year, and area of expertise in economics. The general public was not asked the last two questions.

## 2.2 Survey Target Groups

We surveyed three different groups. First, a group of *"Selected Academic Economists"*, which included current members of the NBER, the CEPR, and the expert panel of the University of Chicago Kent A. Clark Center for Global Markets. We obtained the email of all members by checking their public academic profiles. We successfully contacted 3,051 respondents and the overall response rate was 18%.

Second, a group of *"Ordinary Academic Economists"* used by Javdani and Chang (2023) in their paper on ideology in the publication of economics. The authors kindly shared their contact list with us. It contains contact information for 16,126 PhD economists, mostly based in the US (10,369 people), Canada (2,044 people), the UK (1,903 people), and Australia (1,810 people). We removed the information of those belonging to the "Selected Academic Economists" group, and after bounce-backs, we contacted 12,336 economists. The overall response rate was 8%.

Finally, to survey the *"General Public"* we instructed YouGov, a professional surveying company, to administer our survey to a representative sample of US adults. YouGov interviewed 1,812 respondents, who were then matched with a final data set of 1,500 responses. This group is representative in relation to gender, age, race, education, family income, and political preferences.

These three samples represent different segments. The Selected Economists group consists primarily of full-time scholars who regularly publish, edit, and read articles that appear in the top economic journals. In a way, these are the academic leaders of the profession. The Ordinary Economists group forms the rank-and-file of the economic profession. Their opinions are important because they represent the views of the majority of those who produce and consume economic knowledge. Finally, the General Public group consists of a representative sample of the American population, reflecting the ultimate users in the marketplace of ideas.[5]

[5]We also tried to obtain responses from business journalists, a group that plays an important role in reading, interpreting, and communicating to the public the results of academic research in economics in the relevant context.

## 2.3 Selection and Priming

In recruiting the economists' sample and administering the survey in all three samples, we faced the delicate issue of how transparent we wanted to be. Given the large number of CoI questions present in the survey, there was no chance that the respondents would eventually not understand what the survey was about. Thus, rather than letting the subjects discover the purpose of the survey during the survey itself (probably at different speeds), we decided to put all the subjects on equal footing by being transparent about the goal of the survey. This transparency may have incurred two potential costs: selection bias and priming. Both biases are live possibilities, meaning that their presence (or lack thereof) must be tested for empirically.

We begin by examining the risk of inducing selection bias. Because answering our survey was voluntary, the economists' sample may be biased in the direction of people who are more sensitive to conflicts of interest. It is reassuring that 18% of the top economists answered. This is an unusually high percentage in this type of survey. Even if our results were read as the average change in trust between one out of five top economists, it would still be interesting and relevant. However, additional tests do not indicate the presence of selection bias. We collected data on the year of PhD graduation for our selected economist sample (in five-year intervals). We can then check that data with the underlying real distribution of PhD graduation years for NBER, CEPR and IGM economists. As Figure 1 shows, our sample is slightly younger than the real distribution, but this difference is small. This helps ease fears of selection, at least for that group. In addition, the YouGov sample does not face any sample selection. YouGov chose the participants among its regular panelists based on sample characteristics, and these participants were not informed about the content of the survey when they were invited. Therefore, we can test whether the economists' sample is selected by comparing the answers with the YouGov sample in our sciences vignette—a neutral topic where economists are not particularly self-interested. The average reduction in trust in science among economists is identical to the average reduction in trust among average Americans.

Turning to potential priming connected to the fact that respondents were informed that the

However, despite multiple attempts, we were unable to elicit many responses (only 184 in total). Given the lack of representativeness of the sample, we elected to exclude this group from the paper.

survey focuses on CoI. It is not theoretically clear how priming would affect our results. Arguably, because we first ask respondents about their initial trust in the conclusions of a paper and then rely on disclosure of a CoI as the relevant treatment, priming could lead us to underestimate the impact of CoI on trustworthiness. If respondents expect conflicts in most vignettes even before treatment, they would tend to trust the papers less to begin with, and be less surprised by the subsequent CoI revelation. Because each of the *vignettes* were presented to respondents in random order, we can use that variation to empirically test for priming across questions. Figure 2 shows the average CoI Trust Reduction per vignette position for both the prior trust levels and the CoI Trust Reduction. The differences are not statistically significant, diminishing concerns that we elicited an experimenter's effect. This exercise mitigates some concerns for priming in our *within-subject* survey setting whereby the same respondent reacts to a given vignette first with no mention of conflict and then with conflict. Additionally, we also show robustness checks for this when we run our survey differently and do the analysis *between-subjects*, by showing some respondents a result with a conflict and comparing responses to those of respondents who are shown the same result without the conflict (see section 6.4 for more details).

### 2.3.1 Administration and Summary Statistics

Our survey was pre-registered on the AEA RCT Registry under the number AEARCTR-0012017 (Barrios et al., 2023). We used Qualtrics to administer the survey to both Selected and Ordinary Economists. We used two distinct methods to contact the various survey participants. For the NBER and IGM members of the Selected Economists, and for the Ordinary Economists Group, we sent emails on behalf of Luigi Zingales from an official University of Chicago email address. CEPR members were directly emailed by CEPR on behalf of Tommaso Valletti. Four waves of emails were sent, starting on September 27, 2023, with subsequent follow-up emails on October 10, 23, and 30.

The survey was conducted anonymously to encourage truthful responses. Aside from differentiating between target groups (NBER, IGM, CEPR, and Ordinary Economists)—we did not collect or record any personal identifiers. However, some respondents did reach out to us with comments or questions, allowing us to confirm that our sample included a range of participants, from Nobel

Prize-winning economists to non-tenure-track economics researchers. Importantly, we are unable to link any identities to specific survey responses. As for the general public sample, YouGov administered the survey through their proprietary platform between October 20 and 30, 2023.

The survey included an attention test question.[6] For our main analysis, we only use respondents who passed the attention test. Table 2 includes information on outreach, response rates, and attention test pass rates for each of the three target groups.[7]

Table 3 provides detailed summary statistics for the three surveyed groups: Selected Academic Economists, Ordinary Academic Economists, and the General Public. These tables include demographic information such as age, gender, and educational background, as well as respondents' political and economic orientations. Additionally, they offer insights into respondents' prior experiences with conflicts of interest, their trust in large corporations, voting behavior, and, for academic economists, their year of PhD completion and area of expertise.

# 3 Trust Reductions from Conflicts of Interest

## 3.1 Summary Statistics: Baseline Trust Levels

Before studying the specifics of how CoI influence trust and value, we examine individuals' baseline trust levels toward science and economics. Figure 3 (Panel A) illustrates the baseline trust levels toward the scientific and economics results presented in the nine vignettes, broken into the three sub-samples. We observe that, on average, individuals trust the results in the science vignettes more than in the economic vignettes (difference = 0.49 likert points, $p < 0.01$). This is true for all three subsamples, albeit selected economists are always more trusting (difference = 0.44 likert

[6]To check if the respondents were paying attention and not haphazardly responding to questions, we incorporated an attention test question, randomly in between vignettes (though never in the first position), with the same structure as our vignettes: *You have just finished reading a paper about the role of tariff protections for infant industries in developing countries. Sometimes survey respondents don't pay much attention to the survey. If you are reading closely, please select A little and Completely.*

[7]In addition to the attention question, respondents could exit the survey at any time. However, only 10 Ordinary Economists and 10 Selected Economists failed to complete the survey after having passed the attention test.

points, $p < 0.01$), followed by ordinary economists and the average Americans last. [8]

In Table 4, we compute the average trust by type of vignette, breaking down the different treatments in the initial vignette. The most trusted result is the scientific test on pollution, with an average of 2.21. The least trusted result is the vignette reporting a study finding that CEOs are underpaid in relation to the value they deliver to companies, with an average trust level of 0.98.

## 3.2 Measuring the Reduction In Trust

Next, we turn our attention to the reduction in trust in the presence of conflicts of interest. Once respondents were informed of a potential conflict of interest, they were asked to indicate how it impacted their level of trust using a symmetrical 7-point Likert scale. We opt for Likert scales because they are more easily interpretable by respondents relative to numerical scores.

However, while Likert scales are useful for gathering data, the resulting scores can be difficult for investigators to handle. To address this, we converted the 7-point Likert scale used in our treatments to numerical values corresponding to percentage changes in the trust level (see Table 5). The conversion is straightforward for the extreme and midpoint responses. For example, selecting "[It] completely makes me distrust the results" corresponds to a 100% decrease in trust, while "It does not impact my trust" corresponds to a 0% change in trust. For symmetry, we assume that "It completely makes me trust the results" corresponds to a 100% increase in trust.[9] The intermediate responses are more subjective. We map "It seriously increases my trust in the results" to a 50% increase in trust and "It somewhat increases my trust in the results" to a 20% increase, with corresponding mappings for decreases in trust.

We validate our approach using two methodologies. First, we administer our survey to a representative sample of Americans on Prolific, an online platform that conducts surveys, randomizing the type of trust-reduction scales (half of the respondents were shown the Likert scale and the other

[8] We report the results of the t-tests and their respective p-values in the online appendix, tables 8,9 and 10.

[9] This assumption may not hold for cases of very high initial trust, but this issue has little impact on our main results because this option was selected infrequently—only 190 out of 22,635 responses, or 0.8% of the sample. Even if we replace the 100% increase with a more conservative adjustment based on pre-disclosure trust levels, our results remain virtually unchanged.

half instead the percentage-mapped scale). On average and across all vignette randomizations, we find no significant differences in responses when comparing the two different scales (see online appendix, Section 1.3.1 for detailed methodology and results).

Next, we use GPT-4 Omni to generate a semantic mapping of the Likert scale to numerical values. We run a 1,000 simulations per Likert-scale option. GPT-4 Omni assigned on average a larger reduction compared to ours when there is a negative change in trust, and a smaller increase when trust goes up post-disclosure. Hence our results could be biased downwards which would strengthen our conclusions. These different mappings do not affect the cross-sectional variation or statistical significance of our findings, even if they impact the absolute levels (see online appendix, Section 1.3.2 for the detailed methodology and results).

Finally, in light of the fact that trust levels generally diminish under most treatments upon the disclosure of a CoI, we define CoI Trust Reduction as the negative change in trust. Specifically, a positive trust reduction signifies a decrease in trust following the disclosure, whereas a negative trust reduction indicates an increase in trust. This methodological convention facilitates the interpretation of our graphs and tables.

## 3.3 Do CoI Reduce Trust in Research?

We now turn to the question: to what extent do CoI reduce trust in research across various fields and conflict types? Figure 3 (Panel B) presents the CoI Trust Reduction averaged across groups of vignettes and in different sub-samples.

The three groups of respondents exhibit a similar reduction in trust in the science vignettes (42-44%). Overall, we see a 30% trust reduction in economics vignettes, but with different levels in the sub-samples. The average is 34% among ordinary Americans, 29% among ordinary economists, and 25% among selected economists. This smaller reduction in trust among the selected economists can be due to their knowledge that conflicts are more pervasive (and thus they are less surprised) or their perception that biases are less severe in economics. We will return to this issue below.

### 3.3.1 CoI Trust Reduction and Initial Trust Levels

CoI Trust Reductions can be associated with an overall distrust of the initial results of a given academic article. That is, people discount more CoI for articles that they do not trust to begin with.

Figure 4 shows the CoI Trust Reduction plotted against the Prior Trust levels across all vignettes and respondent types. The size of the circle is directly proportional to the number of observations. There is a negative relationship between the two. Table 6 confirms this result more formally. The dependent variable is the CoI trust reduction for each vignette. The explanatory variables are the initial level of trust alone as well as the initial trust interacted with an indicator variable for the two subgroups: ordinary economists and selected economists.

In all the vignettes, there is a negative correlation between initial trust levels and the subsequent CoI trust reduction. As the CoI discount is the *negative* of the drop in trust, the results indicate that the average discount is more associated with the strength of the original beliefs—subjects revise their beliefs less when they hold them strongly to begin with. This effect is smaller for ordinary economists and especially so for selected economists.

This smaller reduction in trust can be due to a higher expectation that CoI will be present, leading to a lower initial trust in the result, or to a more benign view of the potential degree of bias generated by a given conflict. Yet, the fact that economists trust the baseline results more than average Americans before our vignettes reveal the presence of a conflict contradicts the first hypothesis. If economist respondents expected a CoI to begin with, they should have trusted the baseline results less. Thus, economists seem to treat conflicts as generating less bias than average Americans. This relatively benign treatment of what is ostensibly a colleague's work can be objective or behavioral. Our analysis will disentangle these two plausible explanations by looking at specific economic sub-fields. We will also return to this point with individual-level regressions after the introduction of a formal model in Section 4.

## 3.4 Evidence from Randomized Vignettes

The above results show that, on average, awareness about authors' CoI significantly decreases the trustworthiness of economic articles. However, the seven economic vignettes present different

types of conflict and are subjected to different treatments. We now provide better context to these discussions by studying each type of CoI separately, namely: (i) monetary incentives, (ii) career incentives, (iii) access to data, (iv) academic conflict, and (v) ideological conflict.

### 3.4.1 Monetary Incentives

Monetary incentives are prominently featured in the American Economic Association's (AEA) Disclosure Policy, which requires authors to disclose *"significant financial support, summing to at least $10,000 in the past three years, in the form of consultant fees, retainers, grants and the like."* Thus, we start by showing results for vignettes dealing with CoI related to such monetary incentives.

Figure 5 Panel A presents the findings related to a vignette that summarizes a study that investigates the profitability of a financial trading strategy based on investors' responses to news events. Two independent randomizations are used. The first pertains to the nature of financial support, where we distinguish between a grant and a consulting fee. The second dimension concerns the monetary amount: subjects are informed that the study's author received either $10K, $100K or $1M from a company that utilizes such strategies.

Financial support through a grant produces less mistrust than that of consulting fees. If we pool across samples, the average CoI trust reduction for research grants is smaller (34%) than the one for consulting fees (40%), a difference that is statistically significant (difference = 6 p.p., $p < 0.01$). Notably, even among selected economists, receiving a research grant still results in a substantial 28% reduction in trust. Although higher compensation amounts correlate with greater reductions in trust, the incremental reductions are relatively modest. For example, the average trust reduction for a $10K consulting fee is 37%, while a $1M consulting fee results in a 44% reduction. Selected economists, however, appear comparatively indifferent to the magnitude of the compensation when the payment is structured as a consulting fee. In contrast, when compensation is in the form of a research grant, the trust reduction among selected economists increases significantly, from 21% for a $10K grant to 32% for a $1M grant (difference = 11 p.p., $p = 0.04$).

Panel B further examines how expertise influences trust reductions, distinguishing between industrial organization (IO) economists (Sub-Panel B1) and financial economists (Sub-Panel B2).

Given the financial nature of the vignette, focusing on profitable trading strategies, differences in trust reductions across these sub-fields are particularly relevant. Trust among IO economists drops by 38% when the author receives $1M in compensation (irrespective of whether it is provided as a research grant or consulting fee). By comparison, financial economists exhibit a significantly smaller reduction in trust of only 19% (difference = 19 p.p., p = 0.03). This raises an important question. Do financial economists generally perceive financial conflicts as less problematic, or is their lower sensitivity specific to conflicts within their own area of expertise? To address this, we examine an additional vignette in the next section.

In Figure 6, we plot the CoI trust reduction for the four treatments in the other vignette that analyzes monetary conflicts. This vignette describes an article studying the connection between a merger involving supermarket chains and consumer welfare when one of the authors of the article had been an expert witness for either the government (DoJ) or a private party litigating a similar transaction. The vignette is arranged so that the presented article's result is always in the direction favored by the sponsor (the DoJ expert finds that supermarket mergers reduce welfare, while the defendant expert witness finds welfare increases). The two randomizations focus on the nature of the sponsor (private party or government) and whether the respondent is made aware that the author received $400K for his expert testimony (not an unusual amount in high-profile cases). Our results show that both awareness of the amount paid and the ultimate employer matter.

First, the trust reduction is smaller for the results obtained by a former DoJ expert witness compared to those of a former expert witness for a private defendant (difference = 13 p.p., $p < 0.01$). This pattern is consistent across all three respondent groups, regardless of compensation disclosure. Specifically, when compensation is not mentioned, the trust reduction toward a former DoJ expert witness averages only 8%, whereas this reduction significantly increases to 21% when the $400K compensation is explicitly disclosed (difference = 13 p.p., $p < 0.01$). This larger reduction is mainly driven by ordinary Americans, who exhibit a significant trust reduction of 25% in this scenario (difference = 15 p.p., $p < 0.01$). When respondents learn that the expert received $400K in compensation, the trust reduction increases further, to 36% if the sponsor is a private defendant, a difference that is significant compared to the DoJ scenario (difference = 15 p.p., $p < 0.01$).

Panel B further decomposes the sample of economists according to their area of expertise. In this IO-focused vignette, IO economists demonstrate significantly smaller trust reductions (23%) compared to financial economists (39%) when informed that the author was previously an expert witness for a private defendant that received $400 K (difference = 16 p.p., p = 0.04). These results suggest that financial economists are not broadly indifferent to CoIs; rather, they display significantly smaller sensitivity only toward conflicts directly relevant to their own specialization.

This differential behavior among economists across sub-disciplines may stem from two possible mechanisms. First, economists familiar with a particular field might better anticipate conflicts of interest, reducing their surprise when such conflicts are disclosed. Second, experts may underestimate biases within their own field, as acknowledging these biases could imply their own susceptibility to bias.

We evaluate the first hypothesis using initial trust ratings collected prior to disclosing the conflicts. When evaluating findings suggesting that supermarket mergers decrease welfare, initial trust is similar among IO economists (1.97), financial economists (1.91), and other economists (1.98), with no statistically significant differences (all pairwise comparisons $p > 0.50$). However, for findings indicating that mergers increase welfare, IO economists display significantly higher initial trust (2.56) than both financial economists (2.15; difference = 0.41 likert points, p = 0.02) and other economists (2.28; difference = 0.27 likert points, p = 0.05). Similarly, in the profitable trading strategy vignette, financial economists exhibit higher initial trust (2.01) compared to IO economists (1.86; difference = 0.15 likert points, p = 0.26), and significantly higher trust compared to other economists (1.82; difference = 0.19 likert points, p = 0.02). These results demonstrate that experts do not initially display lower trust in research findings from their own specialization, contradicting the first hypothesis. By exclusion, the smaller trust reduction among experts aligns more strongly with the second hypothesis: economists underestimate biases within their own areas of expertise. Overall, monetary incentives reduce trust by an average of 27%, computed from trust reductions observed in the trading strategy vignette (37%) and the supermarket merger vignette (21%).

### 3.4.2 Career Incentives

CoI are not limited to consulting fees and grants. Career incentives—such as the prospect of higher salaries, increased prestige, or greater professional influence—can also distort the incentives of academics. Unlike direct financial incentives, career-related biases are often subtle, delayed, and difficult to disclose. Given this complexity, survey methods provide clearer insights than archival data analysis alone.

Our vignette described an article that found that CEOs are underpaid relative to the value they render to their companies. The main variation in informational treatment was timing, as respondents were made aware that the author was (i) contemporaneously seeking a position on the board of publicly listed companies, (ii) had been initially nominated to such a board, or (iii) was nominated two years after publishing the study. The respondents distrusted an author with these career goals (an average trust reduction of 36%, see Figure 7), but the temporal gap between the publication of the study and the nomination mattered. The CoI trust reduction for authors seeking a position was 42%, versus 30% for authors nominated two years after publication (difference = 12 p.p., $p < 0.01$). Overall, the trust reduction in the presence of career incentives is consistently larger for average Americans than for economists (difference = 11 p.p., $p < 0.01$) and is also larger for ordinary economists than for selected ones (difference = 7 p.p., $p < 0.01$). However, even selected economists substantially reduce their trust (by 36%) in the results of an author potentially motivated by certain career aspirations.

Next, we explore how the goals of economists' own career affect their responses. Among economists who expressed an interest in future private consulting or board appointments, the reduction in trust in authors actively seeking board positions was significantly smaller (31%) compared to economists without such interests (43%), a difference of 12 percentage points ($p < 0.01$); see Figure 21, Panel D in the online appendix).

### 3.4.3 Access to Data

The AEA Disclosure Policy requires authors to disclose if an interested party provided data for a research project. Hence, we examine whether the source and access conditions of confidential

data affect readers' trust in research findings. We first consider the role of data access using a ride-sharing vignette. The specific vignette concerns a paper that studies the welfare effects of the entry of a ride-sharing company in a given city. We randomize two elements in the vignette. The first randomization concerns the content of the paper: the results could be favorable to the ride-sharing industry (the consumer surplus generated exceeds the social costs) or unfavorable (the social cost exceeds the consumer surplus generated). The second randomization concerns the content of the disclosure: the data could come from the government or one of the industry operators.

Figure 8 illustrates the changes in trustworthiness. Regardless of the paper's results, learning that the results relied on publicly available data *increases* the trust in the conclusions of the article by an average of 4%. The update is slightly larger when the results were favorable to the industry (5% trust premium versus a 3% trust premium), but the difference is small and only borderline statistically significant ($p = 0.09$). In contrast, disclosing that the article relied on proprietary data (privately owned and not publicly available) leads to an average CoI trust reduction of 20%. This decrease in trust is sensitive to the direction of the results: a paper with results unfavorable to the industry suffers a 13% trust reduction if the data are proprietary, while one with results favorable to the industry suffers twice as much—a 27% drop in trust (difference = 14 p.p., $p < 0.01$). The trust result here is driven by economists. Average Americans do not appear to significantly increase their trust when the results come from publicly available data.

Potentially, trust in public data might be driven by ideological factors. To investigate this aspect, in Figure 9, we decompose the trust revision according to the economic ideology of the respondents (the results are similar when we use the social ideology). We present the results only for selected economists (the other groups are in the online appendix). Across ideological lines, there is a consistent pattern: all groups demonstrate greater trust in public data and lower trust in private data, regardless of the study's conclusions. However, the most liberal selected economists exhibit reduced trust in studies with positive welfare effects from ride-sharing, even when the data is public, as they are more skeptical of results that align with industry interests. Further analysis in the online appendix (Section 1.2) reveals that economists with prior experience using proprietary data show the largest trust increases in public data scenarios (difference = 5 p.p., $p = 0.03$), indicating their greater awareness of potential biases induced by the use of proprietary data.

The previous vignette compares publicly owned accessible data with privately owned proprietary data. Yet, while data ownership and access can go hand-in-hand, there are many instances of publicly owned but not publicly accessible data (for example, the CAMELS financial stability rating data is owned by the NCUA—a government agency—but revealed only to bank managers and Fed employees) and privately owned but publicly accessible data (for example, Compustat data that is available to anyone upon payment).

We hypothesize that data accessibility, specifically whether an interested party can restrict researcher access, is a critical factor that affects trust. Restricted access can enable quid-pro-quo arrangements. To test this, we introduce another vignette in which the initial prompt describes the results of a paper that studies the efficiency of a digital payment platform. We state that the data were provided by a credit card company to make it clear that the data are proprietary and focus on the effect of different access policies. In one treatment arm, we reveal that the data provider has waived the right to review the paper; in the other, the data provider has the right to review the paper before the results are made public.

In Figure 10, we plot the CoI Trust Reduction for each treatment. When a subject learns that the data provider has waived the rights, trust in the results of the paper *increases* by an average of 12%. This increase is driven by economists: average Americans increase their trust by a mere 3%, which is not statistically different from zero, ordinary economists increase their trust by 23%, and selected economists by 19%. When a subject learns that the data provider has the right to block circulation, her trust in the results of the paper drops by 52%, the largest drop observed in our survey. This trust reduction is more accentuated for economists (57%) than for ordinary people (46%); (difference = 11 p.p., $p < 0.01$). The online appendix (Section 1.2) shows that trust increases significantly more among economists who had previously engaged in private consulting (20%) or used proprietary data themselves (19%) compared to those without such experiences (differences 11 p.p. and 9 p.p. with $p < 0.01$ and $p < 0.01$, respectively) when the right to review is waived, suggesting greater sensitivity to potential biases inherent in proprietary data.

These findings highlight the impact of data-access terms on perceived research credibility. However, current disclosure practices often lack specificity in terms of data access rights/standards. Standard disclosures typically indicate only the potential presence of a review right without detailing

the exact terms. This general approach fails to address the considerable influence that detailed disclosure of data access agreements could have on trust in academic research.

### 3.4.4 Academic Conflicts

So far, we have focused on economic CoI, as these are more regularly subject to disclosure requirements. However, there are other sources of pressure that can lead to bias. One such source, which we term "academic conflicts," concerns a scholar who has associated her name and reputation with a particular set of results so that she is not interested in contradicting her previous results in future studies.

To test this hypothesis, we present respondents with a vignette describing a paper studying the effect of higher income taxes on economic growth. Half of the respondents are told that the paper finds some growth benefits of higher taxes, while the other half are told that the paper finds no benefits. Subsequently, we reveal to all respondents that the author of the paper had previously written a book about the beneficial effects of higher taxes on growth.

Figure 11 illustrates the differences across groups. Average Americans do not differentiate significantly between treatments; they reduce trust by roughly 19% in both scenarios, with no statistically significant differences ($p = 0.77$). Economists, however, respond significantly to aligning the current results with the previous work of the author. The trust in the paper *increases* by 4% when the findings contradict the previous position of the author and decreases by 5% when the findings align with previous views, a statistically significant difference ($p < 0.01$). Nevertheless, the magnitude of trust reduction in academic conflicts is substantially smaller than that observed in economic conflicts. Overall, academic conflicts reduce trust by approximately 12% when the findings align with previous work by the author.

### 3.4.5 Ideological Conflicts

Our final vignette examines the effect of ideological conflicts of interest. Previous survey studies have shown that ideology can bias perceptions in economics (e.g., Javdani & Chang, 2023). Our randomized survey allows us to more precisely compute the trust reduction engendered by an

ideological conflict, as well as the extent to which this magnitude reflects the political views of the respondent. Note that there is no mandatory requirement to disclose this conflict in academic research, and such conflicts are rarely voluntarily disclosed by authors.

This vignette includes two levels of randomization. First, respondents were shown an article finding that access to abortion had either a positive effect on women's lifetime earnings or no effect. Second, respondents were informed via a disclosure that the author was either a Democrat or a Republican, having made donations to candidates of their respective party in recent elections.

As Figure 12 illustrates, trust in the results decreases by an average 32% when a Republican author finds no impact of abortion access on women's lifelong earnings (a "Republican" finding). In contrast, the trust reduction is smaller, an average of 19%, when a Democratic author finds a positive effect (a "Democratic finding"); this difference is statistically significant (difference = 13 p.p., $p < 0.01$). Surprisingly, a Democratic author presenting a Republican finding still experiences a 13% trust reduction, while a Republican author presenting a Democratic finding faces a negligible reduction, statistically indistinguishable from zero.

Further analysis of trust reductions across subgroups further clarifies these differences. The average American exhibits substantial trust reductions, ranging from 11% when a Republican author presents Democratic findings to 33% when a Republican author presents Republican findings; this difference is statistically significant ($p < 0.01$). In contrast, trained economists slightly *increase* their trust when a Republican author presents Democratic findings, a response significantly different from that of the general public (difference = 16 p.p., $p < 0.01$). Economists reduce trust markedly (32%) when a Republican author presents Republican findings. For findings authored by Democrats, economists show a 11% trust reduction when results align ideologically (Democratic findings), and only a negligible and statistically insignificant reduction when findings do not align ideologically (Republican findings).

Panels A and B in Figure 13 decompose these results by the political orientation of the respondent selected economist. For brevity, we report only the results for Selected Economists.[10] When findings suggest that abortion access improves women's outcomes, having a Republican author

[10]For Ordinary Economists and Average Americans, please consult Section 2.1 in the online appendix

significantly increases trust among all ideological groups, especially conservatives (difference = 17 p.p., p = 0.08 (w.r.t all other survey populations)). Conversely, having a Democratic author significantly reduces trust, especially among conservative respondents (difference = 7 p.p., p = 0.13 (w.r.t all other survey populations)). When findings suggest abortion access has no effect (Panel B), a Republican author significantly reduces trust across ideological lines (average reduction 30%, $p < 0.01$), while a Democratic author does not significantly enhance trust, even among liberals (p = 0.66). Thus, the effect is not perfectly symmetric: conservatives mistrust Democrat-authored studies regardless of results, whereas strong liberals increase trust significantly when Republicans report results favorable to Democratic views (p = 0.08). Overall, ideological conflicts reduce trust by an average of 17%.

In sum, ideological biases substantially influence trust independently of direct author conflicts. Republicans tend to distrust Democrat-authored results irrespective of outcomes, while Democrats selectively trust results aligned with their ideological preferences, even in the presence of potential conflicts. Consequently, ideological conflicts appear to be relatively more important to Republicans, whereas Democrats prioritize economic conflicts.

# 4 Theoretical Framework

We introduce a theoretical framework that allows us to further interpret our empirical analysis, formalizing how CoIs reduce the perceived trustworthiness of research findings and diminish the overall value of conflicted research papers. Our framework builds upon the work of Ioannidis (2005) and Maniadis et al. (2014).

At the core of the framework is the concept of post-study probability (PSP)—the probability that a research finding is true after accounting for both the statistical power of the empirical test and any biases introduced by a CoI. The PSP provides a structured way to quantify the credibility of research findings in the presence of potential conflicts. In addition, we introduce the concept of the *CoI Discount*, which quantifies the extent to which a paper's value is diminished due to the presence of a CoI. By teasing out the expected bias from the likelihood of conflicts, the CoI Discount allows us to evaluate the broader impact of CoI on the credibility and influence of academic research.

Following Ioannidis (2005) and Maniadis et al. (2014), we define the PSP as the posterior probability that a result is true after observing an (unbiased) empirical study. The PSP equals the number of *true* associations that are declared true (true positive) divided by the *total* number of associations that are declared true. Formally, we can express it as:

$$PSP = \frac{(1-\beta)\pi}{(1-\beta)\pi + \alpha(1-\pi)}, \tag{1}$$

where $\alpha$ is the typical significance level (usually $\alpha$ = 0.05), $1-\beta$ denotes the typical power of an experimental design, and $\pi$ represents the prior belief that a particular association is true. The PSP serves as a baseline measure of trust in research findings in the absence of bias.

In the marketplace of ideas, the value of a research paper is determined by its ability to shift people's priors. This shift can be either positive (e.g., a paper strengthens my prior that quantitative easing increases GDP) or negative (e.g., a paper weakens my prior that quantitative easing increases GDP). For simplicity, rather than dealing with the absolute value operator, we define the value of a paper as its ability to *increase* a prior. This is, without loss of generality, a decrease in the belief that quantitative easing increases GDP can be cast as an increase in the belief that quantitative easing does not increase GDP.

Before learning about the result, an individual's prior belief is $\pi$. After seeing the result, the posterior belief becomes the PSP. Thus, the value of a non-conflicted paper is given by:

$$\text{Value Non-Conflicted} = \frac{PSP}{\pi}. \tag{2}$$

Research characterized by a conflict of interest, in which researchers have financial, professional, or ideological stakes in the study's outcome, tend to report significant findings not due to stronger underlying effects, but as a result of subtle biases introduced at various stages of the research process. This bias emerges from a combination of factors related to study design, data collection, analysis, and presentation, which may generate significant findings even in cases where none should exist (Ioannidis, 2005). Equation (1) does not account for these potential biases and so needs to be adjusted. To address this issue, we assume that conflicted studies are $u$ percent more likely to report a significant result, even in the absence of a true underlying effect (Ioannidis, 2005; Maniadis et al.,

2014)). Hence, we modify the model to reflect this increased likelihood as follows:

$$PSP^{\text{bias}} = \frac{(1-\beta)\pi + \beta\pi u}{(1-\beta)\pi + \beta\pi u + [\alpha + (1-\alpha)u](1-\pi)}. \tag{3}$$

Equation (3) demonstrates that as research bias $u$ increases, the likelihood that participants will trust the result of a study decreases, as $PSP^{\text{bias}} \leq PSP$.

In the first question of each vignette in our survey, we did not ask about the PSP of a research finding in the absence of CoI. Rather, participants were informed upfront that the study involved potential CoI, which they took into account when forming their beliefs. Let $\lambda$ be the expected fraction of conflicted studies. Before the specific conflict status of a study is revealed to the participants, the overall prior trust in the finding can be modeled as a weighted average between the biased and the unbiased PSP values:

$$P_0 = \lambda PSP^{\text{bias}} + (1-\lambda)PSP. \tag{4}$$

$P_0$ corresponds to the first response to each vignette in our survey. Once we disclose that the study is indeed conflicted, participants' trust should shift to the biased PSP value as derived in Equation (3). The reduction in trust due to the CoI, denoted as $r$, can defined as:

$$r = 1 - \frac{PSP^{\text{bias}}}{P_0} = 1 - \frac{PSP^{\text{bias}}}{\lambda PSP^{\text{bias}} + (1-\lambda)PSP}. \tag{5}$$

$r$ corresponds to what we earlier called the CoI Trust Reduction and is the second response to each vignette. It measures participants' perceived loss of trust when informed about the specific conflict. This trust reduction combines the impact of CoI on the posterior ($PSP^{\text{bias}}$) with the expected frequency of these conflicts ($\lambda$). Ideally, we want to separate the two.

To assess the relationship between trust reduction and the overall value of a research paper, we begin by considering the value of a non-conflicted paper, as defined in Equation (2). This value reflects how much empirical evidence of a non-conflicted paper increases readers' prior beliefs. For non-conflicted research, the value is proportional to the degree to which the PSP shifts the reader's priors. However, for a conflicted paper, the presence of a conflict introduces biases that can distort

the findings. The value of a conflicted paper must therefore account for these biases, which are captured by the biased post-study probability $PSP^{\text{bias}}$. The value of a conflicted paper is thus:

$$\text{Value Conflicted} = \frac{PSP^{\text{bias}}}{\pi}. \tag{6}$$

We then calculate the value of a conflicted paper relative to that of a non-conflicted paper (the ratio of Equation (6) to Equation (2)):

$$\text{Conflicted Paper Relative Value} = \frac{\text{Value Conflicted}}{\text{Value Non-Conflicted}} = \frac{PSP^{\text{bias}}}{PSP}. \tag{7}$$

This ratio captures how much less a conflicted paper is worth compared to a paper free of conflicts. From this, we can define the CoI Discount as the reduction in value caused by the conflict:

$$\text{CoI Discount} = 1 - \frac{\text{Value Conflicted}}{\text{Value Non-Conflicted}} = 1 - \frac{PSP^{\text{bias}}}{PSP}. \tag{8}$$

The CoI Discount quantifies the proportion of paper value lost as a result of the presence of a CoI. This discount is directly related to the reduction in trust $r$. From Equation (5), we can rewrite the CoI Discount as:

$$\text{CoI Discount} = \frac{r}{1 - \lambda(1 - r)} \tag{9}$$

where $\lambda$ represents the expected frequency of conflicted papers in the relevant academic literature.

The CoI Discount starts at 0 (no discount) when there is no reduction in trust ($r = 0$), meaning that the conflicted paper is perceived as equally valuable as a non-conflicted one. Then it increases monotonically as $r$ grows, reaching a maximum at 1 (full discount) when trust is completely eroded ($r = 1$), meaning that the paper is considered worthless in terms of its ability to shift prior beliefs.

A key insight from our framework is that the CoI Discount typically exceeds the direct reduction in trust ($r$), as long as conflicted papers appear with some positive frequency ($\lambda > 0$). In other words, even small reductions in trust can substantially diminish the perceived value of research in fields where conflicts of interest are frequent. As conflicts become more common, each conflicted paper is discounted more heavily, compounding the overall loss of credibility in the literature. Our survey directly measures these trust reductions ($r$). By combining these measured reductions with

estimates of how frequently conflicts occur ($\lambda$), we can quantify the CoI Discount across various types of conflicts. We present these calculations in Section 5, following an empirical test of our model's predictions.

## 4.1 Bringing the Data to the Model

To empirically validate our theoretical model, we map responses from our survey directly onto its key variables. Specifically, the initial trust rating provided by respondents corresponds to the prior trust level ($P_0$), while the subsequent trust revision upon revealing a conflict of interest corresponds to the trust reduction ($r$), as formally defined by Equations (4) and (5), respectively. Because we elicit multiple responses from each individual across different vignettes, we can leverage within-respondent variation to estimate the relationship between our key observables ($P_0$ and $r$) while comprehensively controlling for individual fixed effects.

Our model clearly specifies the potential sources of variation across vignettes. Differences could arise from variations in the underlying quality parameters ($\alpha$ and $\beta$), prior beliefs about the likelihood of true associations ($\pi$), the severity of bias when conflicts are present ($u$), or the expected frequency of conflicted papers ($\lambda$). Given that respondents were consistently informed that each study appeared in a highly reputable journal under rigorous refereeing standards, parameters reflecting journal quality ($\alpha$ and $\beta$) and the overall frequency of conflicts ($\lambda$) are plausibly constant for a given respondent. Therefore, within-respondent variation across vignettes primarily originates from two sources: the perceived severity of disclosed conflicts ($u$), and prior beliefs about the likelihood of the result ($\pi$).

The model delivers a testable prediction regarding the correlation between initial trust ($P_0$) and trust reduction ($r$). Specifically, an increase in the severity of the conflict ($u$) reduces the probability that the study is perceived as unbiased ($PSP^{\text{bias}}$, see Eq. (3)). Through Eq. (4) and (5), this leads to a lower initial trust ($P_0$) and a larger subsequent trust reduction ($r$), thereby generating a negative correlation between these two variables. Similar logic applies if variation arises from prior beliefs ($\pi$): a stronger prior belief against an association decreases initial trust and amplifies subsequent trust reductions upon learning of a conflict. Because our vignette randomization ensures orthogonality between $u$ and $\pi$, the model predicts that $P_0$ and $r$ should be negatively correlated.

While earlier aggregate results hinted at this negative correlation (Section 3.3.1), our model formalizes this relationship explicitly. Table 7 presents a robust empirical validation of this prediction. Consistent with the model, we find a strongly significant negative correlation between initial trust ($P_0$) and trust reduction ($r$), a result robust to the inclusion of comprehensive individual fixed effects. This negative correlation holds across all respondent groups, with the magnitude strongest among average Americans, followed by ordinary economists, and smallest—but still statistically significant—among selected economists.

# 5 Quantifying the CoI Discount: Measuring the Impact of Conflicts on Research Value

Up to this point, we've shown that the presence of a CoI significantly reduces the perceived trustworthiness of a paper—typically by a one-third. We now move beyond trust reductions alone and quantify how much these conflicts diminish the overall value of research papers, a measure we refer to as the CoI Discount. As detailed in our theoretical framework (Section 4), this requires estimates of both the trust reductions elicited by conflicts and the expected frequency with which these conflicted papers appear in the literature.

Since participants completed our survey in October 2023, it seems reasonable to assume that their expectations about the prevalence of conflicts are informed by recent publication patterns. To estimate the frequency of conflicted articles, we analyze publications from 2019–2023 in two influential economics journals that were chosen also for their notably different editorial approaches toward conflicts: the *American Economic Review* (AER), known for its rigorous CoI disclosure and data management policies, and the *RAND Journal of Economics* (RAND), which explicitly states that conflicts of interest do not factor into publication decisions.

We rely on two complementary sources to identify CoIs. For the AER, we manually review 1,280 individual disclosure statements, authors' replication-package *Readme* files, and the first footnotes of papers published during our sample period. Specifically, we classify conflicts conservatively, only identifying research grants as conflicted if they come from entities with a direct stake

in the research outcome (e.g., Google on digital markets, the Kaiser Family Foundation on healthcare), thus excluding neutral sources such as the NSF or J-PAL. For data conflicts, we identify only "gated" sources—those involving explicit discretionary access—excluding widely accessible commercial databases (like Compustat) but including sources like Uber, Microsoft, Meta, and restricted government data (e.g., Census Bureau or the Federal Reserve). For RAND, due to fewer publicly available disclosures, we primarily analyze first footnotes, supplemented by the main text when footnotes are incomplete. Finally, to quantify ideological conflicts, we collect data on economists appointed to senior federal offices.[11] We then collected data on individual donations by academic economists to candidates running for federal office from Federal Election Commission records of individual donations exceeding $10,000 (cumulatively, between 2000 and 2024). Together, these two sources provide a reliable indicator of authors' political affiliations.

Table 8 presents our findings. Approximately 39% of AER articles and 29% of RAND articles exhibit at least one form of potential conflict. The most frequent conflicts involve reliance on gated data (21% for AER, 17% for RAND) and author employment with organizations directly interested in the research outcomes (16% for AER, 13% for RAND).

Armed with these estimates, Table 9 translates the trust reductions into CoI Discounts using the theoretical framework from Section 4. On average, disclosed conflicts reduce the perceived value of a document by 39%, although there is considerable variation by conflict type. Research grants from interested parties result in a discount of 35%, while consulting arrangements incur a larger penalty of 41%. When considering expert testimony, papers authored by antitrust experts who previously testified for the Department of Justice retain about 78% of the value of comparable non-conflicted papers. In contrast, papers by authors who testified for private defendants, especially when substantially compensated, retain only about 63% of their non-conflicted value.

[11] To be conservative, we considered only a restricted list of offices, such as the Council of Economic Advisors, National Economic Council, Office of Management and Budget, US Trade Representative, Department of Agriculture, Department of Commerce, Department of Labor (Bureau of Labor Statistics), , Department of Transportation, Department of the Treasury, Commodity Futures Trading Commission, Federal Reserve System, Federal Communications Commission, Federal Trade Commission, Government Accountability Office, Division of Social and Economic Sciences at the National Science Foundation, US International Trade Commission, Division of Economic and Risk Analysis at the Securities and Exchange Commission,Small Business Administration, Social Security Advisory Board, Congressional Budget Office, World Bank, and International Monetary Fund.

Career- and data-related conflicts are particularly detrimental to perceived credibility. For example, a paper on executive compensation written by someone actively pursuing a board seat retains only 53% of its original value. Similarly, papers based on proprietary data retain 76% of the value of those utilizing publicly accessible data. The penalty grows dramatically—reducing the value of a conflicted to a mere 42% of that of an unconflicted one—when the data provider retains the right to review results before publication, highlighting the critical role data transparency plays in perceived research credibility. Ideological conflicts, although somewhat less severe, still impose meaningful discounts: papers authored by prominent researchers on reputationally sensitive topics lose around 13% of their value, while politically charged research, such as studies on abortion, face discounts of 16–18%, with slightly higher penalties for authors affiliated with Republican positions.

We further examine citations—a complementary measure of research impact—for conflicted and non-conflicted articles published in these two journals (Table 10). Despite AER's rigorous disclosure policies and RAND's disregard for conflicts in editorial decisions, we observe no statistically significant differences in citations between conflicted and non-conflicted articles.

This result permits two interpretations. A pessimistic view suggests even stringent disclosure policies fail to influence how research is perceived and disseminated. A more optimistic interpretation is that editors, irrespective of formal policies, effectively manage conflicts through internal processes—by identifying potential biases, selecting impartial referees, and publishing only high-quality studies. According to this perspective, the lack of citation differences represents a success of editorial oversight rather than a failure of disclosure policy. However, our survey complicates this optimistic narrative: readers experience considerable trust erosion upon learning of conflicts. Thus, effective editorial efforts, if present, remain invisible to the broader audience, leaving a troubling disconnect. Current disclosure and editorial practices, at best, do not reassure readers sufficiently and, at worst, fail to address fundamental concerns regarding research credibility. In either case, mere disclosure appears inadequate to fully bridge the trust gap induced by CoI.

# 6 Perception and Reality

In earlier sections, we established that both economists and the broader public substantially reduce their trust in research produced by conflicted authors. Importantly, this reduction is consistent between different groups of respondents, demonstrating that it is not simply a reflection of generalized skepticism toward science. Yet, a crucial question remains unanswered: Is this diminished trust justified by actual bias in research outcomes? To address this, we next investigate whether empirical evidence of bias associated with conflicts—examined both within economics and in the medical literature, which has a longer history of studying these issues—aligns with the magnitude of trust reductions identified in our survey.

## 6.1 Bias in Research with CoI: Economics

Conflict of interest disclosure practices are relatively new in economics, especially compared to medicine, where they date back at least half a century. Among top economics journals, the *Review of Financial Studies* was the earliest adopter, instituting mandatory disclosures in 2006. The *American Economic Review* and associated AEA journals followed in 2012, prompting other leading journals to adopt similar practices soon thereafter. As a result, the literature assessing the impacts of conflicts of interest in economics is considerably smaller and more recent than its medical counterpart.

Nevertheless, existing studies document meaningful effects. Asatryan et al. (2020) find that government-funded research projects report fiscal multipliers that are, on average, 33% larger compared to projects without such funding. Similarly, Fabo et al. (2021) demonstrate that studies co-authored by central bank employees systematically estimate larger effects of quantitative easing (QE) on both output and inflation than studies authored solely by independent researchers. Notably, all papers authored by central bank affiliates find significant QE effects on output, compared to just 50% of independent studies. In quantitative terms, central-bank-affiliated papers estimate an average standardized peak effect of QE on output of 0.28, while independent studies estimate an average effect of only 0.11.

## 6.2 Bias in Research with CoI: Medicine

The medical literature has extensively documented conflicts of interest and their consequences, reflecting a much longer history of disclosure practices compared to economics. By 2006—the year the *Review of Financial Studies* first introduced mandatory CoI disclosures—approximately 90% of high-impact medical journals already required authors to disclose conflicts (Blum et al., 2009). These disclosures have revealed significant effects of industry sponsorship on research findings related to drug efficacy, medical devices (Lundh et al., 2017), clinical guidelines, advisory committee reports, and narrative reviews (Nejstgaard et al., 2020).

We systematically examine this extensive literature by conducting a structured review following established methodological guidelines. We begin by identifying a set of "seed articles" using Google Scholar and the citation network of Oostrom (2024).[12] We then leverage the PubMed API to perform a comprehensive search based on the Medical Subject Headings (MeSH) keywords extracted from these seed articles, initially retrieving 27,140 publications.

We applied a three-step filtering procedure to refine our sample. First, we excluded articles published in journals with a Scimago rank below that of the *British Medical Journal* (BMJ), reducing our sample to 4,394 articles. Second, we used GPT-4 Omni to identify articles explicitly investigating the link between industry sponsorship and research bias, narrowing the pool to 224 articles.[13] Third, we manually verified the relevance of these articles, yielding a final set of 38 studies specifically addressing the impact of industry sponsorship on research findings. Using these 38 papers as a new seed, we repeated the MeSH keyword search until no additional relevant articles emerged. This iterative approach produced a final set of 53 studies, including Oostrom (2024), which to our knowledge constitutes the largest systematic meta-analysis of CoI in medical research to date.

The 53 selected studies fall broadly into four categories: Observational Studies (ObS), Meta-

[12] See also foundational studies by Bero et al., 2007; Bourgeois et al., 2010; Cho and Bero, 1996; Kjaergard and Als-Nielsen, 2002; Knox et al., 2000; Lexchin et al., 2003.

[13] We provided GPT-4 Omni with the following prompt: "Based on the title and abstract provided, determine whether the paper investigates the relationship between industrial funding/sponsorship and bias in research outcomes. If the paper studies this correlation, return 1; otherwise, return 0."

analyses of Meta-analyses (MoM), Pharmacoeconomic Studies (Ph-E), and Randomized Clinical Trials (RCTs). Within each category, studies cover various medical fields, including oncology, cardiovascular medicine, psychiatry, and diabetes.

Figure 14 summarizes the findings, showing that conflicts of interest in medicine predominantly involve two sources: industry funding of studies or authors' direct financial stakes in a drug's success. We treat these two forms of conflict equivalently in our subsequent analysis. Specifically, Figure 14 reports odds ratios (ORs) for each category, defined as:

$$OR = \frac{O(S|C)}{O(S|NC)}, \tag{10}$$

where

$$O(S|C) = \frac{P(S|C)}{1 - P(S|C)} \tag{11}$$

is the odds of a significant finding given the study is funded by a drug manufacturer (it is "conflicted"). In contrast,

$$O(S|NC) = \frac{P(S|NC)}{1 - P(S|NC)} \tag{12}$$

is the odds of a significant finding given that the study is *not* funded by a drug manufacturer (it is not conflicted). $P(S|C)$ is the probability that a study funded by a drug manufacturer will find a significant result, and $P(S|NC)$ is the same probability for a study not funded by the industry.

For observational studies (ObS), the combined odds ratio is substantial at 22.4. Assuming a baseline 50% likelihood of significant findings in non-funded studies, this odds ratio implies industry-funded studies report significant findings with a 95.7% probability—nearly a 46 percentage-point increase.[14] For meta-analyses of meta-analyses (MoM), the odds ratio is more modest at 1.37, translating to a 58% probability of significant industry-favorable findings (an 8 percentage-point increase over non-funded studies). Pharmacoeconomic studies (Ph-E) exhibit a higher bias, with an odds ratio of 2.09, implying a 67% probability that funded studies report more favorable cost-effectiveness outcomes.

[14] To illustrate, $OR = \frac{\frac{P(S|C)}{1-P(S|C)}}{1} = 22.4$, implies $P(S|C) = 0.957$.

The bias is less pronounced, though still present, in randomized clinical trials (RCTs). Here, the combined odds ratio is 1.26, equivalent to a 56% probability of industry-favorable findings—just 6 percentage points higher than non-funded trials. This relatively lower bias in RCTs likely reflects their stricter methodological controls and mandatory pre-registration, reducing opportunities for subjective framing or selective reporting.

Upon aggregation of all studies in these categories, the overall combined odds ratio is approximately 1.39.[15] On average, industry-funded studies are thus about 8 percentage points more likely to produce results favorable to sponsors than non-funded studies, underscoring the pervasive influence of conflicts of interest in medical research.

## 6.3 Bayesian Updating

One way to further validate our results is to compare the CoI Trust Reduction in our survey with the updating of rational Bayesian subjects who know the literature. Thus, in this subsection, we apply the theoretical framework of Section 4 to the data and results from Fabo et al. (2021).

In this paper, the conflicts arise from researchers affiliated with central banks who may have incentives to produce results that support the effectiveness of QE. For our Bayesian analysis, we set the prior probability ($\pi$) that QE is effective at 0.5—reflecting the proportion of non-conflicted studies reporting significant results. We assume a bias factor ($u$) of 0.5, i.e., the difference between the fraction of significant results in conflicted studies (1) and non-conflicted ones (0.5). The frequency of conflicted studies ($\lambda$) is set at 0.6, based on the fraction of conflicted papers reported in Fabo et al., 2021. Finally, we assume a test power ($1-\beta$) of 0.8 and a statistical significance level ($\alpha$) of 0.05. Substituting these values into Equation (5), we calculate a CoI trust reduction equal to 16%.[16]

In Fabo et al., 2021, the CoI stems from public employees (researchers at central banks) who

[15] Note that odds ratios are combined using inverse-variance weighting via the metan package in Stata (Borenstein et al., 2009). This method assigns greater weight to studies with higher precision, ensuring that larger, more reliable studies have greater influence on the pooled effect size.

[16] Conducting a robustness check based on the values of $\beta$ reported by Maniadis et al. (2014), we obtain CoI Trust Reductions in the range of 16% and 18%.

may be motivated to find results that align with the preferences of their superiors. The most comparable vignette we have in the survey involves an expert witness for the DoJ with unspecified compensation, where we observe an average CoI Trust Reduction of 14.5%. This is close to what would be predicted by optimal Bayesian updating.

## 6.4 GPT Updating

Recent research in economics and computer science has explored the potential for GPT and similar large language models (LLMs) to emulate rational economic behavior (Chen et al., 2023; Horton, 2023; Kim et al., 2024; Korinek, 2023). Notably, these studies suggest that, when appropriately prompted, GPT models exhibit largely rational decision-making—even outperforming humans in areas involving risk, temporal preferences, social interactions, and consumption decisions (Chen et al., 2023). Motivated by this, we use GPT-4 Omni as an alternative benchmark to assess the robustness of our empirical findings.

We first evaluate GPT-4 Omni's CoI Trust Reductions by administering exactly the same survey (limited to the CoI vignettes) presented to our human participants. To ensure methodological consistency, GPT-4 Omni receives identical vignette texts, background assumptions, and randomizations as described previously in Section 2.1. After 1,000 simulations, GPT-4 Omni's average CoI Trust Reduction is slightly above that observed among economists but aligns closely with—and is not statistically distinguishable from—the reduction reported by average Americans (Table 11).

To further probe GPT-4 Omni's responses, we construct an additional scenario inspired by Fabo et al. (2021). Specifically, we present GPT-4 Omni with a vignette describing a paper reporting a significant impact of quantitative easing (QE) on economic output, first without mentioning any conflict. We elicit GPT-4 Omni's initial beliefs about the validity of these results using the same 5-point Likert scale from our original survey, indicating the paper's quality (i.e., published in a reputable, peer-reviewed journal). Next, we disclose to GPT-4 Omni that the study's authors were affiliated with the central bank responsible for implementing QE, subsequently eliciting the updated trust level on our 7-point Likert scale. After 1,000 simulated iterations of this exercise, GPT-4 Omni's average trust reduction is approximately 20%, corresponding to "somewhat undermines my trust," closely aligned with the 16.4% Bayesian updating benchmark calculated from the original

Fabo et al. results.

Lastly, we use GPT-4 Omni to address potential methodological concerns arising from priming effects in our original human survey (Section 2.3). The original survey used a *within-subject* design and explicitly informed participants of its CoI focus, potentially priming their responses. Ideally, to address this concern, we would have liked to administer the survey again with a *between-subjects* design. For the sample of economists, this second survey would be contaminated by the exposure to the first one. For the sample of ordinary Americans, it would have more than double the cost.

To mitigate this concern, we implement a complementary *between-subjects* design with GPT-4 Omni. In this alternative design, each GPT-simulated respondent views a single vignette, with half of the cases having no mention of conflicts of interest and the remaining half having a full disclosure of the conflicts. Responses are again measured on our 5-point Likert scale, allowing us to calculate trust reductions using Equation (5) between subjects rather than within subjects.

After 1,000 GPT-4 Omni simulations, the average trust level for the findings of a published article was 55% in the absence of any mention of conflicts and 37% when a conflict was reported. Consequently, measured between subjects, the average trust reduction caused by the presence of a conflict of interest is 32% (Table 12). This estimate is marginally larger, though statistically indistinguishable from the within-subject trust reductions of GPT-4 Omni. Thus, there is no difference in the reduction in trust measured between and within subjects. Recall that GPT-4 Omni provided a within-subject estimate virtually identical to the within-subject estimate obtained with human subjects. Thus, this evidence reassures us that our results are not specific to the methodology used.

## 6.5 Comparing Trust Reductions with Citation Impacts

As a final validation, we compare our findings to those of Leuz et al. (2023), who examine whether disclosures of industry financial ties influence the citation rates and perceived importance of medical journal articles. Their analysis uses an extensive dataset of over 17,000 research and review articles published between 1988 and 2008 in seven leading medical journals. To address potential selection bias—given that higher-quality researchers might be more likely to have industry ties—they adopt

several empirical strategies, including controlling for observable indicators of article quality and exploiting within-author variation in disclosure.

One of their key analyses involves evaluating how disclosure of industry conflicts affects an article's probability of being recommended by the University of Chicago's PURL (Priority Updates from the Research Literature) program, which identifies and disseminates influential studies to practicing family physicians. Leuz et al. (2023) find that disclosing industry ties significantly reduces the likelihood of a paper's inclusion in PURL by between 7% and 16.5%, depending on the controls employed. Assuming that PURL evaluators adjust their trust in response to CoI disclosures similarly to our survey respondents—which include expert academics—we interpret this reduction as indicative of a CoI-induced trust discount ranging from 7% to 16.5%. These figures are somewhat lower than our survey-based estimates, which set the average CoI discount for economic research around 39%.

However, several factors suggest that Leuz et al.'s (2023) results may understate the true impact of CoI disclosures. First, although the PURL recommendation process is a clever method to account for quality differences between papers, it does not perfectly control for unobservable quality attributes; non-recommended articles might still systematically differ in subtle, unmeasured ways. In contrast, our survey design directly controls the quality of the article by randomly excluding the disclosure of conflicts. Second, the decision to exclude a paper from the PURL list is publicly observable, potentially exposing evaluators to retaliation from influential conflicted authors. Anticipating this risk may cause evaluators to adopt a more cautious stance, reducing the observable effect of conflicts on recommendation decisions. Given these considerations, it is reasonable to infer that the true impact of CoI disclosures on perceived research trustworthiness could exceed the effects identified in Leuz et al.'s (2023) study.

# 7 Implications

Our findings have implications extending far beyond academia. Scholars have extensively documented how conflicted research shapes policy in critical areas, such as drug approval processes (McGarity & Wagner, 2008), and the regulation of tobacco, diesel, and alcohol, among others

(Michaels, 2020). Courts, including the US Supreme Court, have relied on sponsored economic research when issuing judicial decisions with significant social impacts (e.g., McIntire & Kantor, 2024). Moreover, companies increasingly sponsor judicial training and engage academics explicitly to influence judicial outcomes and public policies (Ash et al., 2022; Lancieri et al., 2023).

Given the policy impact of academic research, strategic attempts to influence its outcomes should not be surprising. For example, a recent congressional investigation revealed that oil companies strategically partnered with universities to shape research outcomes favorable to their interests. Internal communications from Shell, disclosed in the investigation, revealed that the company funded Imperial College London's Futures Lab as part of a "Global Methane Communications Plan" designed explicitly to generate research supporting natural gas's role in energy transitions.[17]

Such influence increasingly extends beyond financial support, often involving selective data access. Companies and public institutions, including the Federal Reserve, often grant restricted access to proprietary datasets, allowing them to shape research narratives aligned with their interests (Berg & Johnston, 2019), while simultaneously withholding data from researchers whose findings might challenge these narratives (Horan, 2019; Wagner, 2023; Zingales, 2019). For example, a whistleblower's leak of Uber's internal communications revealed that company executives deliberately controlled data sharing with academic researchers to retain editorial influence, explicitly stating that such data restrictions allowed Uber to "*work with [. . . ] on framing the study and decide what data we share*" (Lawrence, 2022).

The key concern, therefore, is not whether strategic manipulation occurs—this is already well-documented—but rather the extent to which it penetrates peer-reviewed academic research in economics, as has long been observed in medicine. Our analysis reveals that both economists and the general public perceive bias as widespread, leading to a systemic erosion of trust in scholarly research. Former US Assistant Attorney General for Antitrust, Jonathan Kanter, recently highlighted this erosion of trust, noting that conflicts of interest among antitrust scholars have blurred the line between "*expertise and advocacy*" in competition policy, fueling "*a seeping distrust of expertise by courts and law enforcers*" (Kanter, 2024). This distrust disproportionately affects

[17] Exhibit, Joint Staff Report of the Committee on Oversight and Accountability: Democrats and Senate Committee on The Budget, April 2024.

independent scholars who produce unbiased research. They receive none of the benefits (funding, data access) associated with CoIs, yet suffer from the broader erosion of scientific credibility.

What measures, then, can mitigate conflicts of interest and their biases? Our paper does not conduct a comprehensive cost-benefit analysis of conflicted research. Indeed, some conflicted relationships can benefit science—certain studies require extensive funding, proprietary datasets, or specialized knowledge available only through such collaborations. For instance, private data access can yield substantial discoveries otherwise unattainable to the scientific community (e.g., González-Bailón et al., 2023).

Nevertheless, even without a full cost-benefit evaluation, our empirical results and analytical framework offer concrete avenues for reducing the negative consequences of conflicts of interest on the credibility of academic research.

## 7.1 Credibility of Disclosure

Our findings suggest that detailed and credible disclosure of CoI can significantly reduce—and in some cases entirely eliminate—the implicit cross-subsidization of conflicted papers by non-conflicted ones. For example, in our credit-card vignette, respondents were initially unaware of the conditions under which the company provided data to researchers. Revealing that the company waived its review rights increased respondents' trust by 17%; conversely, disclosing that the company retained review rights reduced trust by 52%. Thus, the lack of precise disclosure artificially inflates trust in conflicted papers by implicitly drawing credibility from non-conflicted research. Indeed, conflicted parties themselves recognize this dynamic. A recent Wall Street Journal investigation revealed instances of authors deliberately withholding conflict disclosures to maintain their credibility and policy influence (Mullins, 2024).

In theory, non-conflicted researchers could deviate from this equilibrium by explicitly disclosing their lack of conflicts (e.g., indicating no review rights for data providers). However, this strategy is effective only if disclosures are perceived as credible. To assess this, we surveyed economists on two questions. First, we asked, "Do you think your past disclosure of potential conflicts of interest has been adequate?" An overwhelming 95% responded affirmatively ("Yes" or "Mostly"; Figure 15).

Yet, when asked, "Do you think academic economists generally disclose conflicts adequately?" only 2% answered "Yes," with 43% selecting "Mostly" and the remaining responses significantly more skeptical or unsure. This stark contrast highlights an inherent credibility problem: economists trust their own disclosures, yet harbor doubts about the disclosures of their peers. Such skepticism is likely exacerbated by prominent cases of inadequate disclosure, illustrating the negative externality conflicted research imposes on the broader academic community.

We attribute this credibility gap primarily to a lack of enforcement mechanisms. Penalties for inadequate disclosure are rare and seldom publicized, allowing conflicted papers to persistently impose externalities on the credibility of non-conflicted work. To address this, journals and academic institutions must develop credible and enforceable standards for conflict disclosure, ensuring transparency and integrity across research practices.

## 7.2 Transparency in Data Access

The terms under which researchers gain access to proprietary data are frequently opaque. Although some institutions and journals formally prohibit data providers from retaining rights to censor publication, adherence to these rules remains unclear. Agreements governing data access are typically confidential—even to researchers within the same institution—creating opportunities for selective data sharing or manipulation of research outcomes.

Confidential data-access agreements pose a direct threat to research transparency and replicability. We propose that academic journals explicitly refuse submissions from authors whose data providers retain review or approval rights beyond necessary confidentiality protections. Furthermore, journals should require public disclosure of data-access terms (excluding potentially sensitive payment details) as a prerequisite for publication. Such transparency would bolster researchers' negotiating power with data providers and strengthen the overall integrity of research.

## 7.3 Prospective Conflicts

Our paper also highlights an under-recognized yet critical form of conflict: future career prospects. These prospective conflicts are especially problematic as they are difficult to identify and disclose,

particularly when potential employers are concentrated within a few influential institutions. As Kempf (2020) demonstrate, career incentives do not necessarily distort behavior when multiple potential employers observe a candidate's past actions. However, when opportunities are limited, career pressures can strongly bias research outcomes. Such dynamics underscore a hidden cost of industrial concentration: fewer potential employers inevitably narrow the diversity of research outcomes, ultimately impoverishing the marketplace of ideas.

## 7.4 Discounting Publications

Finally, our results demonstrate that conflicted research has a systematically lower value, roughly 61% that of comparable non-conflicted studies. Yet, under current practices, academics fully internalize benefits from conflicted relationships (funding, proprietary data, prestige) without bearing the broader societal costs stemming from reduced public trust. This creates asymmetric incentives and negative externalities that disproportionately harm independent and impartial scholars.

Despite this imbalance, in literature reviews, legal settings, and academic promotions, papers are currently evaluated based mostly on the prestige of the journals and citation count. Our results show that CoI generates distrust even in papers published in prestigious journals and that citations do not fully capture the level of distrust stemming from different conflicts (see Table 10). As a result, conflicted papers and their authors receive the same rewards as those awarded to non-conflicted papers, despite having a lower social value (everything else being equal). This misalignment can lead to an overproduction of conflicted papers from a social perspective.

We suggest a straightforward solution: explicitly discounting the value of conflicted publications in literature reviews, legal contexts, and academic evaluations. Such a policy would align private academic incentives more closely with public interest, reducing conflicted research overproduction. Our findings offer an initial benchmark for effectively implementing such discounting. Nonetheless, more work is needed to refine discount factors and ensure their practical applicability.

# 8 Conclusion

In this paper, we examine how various types of conflicts of interest affect the perceived trustworthiness and value of economic research. Our survey provides clear evidence that the general public and expert economists discount the credibility of conflicted research, even when published in prestigious journals. By formalizing the relationship between trust and value in the presence of conflicts, we quantify this credibility penalty—the CoI Discount—which represents the broader social cost of conflicted research.

This discount might be interpreted as evidence that the market properly discounts conflicted research. This interpretation, however, is incorrect for several reasons. First, inadequate and non-credible disclosure allows conflicted research to impose negative externalities on non-conflicted scholarship, eroding trust across entire fields. Second, cognitive dissonance limits rational discounting precisely among the experts most likely to evaluate research critically. Third, conflicted papers are not discounted in practical settings, such as academic promotions, policy making, and literature reviews. Since we have shown that citations do not properly reflect conflicts, even weighting conflicted papers by citations will not adequately discount conflicted research, resulting in an inflated value in the marketplace of ideas. Given the growing reliance on evidence-based policy (Haskins, 2018), such distortions are likely to have substantial welfare implications.

Moreover, our estimates of social costs assume fully rational and informed actors, underestimating the true societal costs of conflicted research. Conflicted research often deliberately aims to mislead policy makers and the public (Michaels, 2008), suggesting a greater social impact. Future research should aim to quantify these larger costs, providing deeper insight into how CoI affect trust, policy outcomes, and societal welfare.

Finally, conflicts of interest represent only one type of distortion within academia and the broader marketplace of ideas. Identifying and addressing these biases remains complex but critical. If left unresolved, distortions like these will continue to erode public and expert trust in academic research. Ultimately, our goal is to spark a sustained and open conversation about conflicts of interest that motivates scholars, institutions, and society to mitigate these biases effectively, reinforcing the credibility of science.

# References


Asatryan, Z., Havlik, A., Heinemann, F., & Nover, J. (2020). "Biases in fiscal multiplier estimates". *European Journal of Political Economy*, *63*.

Ash, E., Chen, D., & Naidu, S. (2022). "Ideas have consequences: The impact of law and economics on american justice". *NBER Working Paper 29788*.

Barrios, J., Lancieri, F., Levy, J., Singh, S., Valletti, T., & Zingales, L. (2023). "Norms and practices in the disclosure of conflicts of interest: Evidence from the field of economics". *In : AEA RCT Registry*.

Berg, J., & Johnston, H. (2019). "Too Good to Be True? A Comment on Hall and Krueger's Analysis of the Labor Market for Uber's Driver-Partners". *ILR Review*, *72*(1), 39–68.

Bero, Oostvogel, F., Bacchetti, P., & Lee, K. (2007). "Factors associated with findings of published trials of drug-drug comparisons: Why some statins appear more efficacious than others". *PLoS Medicine*, *4*(6), 1001–1010.

Blum, J. A., Freeman, K., Dart, R. C., & Cooper, R. J. (2009). "Requirements and definitions in conflict of interest policies of medical journals". *JAMA*, *302*(20), 2230–2234.

Borenstein, M., Hedges, L., Higgins, J., & Rothstein, H. (2009). "Chapter 13. fixed-effect versus random-effects models". In *Introduction to meta-analysis*. John Wiley  Sons, Ltd.

Borjas, G. J., & Breznau, N. (2024, December). *Ideological bias in estimates of the impact of immigration* (Working Paper No. 33274). National Bureau of Economic Research.

Bourgeois, F. T., Murthy, S., & Mandl, K. D. (2010). "Outcome reporting among drug trials registered in clinicaltrials.gov". *Annals of Internal Medicine*, *153*(3), 158–166.

Chen, Y., Liu, T. X., Shan, Y., & Zhong, S. (2023). "The emergence of economic rationality of gpt". *Proceedings of the National Academy of Sciences*, *120*(51).

Cho, M. K., & Bero, L. A. (1996). "The quality of drug studies published in symposium proceedings". *Annals of Internal Medicine*, *124*(5), 485–489.

Christensen, G., & Miguel, E. (2018). "Transparency, reproducibility, and the credibility of economics research". *Journal of Economic Literature*, *56*(3), 920–80.

Dawson, E., Gilovich, T., & Regan, D. T. (2002). "Motivated reasoning and performance on the was on selection task". *Personality and Social Psychology Bulletin*, *28*(10), 1379–1387.

De Long, J. B., & Lang, K. (1992). "Are all economic hypotheses false?" *Journal of Political Economy*, *100*(6), 1257–1272.

Edelson, L., Graef, I., & Lancieri, F. (2023). "Access to Data and Algorithms: For an Effective DMA and DSA Implementation".

Fabo, B., Jančoková, M., Kempf, E., & Pástor, Ľ. (2021). "Fifty shades of qe: Comparing findings of central bankers and academics". *Journal of Monetary Economics*, *120*, 1–20.

Gallup Inc. (2023, July). "Americans' Confidence in Higher Education Down Sharply".

González-Bailón, S., Lazer, D., Barberá, P., Zhang, M., Allcott, H., Brown, T., Crespo-Tenorio, A., Freelon, D., Gentzkow, M., & Guess, A. M. (2023). "Asymmetric ideological segregation in exposure to political news on Facebook". *Science*, *381*(6656), 392–398.

Hampson, L. A., Agrawal, M., Joffe, S., Gross, C. P., Verter, J., & Emanuel, E. J. (2006). "Patients' views on financial conflicts of interest in cancer research trials". *The New England Journal of Medicine*, *355*(22), 2330–2337.

Haskins, R. (2018). "Evidence-based policy: The movement, the goals, the issues, the promise". *The ANNALS of the American Academy of Political and Social Science*, *678*(1), 8–37.

Horan, H. (2019). "Uber's Academic Research Program: How to Use Famous Economists to Spread Corporate Narratives". *ProMarket*.

Horton, J. J. (2023). "Large language models as simulated economic agents: What can we learn from homo silicus?" *National Bureau of Economic Research*, (Working Paper 31122).

Ioannidis, J. P. A. (2005). "Why most published research findings are false". *PLoS Medicine*, *2*(8), 696–701.

Ioannidis, J. P. A., Stanley, T. D., & Doucouliagos, H. (2017). "The Power of Bias in Economics Research". *The Economic Journal*, *127*(605), 236–265.

Javdani, M., & Chang, H.-J. (2023). "Who said or what said? Estimating ideological bias in views among economists". *Cambridge Journal of Economics*, *47*(2), 309–339.

Jelveh, Z., Kogut, B., & Naidu, S. (2024). "Political Language in Economics". *The Economic Journal*, *134*(662), 2439–2469.

Kanter, J. (2024). "Assistant Attorney General Jonathan Kanter Delivers Remarks for the Fordham Competition Law Institute's 51st Annual Conference on International Antitrust Law and Policy". *Deparment of Justice*.

Kempf, E. (2020). "The job rating game: Revolving doors and analyst incentives". *Journal of Financial Economics*, *135*(1), 41–67.

Kesselheim, A. S., Robertson, C. T., Myers, J. A., Rose, S. L., Gillet, V., Ross, K. M., Glynn, R. J., Joffe, S., & Avorn, J. (2012). "A randomized study of how physicians interpret research funding disclosures". *New England Journal of Medicine*, *367*(12), 1119–1127.

Kim, A. G., Muhn, M., & Nikolaev, V. V. (2024). "Financial statement analysis with large language models". *Chicago Booth Research Paper, Fama-Miller Working Paper Series*.

Kjaergard, L., & Als-Nielsen, B. (2002). "Association between competing interests and authors' conclusions: Epidemiological study of randomised clinical trials published in the bmj". *BMJ*, *325*(7358).

Knox, K. S., Adams, J. R., Djulbegovic, B., Stinson, T. J., Tomor, C., & Bennet, C. L. (2000). "Reporting and dissemination of industry versus non-profit sponsored economic analyses of six novel drugs used in oncology". *Annals of Oncology*, *11*(12), 1591–1595.

Korinek, A. (2023). "Generative ai for economic research: Use cases and implications for economists". *Journal of Economic Literature*, *61*(4), 1281–1317.

Lancieri, F., Posner, E. A., & Zingales, L. (2023). "The Political Economy of the Decline in Antitrust Enforcement in the United States". *Antitrust Law Journal*, *85*(2).

Lawrence, F. (2022). "Uber Paid Academics Six-Figure Sums for Research to Feed to the Media". *The Guardian*.

Leuz, C., Malani, A., Muhn, M., & Jakab, L. (2023). "Do conflict of interests disclosures work? evidence from citations in medical journals". *National Bureau of Economic Research*, (Working Paper 30927).

Lexchin, J., Bero, Djulbegovic, B., & Clark, O. (2003). "Pharmaceutical industry sponsorship and research outcome and quality: Systematic review". *BMJ*, *326*(7400), 1167–1170.

Lianos, I. (2024). "Academic "Capture"? The Hidden Costs of Corporate Funding in Competition Policy Research and Proposed Remedies". *Working Paper*.

Lundh, A., Lexchin, J., Mintzes, B., Schroll, J., & Bero, L. (2017). "Industry sponsorship and research outcome". *Cochrane Database of Systematic Reviews*, *2*.

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.

McGarity, T. O., & Wagner, W. E. (2008). *Bending science: How special interests corrupt public health research*. Harvard University Press.

McIntire, M., & Kantor, J. (2024). "The Gun Lobby's Hidden Hand in the 2nd Amendment Battle". *The New York Times*.

Michaels, D. (2008). *Doubt is their product: How industry's assault on science threatens your health*. Oxford University Press.

Michaels, D. (2020). *The triumph of doubt: Dark money and the science of deception*. Oxford University Press.

Mills, M. A. (2023). "Opinion — Why So Many Americans Are Losing Trust in Science". *The New York Times*.

Mullins, B. (2024). "The Hidden Life of Google's Secret Weapon". *The Wall Street Journal*.

Nejstgaard, C. H., Bero, L., Hróbjartsson, A., Jørgensen, A. W., Jørgensen, K. J., Le, M., & Lundh, A. (2020). "Association between conflicts of interest and favourable recommendations in clinical guidelines, advisory committee reports, opinion pieces, and narrative reviews: Systematic review". *BMJ*, *371*.

Nevo, A. (2024, November). "Disclosures and Accountability: Helping Economic Research Impact Policy".

Oostrom, T. (2024). "Funding of clinical trials and reported drug efficacy". *Journal of Political Economy*, *Forthcoming*.

Rosenbluth, T. (2024). "Nobel Laureates Urge Senate to Turn Down Kennedy's Nomination". *The New York Times*.

Saint-Paul, G. (2012). "Toward a political economy of macroeconomic thinking". *NBER International Seminar on Macroeconomics*, *8*(1), 249–284.

Snyder, P. J., Papp, K. V., Bartkowiak, J., Jackson, C. E., & Doody, R. S. (2009). "Commentary on "a roadmap for the prevention of dementia ii. leon thal symposium 2008." recruitment of participants for alzheimer's disease clinical trials: The role of trust in caregivers, clinical researchers, regulatory authorities, and industry sponsors". *Alzheimer's & Dementia*, *5*(2), 122–124.

Stantcheva, S. (2023). "How to run surveys: A guide to creating your own identifying variation and revealing the invisible". *Annual Review of Economics*, *15*, 205–234.

Valletti, T. (2024). "'Doubt is their product'—The difference between research and academic lobbying". *Public Money & Management*.

Wagner, M. W. (2023). "Independence by permission". *Science*, *381*(6656), 388–391.

Zingales, L. (2013). "Preventing economists' capture". In D. Carpenter & D. Moss (Eds.), *Preventing regulatory capture: Special interest influence and how to limit it*. The Tobin Project.

Zingales, L. (2019). "Uber and the Sherlock Holmes Principle: How Control of Data Can Lead to Biased Academic Research". *ProMarket*.

# Figures

Figure 1: **Selection Bias : Comparison of PhD Years**

Figure 1 compares the PhD years of all selected economists in our sample vs. the selected economists who responded to our survey. By selected economists, we mean the economists affiliated to NBER, CEPR and/or Kent A. Clark Center for Global Markets.

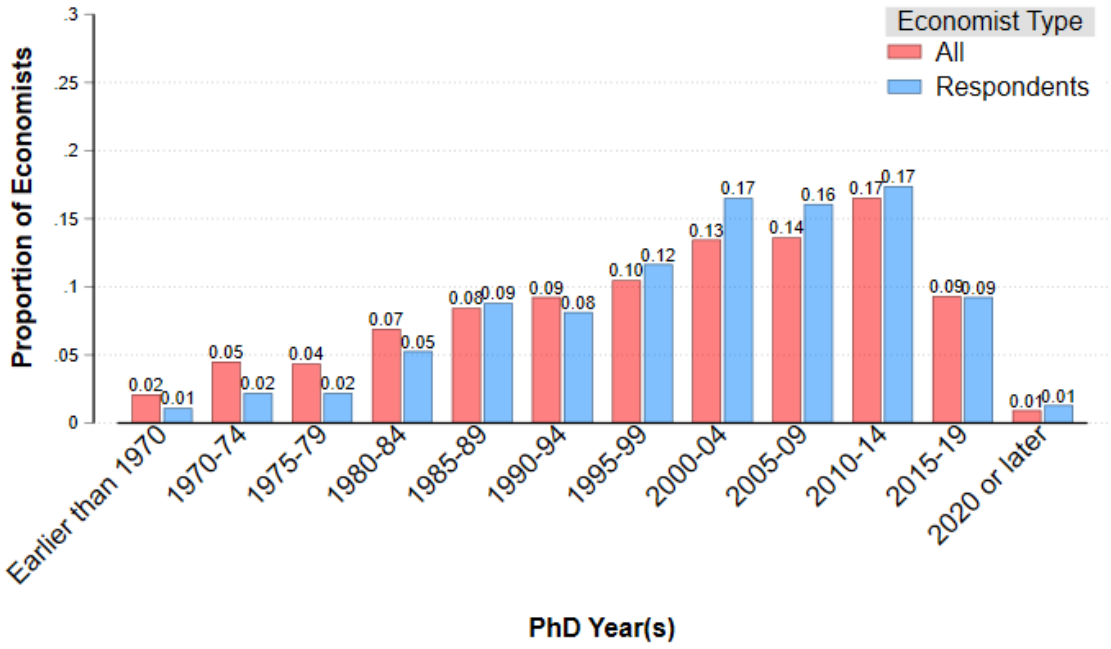


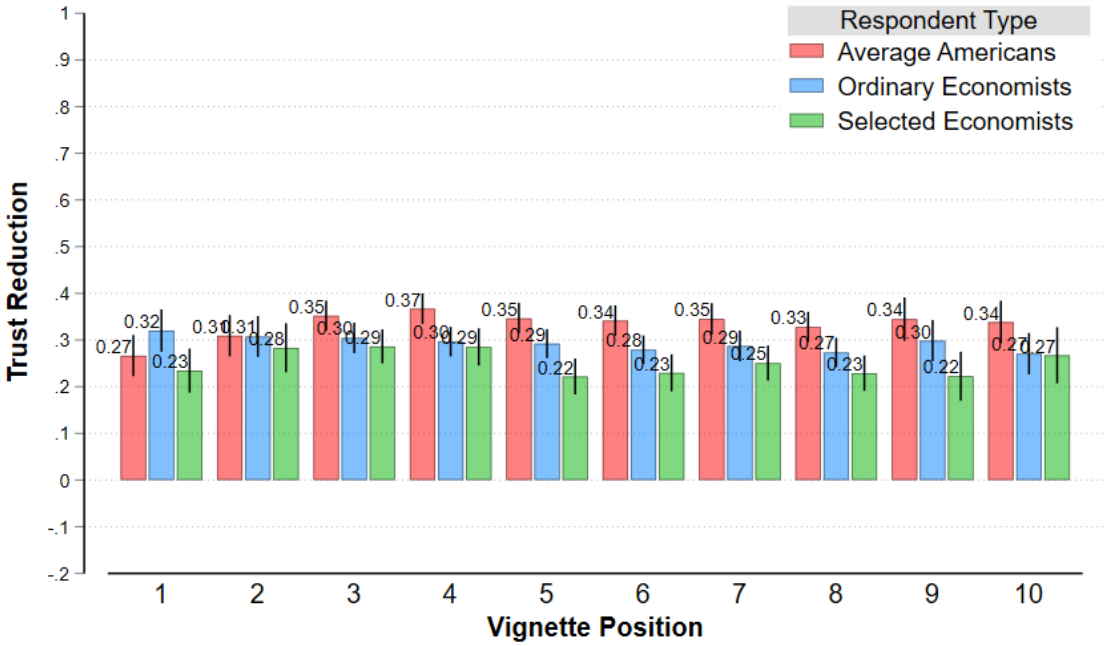
Figure 2: **Priming Bias Conditional on Position of the Vignette**

Figure 2 Panel A shows the average prior trust contingent on vignette positioning in the survey. Panel B shows average trust reduction contingent on vignette positioning. There are 10 possible positions in the survey where a vignette can appear since there are 9 CoI vignettes and 1 attention test vignette.

A. Prior Trust



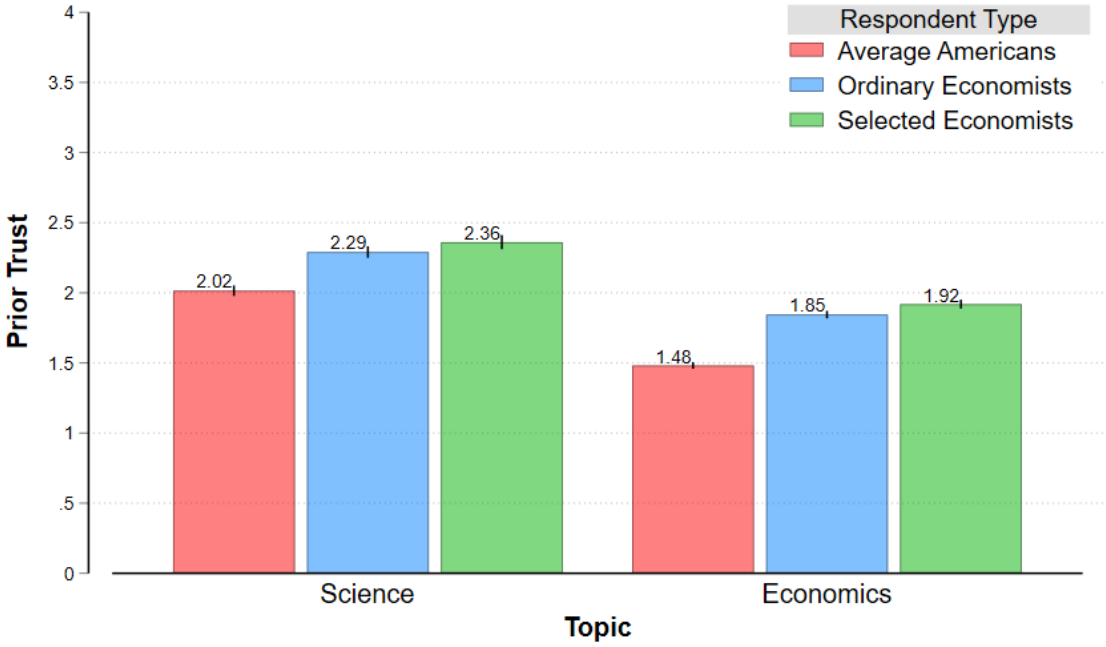

B. CoI Trust Reduction



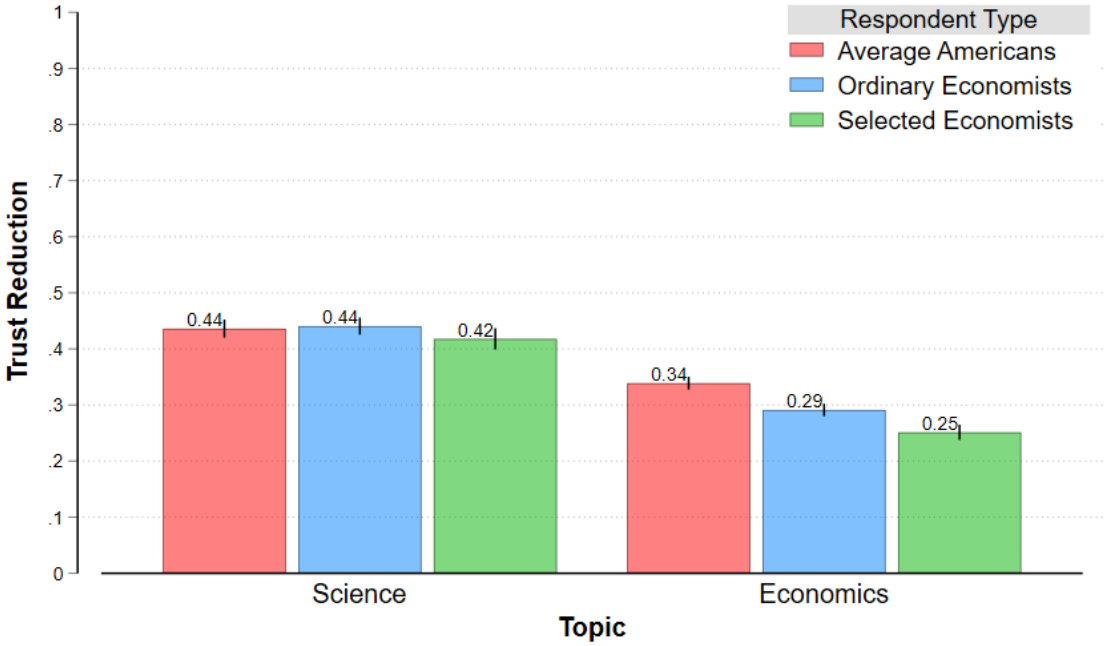
## Figure 3: **Prior Trust and CoI Trust Reduction Across Topics and Respondent Types**

Figure 3, Panel A illustrates respondents' pre-CoI disclosure trust levels (Prior Trust) in the paper's findings, aggregated across all vignettes, for each respondent type. The science vignette category includes Industrial Disaster and Alzheimer's Drug scenarios. The economics vignette category encompasses Trading Strategy, Underpaid CEOs, Supermarket Merger, Rideshare, Credit Card, Abortion, and Tax Policy scenarios. Trust levels were measured using a 5-point Likert scale: 0: Not at All, 1: A little, 2: Moderately, 3: Seriously, 4: Completely. Panel B illustrates respondents' post-CoI disclosure trust reduction in the paper's findings, aggregated across all vignettes. Respondents reported changes in trust using a 7-point Likert scale: 100% Decrease: It completely makes me distrust the results, 50% Decrease: It seriously undermines my trust in the results, 20% Decrease: It somewhat undermines my trust in the results, 0%: It does not impact my trust in the results, 20% Increase: It somewhat increases my trust in the results, 50% Increase: It seriously increases my trust in the results, 100% Increase: It completely makes me trust the results.

A. Prior Trust




B. CoI Trust Reduction



## Figure 4: **CoI Trust Reduction v/s Prior Trust**

Figure 4 shows the CoI Trust Reduction plotted against the Prior Trust levels across all vignettes and respondent types. The size of the circle is directly proportional to the number of observations exhibiting the corresponding level of Prior Trust and CoI Trust Reduction. Trust levels were measured using a 5-point Likert scale: 0: Not at All, 1: A little, 2: Moderately, 3: Seriously, 4: Completely. Panel B illustrates respondents' post-CoI disclosure trust reduction in the paper's findings, aggregated across all vignettes. Respondents reported changes in trust using a 7-point Likert scale: 100% Decrease: It completely makes me distrust the results, 50% Decrease: It seriously undermines my trust in the results, 20% Decrease: It somewhat undermines my trust in the results, 0%: It does not impact my trust in the results, 20% Increase: It somewhat increases my trust in the results, 50% Increase: It seriously increases my trust in the results, 100% Increase: It completely makes me trust the results. Prior Trust levels were measured using a 5-point Likert scale: 0: Not at All, 1: A little, 2: Moderately, 3: Seriously, 4: Completely.

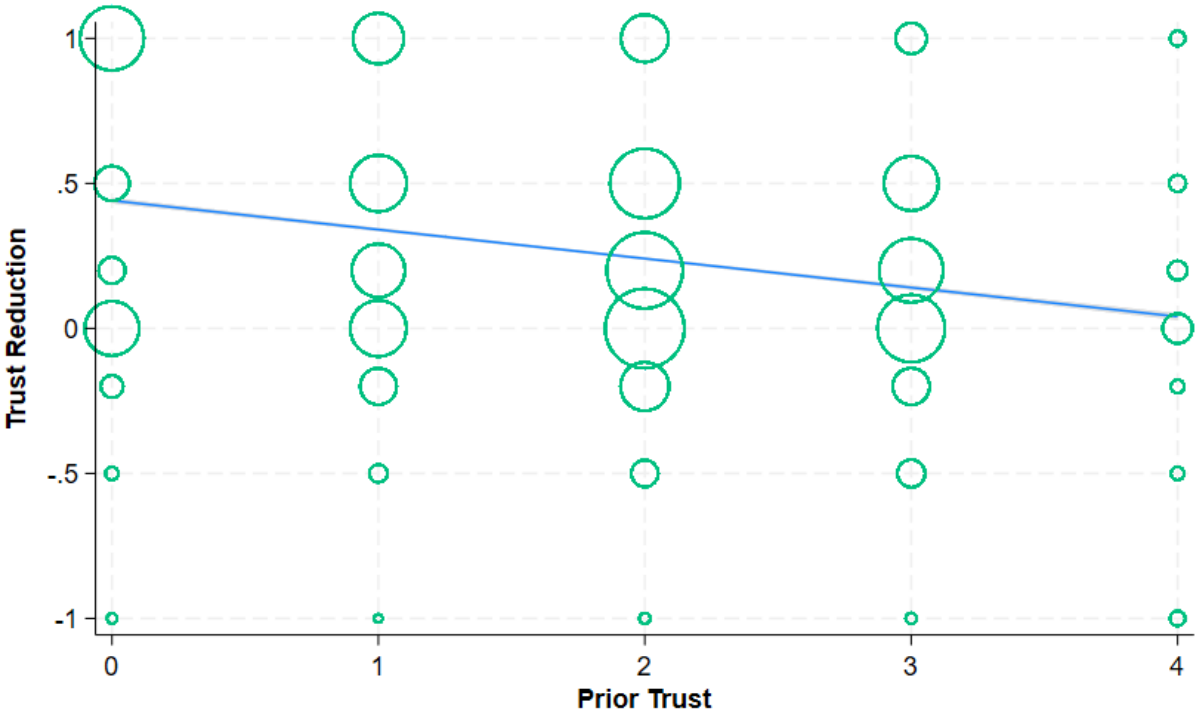
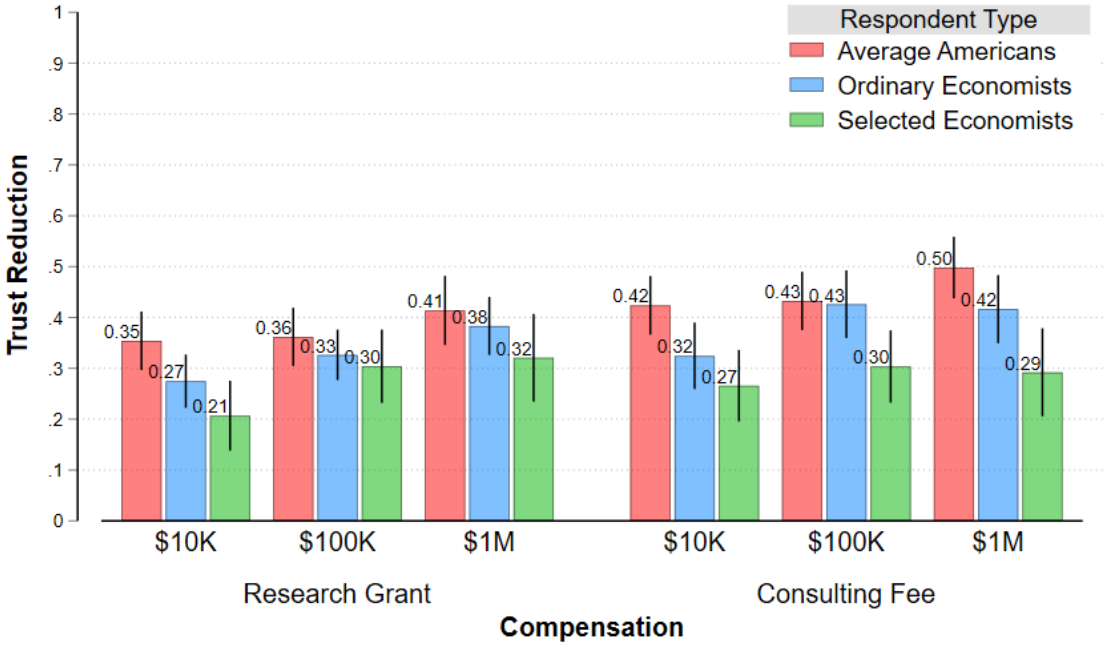


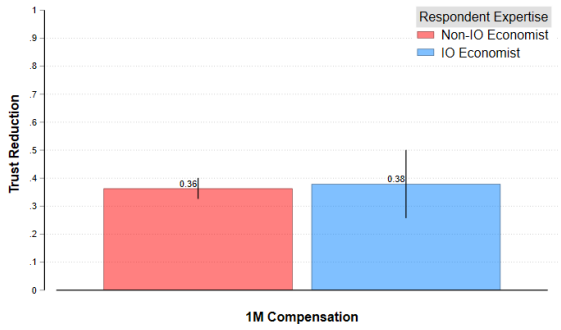
Figure 5: **CoI Trust Reduction in the Presence of Monetary Incentives: Trading Strategy**

Figure 5, Panel A explores the monetary variation in the Trading Strategy vignette across all respondent types. In the Trading Strategy vignette, we reveal to the respondent that the author of a paper on financial trading strategies, which utilize investors' reactions to 'news' related to companies with similar names, found such strategies useful for making money. As part of the disclosure statement, we further reveal that the author was paid $10,000/$100,000/$1,000,000 in the form of a research grant or consulting fee. Panel B decomposes the results for economists by their self reported expertise. We only focus on IO and Finance experts.

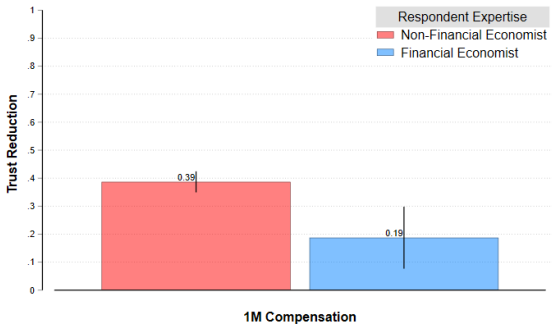
A. All Three Samples




B1. Only IO Economists



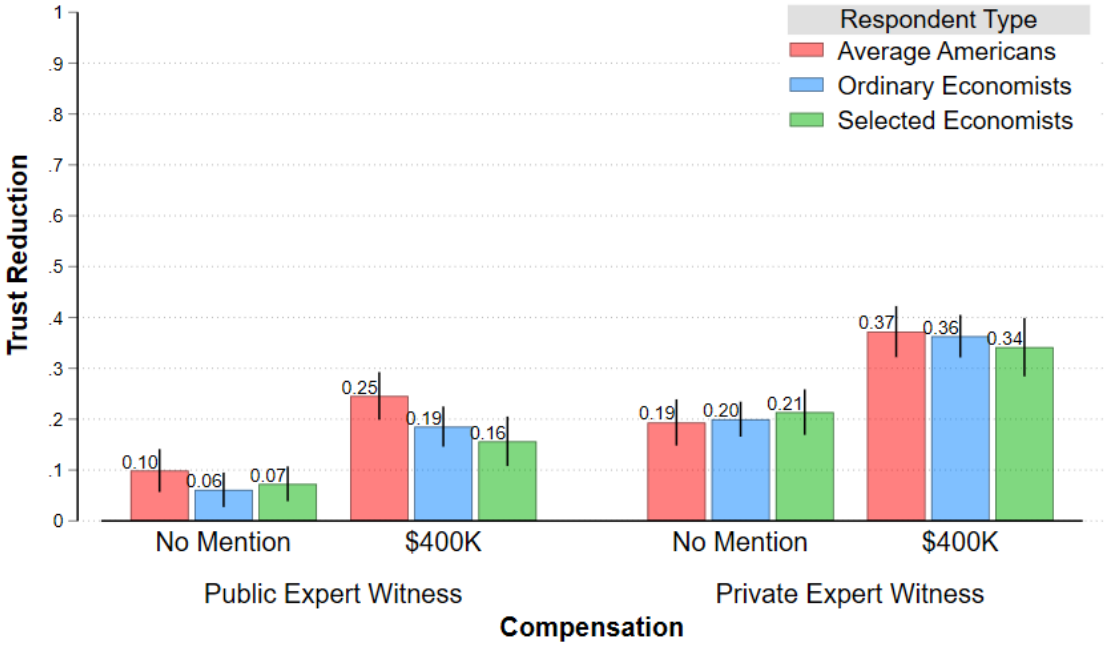

B2. Only Financial Economists



## Figure 6: **CoI Trust Reduction in the Presence of Monetary Incentives: Supermarket Vignette**

Figure 6, Panel A shows the average CoI trust reduction in the Supermarket Merger vignette concerning monetary incentives across respondent types. In the Supermarket Merger vignette, the CoI disclosure reveals to the respondent that the author of the paper, who studied how the growth of supermarket chains affects consumer welfare and found that larger chains are associated with increased (decreased) consumer welfare due to lower (higher) prices, was an expert witness either for the Department of Justice Antitrust Division or for a supermarket chain. Additionally, we inform the respondent that the author was either compensated $400,000 for the testimony or we do not mention anything related to compensation. Panel B explores the monetary variation in the Supermarket Vignette across economists' self-reported expertise/specialization. We only focus on IO and Finance experts.

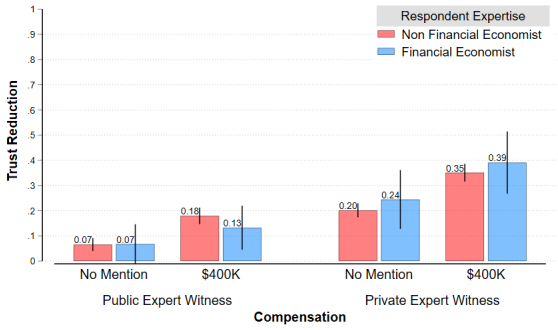
A. All Three Samples




B1. Only IO Economists




B2. Only Financial Economists



## Figure 7: **CoI Trust Reduction in the Presence of Career Incentives**

Figure 7 shows the CoI trust reduction with respect to the Underpaid CEOs vignette, across all respondent types. In the Underpaid CEOs vignette, we first reveal to the respondent that the author of a paper found that CEOs are underpaid relative to the value of the services they render to their companies. We then disclose the CoI that the author was seeking a position on the board of directors of a large US public company before writing the paper, was immediately nominated by the management of a large US public company to its board of directors, or was nominated by the management of a large US public company to its board of directors two years following the publication of the paper.

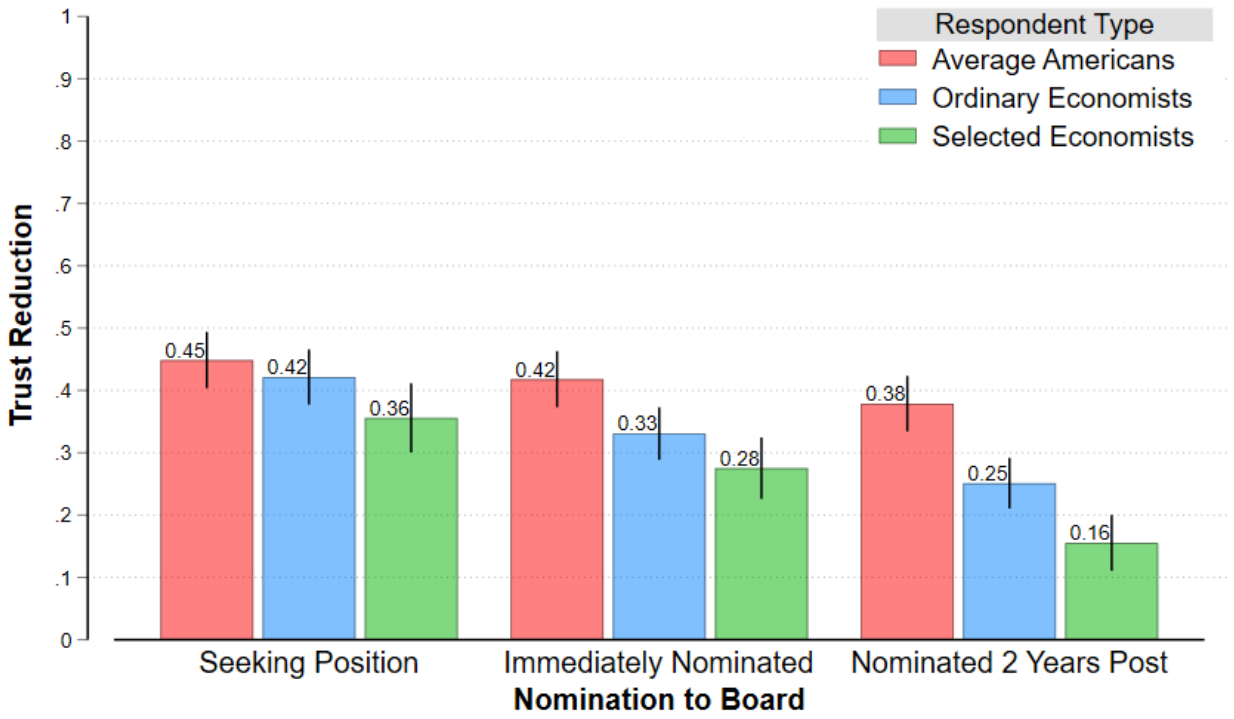


Figure 8: **CoI Trust Reduction in the Presence of Data Restrictions : Ridesharing**

Figure 8 shows the CoI trust reduction related to data incentives from the rideshare vignette, across all respondent types. In the rideshare vignette, we first inform the respondent about the findings of a paper on the benefits of ridesharing services. One randomization mentions that the benefits of ridesharing services outweigh the costs, while the other mentions that the costs outweigh the benefits. We then inform the respondent about the CoI disclosure, revealing the source of the data used in the paper. In one randomization, the data was proprietary and provided by the ridesharing company. In another randomization, the data was publicly available administrative data provided by the city government.

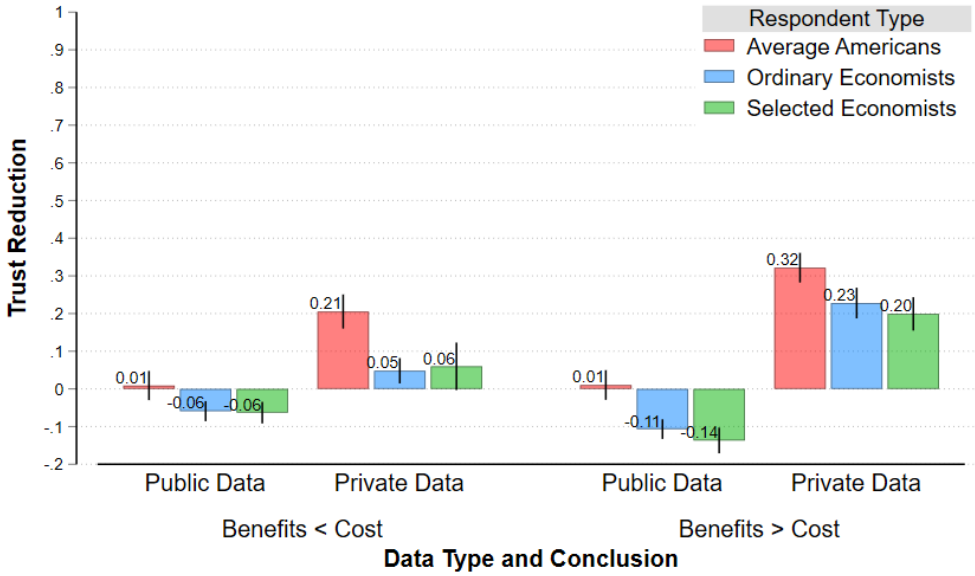
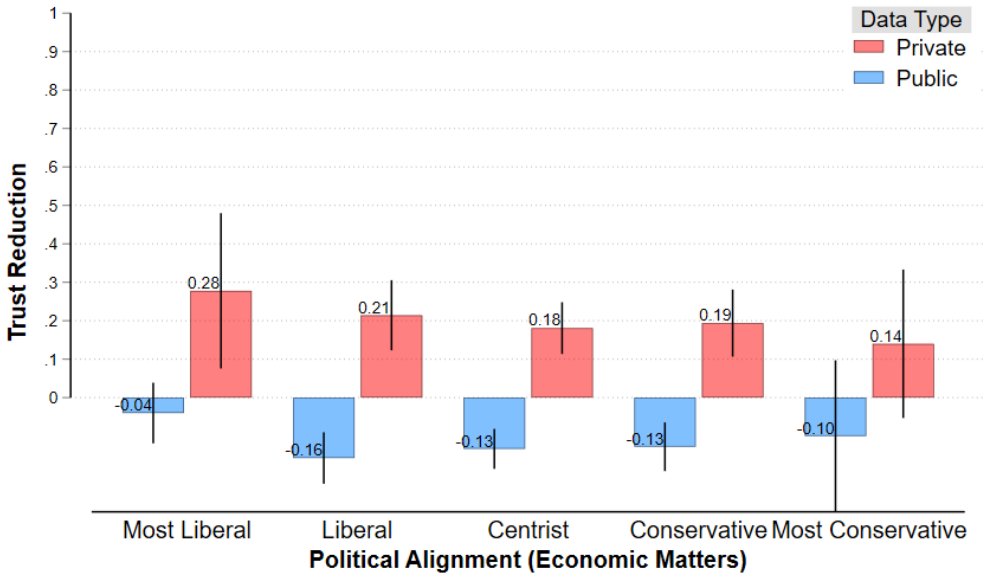


Figure 9: **Political Alignment and CoI Trust Reduction: Ridesharing x Selected Economists**

Figure 9, Panel A illustrates the CoI trust reduction in the ridesharing vignette (see Figure 8) across the political alignment of selected economist respondents on economic matters. Panel B shows the CoI trust reduction across the political alignment of selected economist respondents. Political alignment on economic matters is coded between *most liberal* (e.g. High taxes and government intervention) and *most conservative* (e.g. Low taxes and government intervention).

A. Selected Economists: Data Type X Benefits Greater than Costs Finding




B.Selected Economists: By Paper's Findings

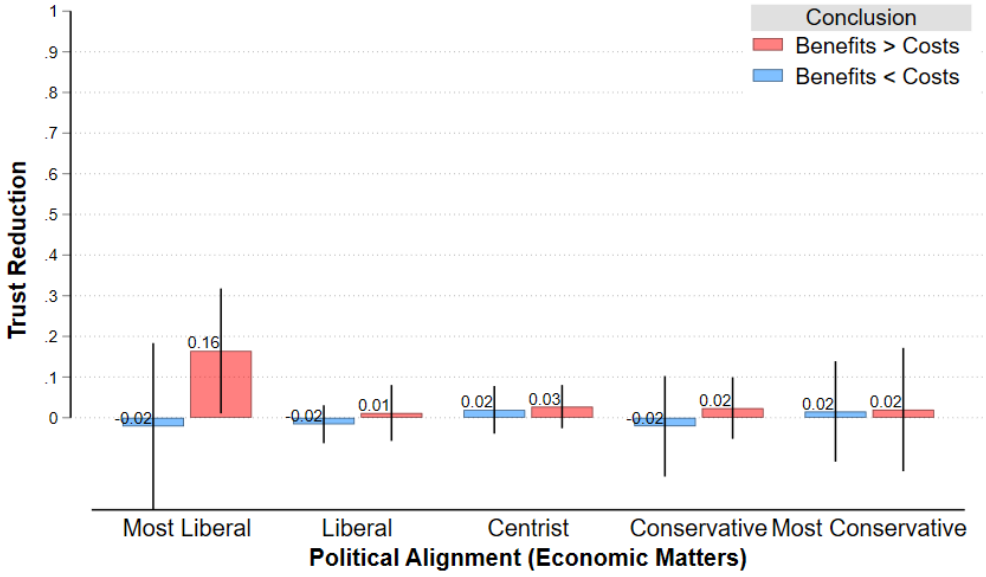


Figure 10: **CoI Trust Reduction in the Presence of Data Restrictions: Credit Card**

Figure 10, exploits the right-to-review randomization of the Credit Card vignette to analyze the average CoI trust reduction across respondent types. In the Credit Card vignette, we first tell the respondents about the findings of a paper that shows that a Credit Card's business practices are efficient and that the Credit Card company provided the data. In the disclosure treatment, we mention that the Credit Card company had the right to review the paper and control its findings. In another randomization of the disclosure, we mention that Credit Card company waived the review rights.

A. Credit Card

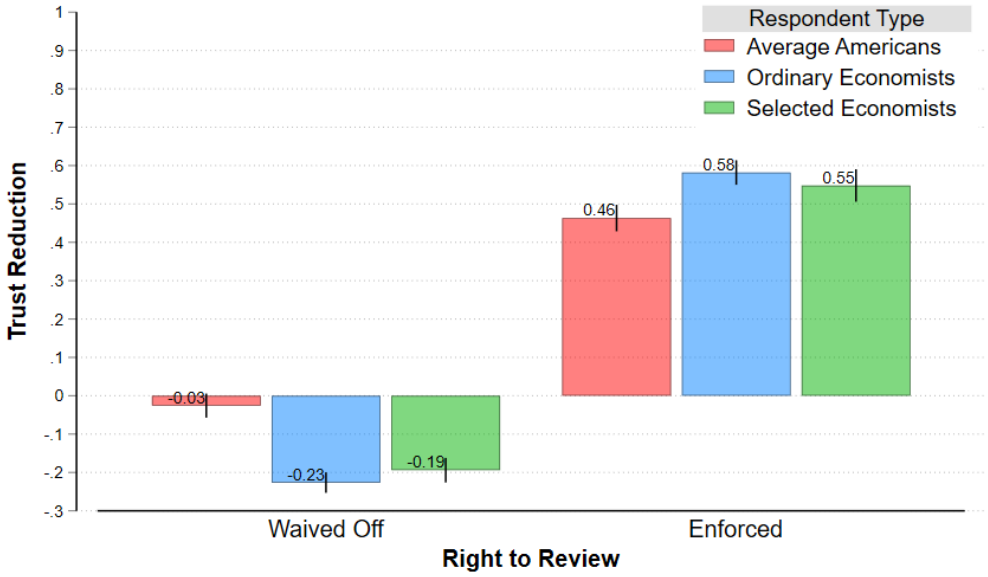
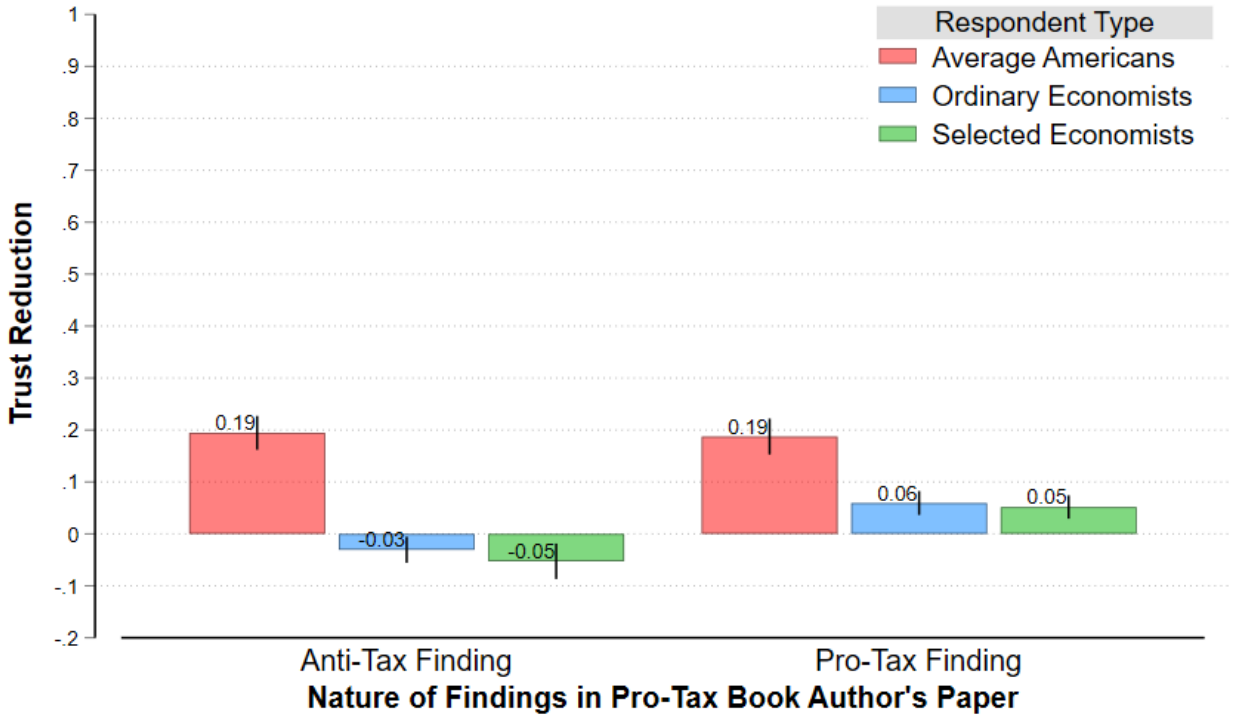


Figure 11: **CoI Trust Reduction in the Presence of Academic Conflicts: Tax Policy**

Figure 11 plots the CoI trust reduction for the Tax Policy vignette across all respondent types. In this vignette, we first inform the respondents about the findings of a paper, which can randomly be pro-tax (higher income tax is beneficial for economic growth) or anti-tax (higher income tax is not beneficial for economic growth). We then reveal to the respondents that the author of the paper wrote a pro-tax book in the past and ask for their change in belief in the findings of the paper.



## Figure 12: **CoI Trust Reduction in the Presence of Political Conflicts**

Figure 12 shows the CoI trust reduction for the Abortion vignette, across different respondent types. In this scenario, respondents are initially presented with the findings of a study that includes two randomizations: (i) Access to abortion increases the lifetime earnings of women, and (ii) Access to abortion has no impact on the lifetime earnings of women. Subsequently, the political affiliation of the study's author is revealed to the respondents, with the affiliation randomly assigned as either Democrat or Republican, regardless of the study's findings.

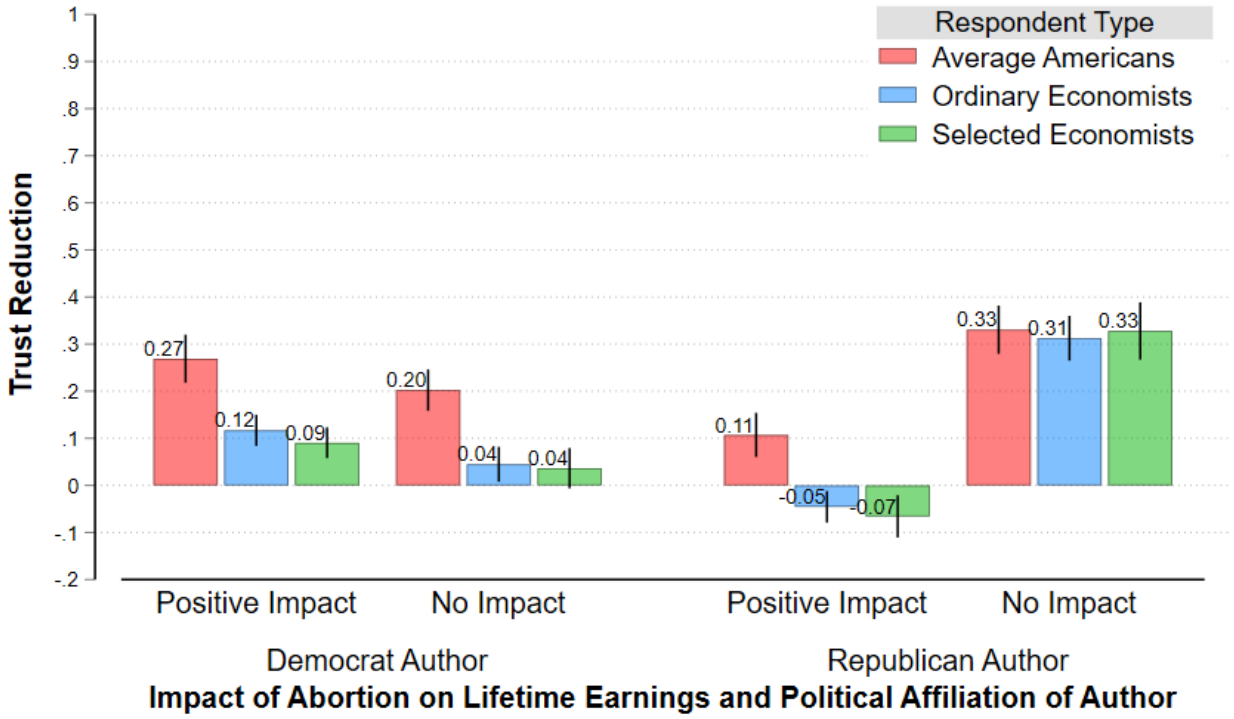
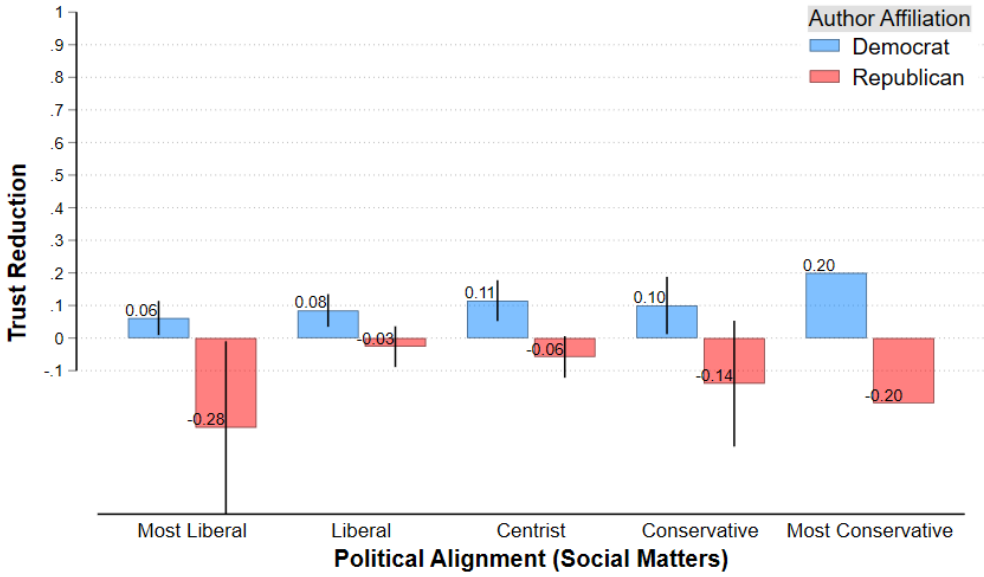


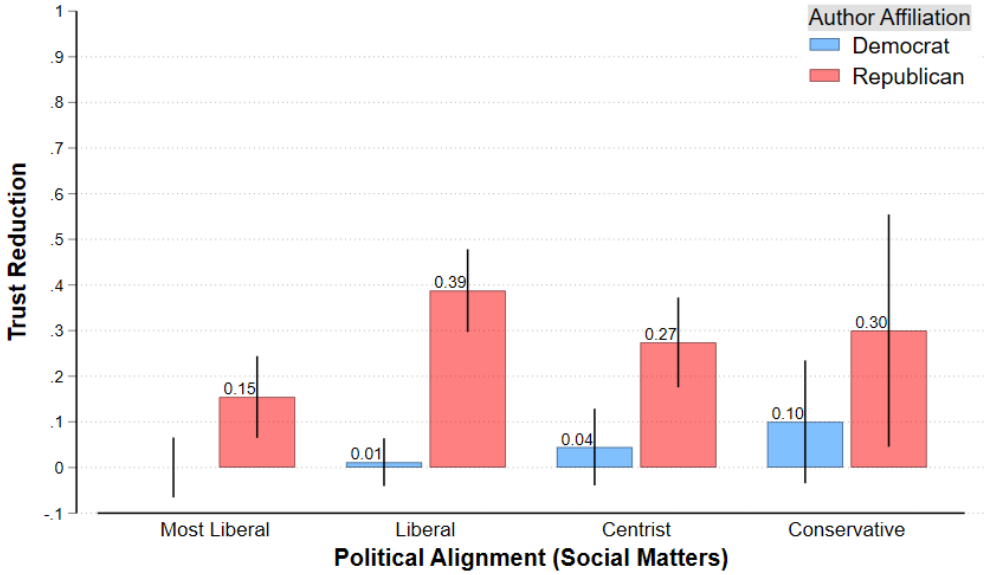
Figure 13: **Political Alignment and Trust Reduction: Abortion Policy, Selected Economists**

Figure 13 illustrates the CoI trust reduction in the abortion vignette (Figure 12) across selected economist respondents' political alignment on social matters. Panel A shows the results for the positive impact on the lifetime earnings of women, while Panel B shows the results for no impact of abortion on the lifetime earnings of women.

A. Positive Impact on Lifetime Earnings of Women

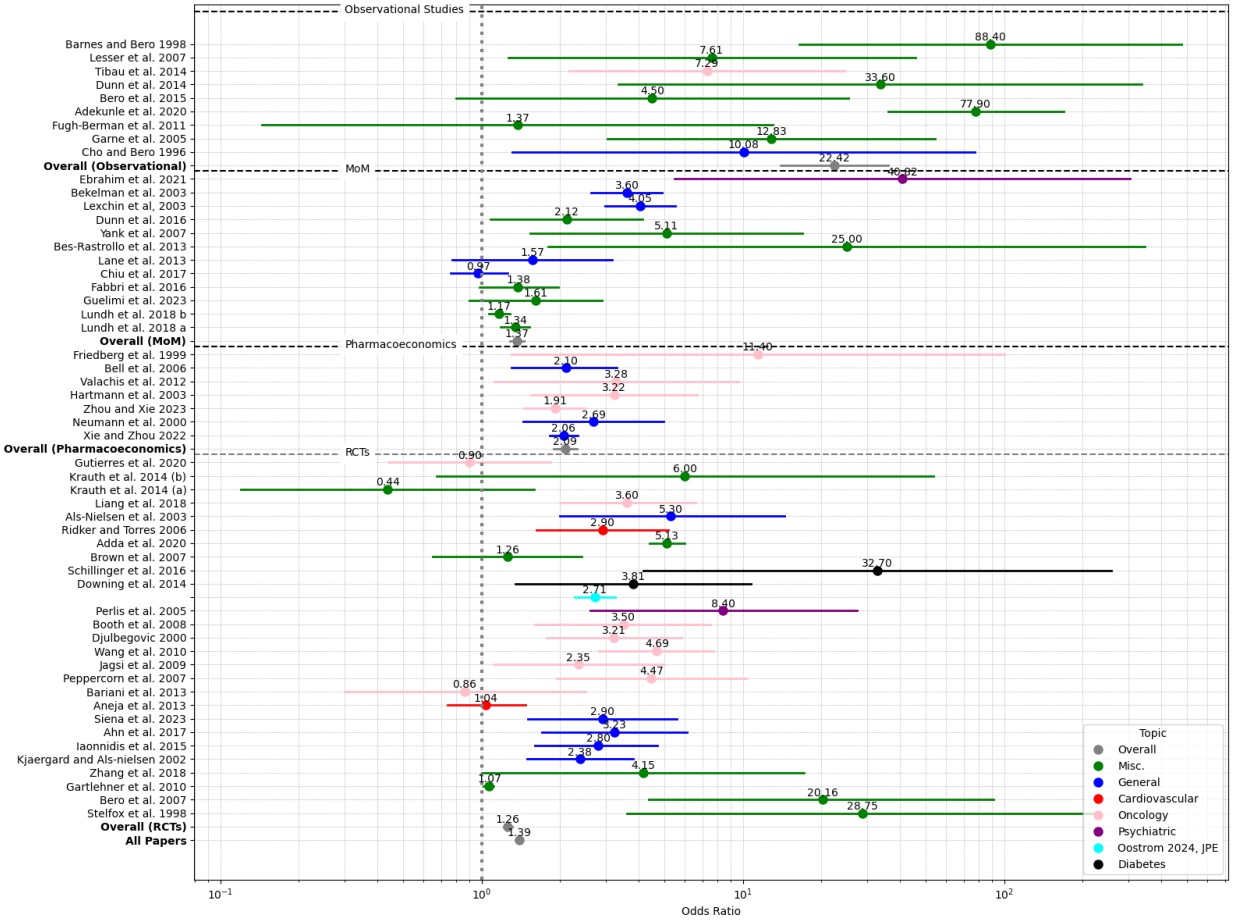



B. No Impact on Lifetime Earnings of Women



# Figure 14: **CoI and Bias in Results: Medicine**

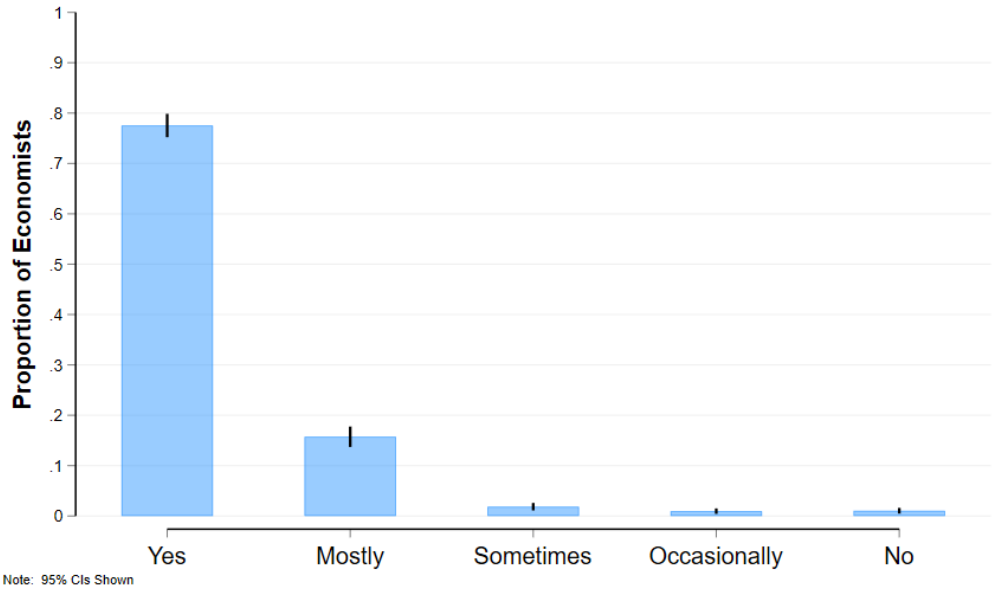
Figure 14 plots the odds ratios and confidence intervals of various studies from Table 14 (online appendix) that analyze CoI and how they bias results in medicine. We grouped niche, single-topic studies into the Miscellaneous category. The topics of the remaining studies are directly mapped from Table 14 (online appendix).



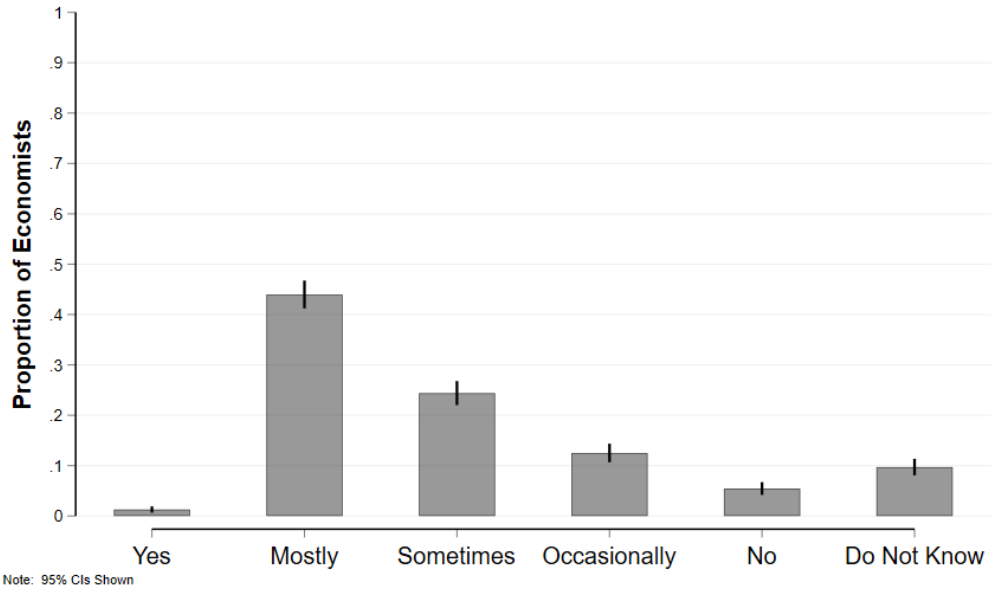
## Figure 15: **Economists' Disclosure**

Figure 15 illustrates economists' responses to our question about CoI disclosure. For Panel, A Self Disclosure, we ask them - *Do you think that your past disclosure of potential conflicts of interest has been adequate?* and elicit their responses on a 5-point scale (Yes, Mostly, Sometimes, Occasionally and No). For Panel B, Others' Disclosure, we ask them - *Do you think that, at present, academic economists generally disclose potential conflicts of interest adequately?* and elicit their responses on a similar scale, with an additional 6th point of "Do Not Know".

A. Self Disclosure




B. Others' Disclosure



# Tables

## Table 1: **Vignettes at a Glance**

Table 1 summarizes our various survey vignettes, divided into Science and Economics topics. We divide the economics vignettes into two categories, namely, Non-ideological and Ideological. The first column *Vignette* highlights the shorthand notation we use for the vignette. The second column *Findings of a Paper/Test* shows the first question we ask the respondent within each vignette to gauge their prior trust. Corresponding to each value in the second column, we report the associated disclosure treatment in the *Disclosure* column. The randomizations within each treatment have been italicized. The respondents only see one randomization per Finding/Disclosure.

| Vignette | Findings of a Paper/Test | Disclosure |
|---|---|---|
| *A. Science* | | |
| Industrial Disaster | A professional testing company tests your house for toxic substances post an industrial accident and finds your house safe | The testing company was hired by the company responsible for the industrial accident |
| Alzheimer's Drug | A medical publication that finds that a drug is effective in slowing the progression of Alzheimer's disease among patients | *Author was a consultant for the company that produces the drug / Author's research was funded by the company that produces the drug* |
| *B. Non-Ideological Economics* | | |
| Intensity of Financial Support: Trading Strategy | A trading strategy that utilizes investor's reactions to 'news' related to companies with a similar name can be used to make money | *$10k / 100k / 1M* in the form of a *research grant / consulting fee* |
| Expert Witness: Supermarket Merger | Larger chains are associated with decreased consumer welfare due to higher prices | Expert witness for DoJ Antitrust Division; *$400K Comp. / No mention of Comp.* |
| | Larger chains are associated with increased consumer welfare due to lower prices | Expert witness for the supermarket chain; *$400K Comp. / No mention of Comp.* |
| Retrospective/Prospective Conflicts: Underpaid CEOs | CEOs are underpaid relative to the value of the services they render to the company | Author was *seeking a position on the board of a large US public company / immediately nominated to the board / nominated to the board 2 years post publishing the paper* |
| Data Ownership: Ridesharing | Costs of ridesharing services outweigh their benefits | *Publicly available administrative data provided by the city government / Proprietary data provided by the ridesharing company* |
| | Benefits of ridesharing services outweigh their costs | |
| Terms of Access: Credit Card | Current business practices of Visa are efficient and the data used in the paper was given by Visa | Visa *had / waived off* the right to review the paper and the right to control which findings to make public |
| *C. Ideological Economics* | | |
| Tax Policy | Higher-income tax is beneficial for overall economic growth. | Author wrote a pro-tax book in the past |
| | Higher-income tax is not at all beneficial for overall economic growth. | |
| Abortion | Increased access to abortion does not affect lifetime earnings | Author was a *Democrat/Republican* |
| | Increased access to abortion positively affects lifetime earnings | |

## Table 2: **Summary of Response Rates**

Table 2 reports the summary statistics of the response rates. Total Reachout refers to the emails we sent that did not bounce back. Response Rate reports the proportion of respondents that completed at least one vignette. Passed Attention Test reports the proportion of respondents that passed the attention test vignette. Response Rate x Passed Attention Test reports the proportion of respondents that answered at least one vignette and passed the attention test. Final Responses denotes the number of respondents that answered at least one vignette and passed the attention test. These are the responses that we use for our analysis. Additionally, we also report Final Responses x Completed Vignettes as the number of respondents who passed the attention test and finished all the vignettes.

| | Economists | | Average Americans |
|---|---|---|---|
| | Ordinary | Selected | YouGov |
| Total Reachout | 12336 | 3051 | 1500 |
| Response Rate | 0.08 | 0.18 | 1.00 |
| Passed Attention Test | 0.76 | 0.84 | 0.85 |
| Response Rate x Passed Attention Test | 0.06 | 0.15 | 0.85 |
| Final Responses | 777 | 469 | 1280 |
| Final Responses x Completed Vignettes | 766 | 459 | 1280 |

# Table 3: **Summary Statistics: Respondents**

Table 3 shows summary statistics of the characteristics of respondents. Political alignment for social and economic matters is rated from 1 to 5, where 1 is Most Liberal (Abortion till the 9th Month/High Taxes and Government Intervention), 2 is Liberal, 3 is Centrist, 4 is Conservative, and 5 is Most Conservative (No abortion whatsoever/Low taxes and government intervention). Education level for the general public is defined from 1 to 6: 1 is No High School, 2 is High School Graduate, 3 is Some College, 4 is 2-year College, 5 is 4-year College, and 6 is Post Graduate. For the Average American sample, YouGov provided age and gender data. Respondents were also asked who they voted for in the 2020 Presidential elections. Paid Private/Public Consulting is coded as 1 if the respondent engaged in paid consulting activities. Used Proprietary Private/Public Data is coded as 1 if the respondent used proprietary or restricted-access data. Formal political affiliation is a dummy variable equal to 1 if the respondent self-reported a formal political appointment or affiliation. For economists, we report specializations that best align with our vignettes: Labor and Demographic Economics (Abortion), Financial Economics (Trading Strategy), Public Economics (Taxes), and Industrial Organization (Supermarket, Rideshare, and Credit Card). Economists were also asked about their PhD year in four-year bins starting in 1970 and ending in 2020. PhD years before 1970 and after 2020 are marked as 1970 and 2020, respectively.

| *Average Americans* | Mean | Median | SD | Min | Max | N |
|---|---|---|---|---|---|---|
| Political Alignment (Social Matters) (1 to 5) | 2.95 | 3.00 | 1.14 | 1.0 | 5.0 | 1280 |
| Political Alignment (Economic Matters) (1 to 5) | 3.17 | 3.00 | 1.23 | 1.0 | 5.0 | 1280 |
| Trust in Large Corporations (1 to 5) | 2.07 | 2.00 | 0.93 | 1.0 | 5.0 | 1280 |
| Education Level (1 to 6) | 3.86 | 4.00 | 1.48 | 1.0 | 6.0 | 1280 |
| Age | 52.21 | 54.00 | 16.52 | 18.0 | 92.0 | 1280 |
| Gender (0-M, 1-F) | 0.55 | 1.00 | 0.50 | 0.0 | 1.0 | 1280 |
| Voted for Biden | 0.41 | 0.00 | 0.49 | 0.0 | 1.0 | 1280 |
| Voted for Trump | 0.35 | 0.00 | 0.48 | 0.0 | 1.0 | 1280 |
| Not a US Citizen | 0.00 | 0.00 | 0.00 | 0.0 | 0.0 | 1280 |
| Did Note Vote | 0.24 | 0.00 | 0.43 | 0.0 | 1.0 | 1280 |
| Paid Private Consultant | 0.07 | 0.00 | 0.26 | 0.0 | 1.0 | 1280 |
| Paid Public Consultant | 0.05 | 0.00 | 0.22 | 0.0 | 1.0 | 1280 |
| Used Proprietary Private Data | 0.07 | 0.00 | 0.25 | 0.0 | 1.0 | 1280 |
| Used Proprietary Public Data | 0.05 | 0.00 | 0.22 | 0.0 | 1.0 | 1280 |
| Formal Political Affiliation | 0.06 | 0.00 | 0.24 | 0.0 | 1.0 | 1280 |
| *Ordinary Economists* | | | | | | |
| Political Alignment (Social Matters) (1 to 5) | 2.34 | 2.00 | 0.83 | 1.0 | 5.0 | 758 |
| Political Alignment (Economic Matters) (1 to 5) | 2.81 | 3.00 | 1.10 | 1.0 | 5.0 | 758 |
| Trust in Large Corporations (1 to 5) | 2.49 | 3.00 | 0.84 | 1.0 | 4.0 | 758 |
| PhD Year | 1995 | 1994 | 12.96 | 1970 | 2020 | 749 |
| Voted for Biden | 0.50 | 1.00 | 0.50 | 0.0 | 1.0 | 777 |
| Voted for Trump | 0.06 | 0.00 | 0.23 | 0.0 | 1.0 | 777 |
| Not a US Citizen | 0.31 | 0.00 | 0.46 | 0.0 | 1.0 | 777 |
| Did Note Vote | 0.11 | 0.00 | 0.31 | 0.0 | 1.0 | 777 |
| Paid Private Consultant | 0.43 | 0.00 | 0.50 | 0.0 | 1.0 | 777 |
| Paid Public Consultant | 0.45 | 0.00 | 0.50 | 0.0 | 1.0 | 777 |
| Used Proprietary Private Data | 0.28 | 0.00 | 0.45 | 0.0 | 1.0 | 777 |
| Used Proprietary Public Data | 0.33 | 0.00 | 0.47 | 0.0 | 1.0 | 777 |
| Formal Political Affiliation | 0.09 | 0.00 | 0.29 | 0.0 | 1.0 | 777 |
| Labor and Demographic Economics | 0.10 | 0.00 | 0.30 | 0.0 | 1.0 | 777 |
| Financial Economics | 0.05 | 0.00 | 0.22 | 0.0 | 1.0 | 777 |
| Public Economics | 0.07 | 0.00 | 0.26 | 0.0 | 1.0 | 777 |
| Industrial Organization | 0.05 | 0.00 | 0.21 | 0.0 | 1.0 | 777 |
| *Selected Economists* | | | | | | |
| Political Alignment (Social Matters) (1 to 5) | 2.31 | 2.00 | 0.73 | 1.0 | 5.0 | 455 |
| Political Alignment (Economic Matters) (1 to 5) | 2.78 | 3.00 | 0.95 | 1.0 | 5.0 | 455 |
| Trust in Large Corporations (1 to 5) | 2.63 | 3.00 | 0.82 | 1.0 | 4.0 | 455 |
| PhD Year | 2000 | 2002 | 12.08 | 1970 | 2020 | 454 |
| Voted for Biden | 0.47 | 0.00 | 0.50 | 0.0 | 1.0 | 469 |
| Voted for Trump | 0.02 | 0.00 | 0.13 | 0.0 | 1.0 | 469 |
| Not a US Citizen | 0.42 | 0.00 | 0.49 | 0.0 | 1.0 | 469 |
| Did Note Vote | 0.06 | 0.00 | 0.24 | 0.0 | 1.0 | 469 |
| Paid Private Consultant | 0.39 | 0.00 | 0.49 | 0.0 | 1.0 | 469 |
| Paid Public Consultant | 0.49 | 0.00 | 0.50 | 0.0 | 1.0 | 469 |
| Used Proprietary Private Data | 0.41 | 0.00 | 0.49 | 0.0 | 1.0 | 469 |
| Used Proprietary Public Data | 0.45 | 0.00 | 0.50 | 0.0 | 1.0 | 469 |
| Formal Political Affiliation | 0.11 | 0.00 | 0.31 | 0.0 | 1.0 | 469 |
| Labor and Demographic Economics | 0.10 | 0.00 | 0.29 | 0.0 | 1.0 | 469 |
| Financial Economics | 0.19 | 0.00 | 0.39 | 0.0 | 1.0 | 469 |
| Public Economics | 0.09 | 0.00 | 0.29 | 0.0 | 1.0 | 469 |
| Industrial Organization | 0.07 | 0.00 | 0.26 | 0.0 | 1.0 | 469 |

Table 4: **Summary Statistics: Respondent Priors**

Table 4 illustrates the prior trust of respondents across all our vignettes. Within each vignette, respondents were first asked about their trust levels in the findings of a paper/test on a 0-4 Likert scale, with 0 being "not at all" and 4 being "completely".

| *Full Sample* | Mean | Median | SD | Min | Max | N |
|---|---|---|---|---|---|---|
| Industrial Disaster | 2.21 | 2.00 | 0.98 | 0.0 | 4.0 | 2515 |
| Alzheimer's Drug | 2.12 | 2.00 | 0.84 | 0.0 | 4.0 | 2515 |
| Trading Strategy | 1.64 | 2.00 | 1.00 | 0.0 | 4.0 | 2515 |
| Supermarket Merger : Welfare Increasing | 2.04 | 2.00 | 0.96 | 0.0 | 4.0 | 1227 |
| Supermarket Merger : Welfare Reducing | 1.93 | 2.00 | 1.03 | 0.0 | 4.0 | 1287 |
| Underpaid CEOs | 0.98 | 1.00 | 1.04 | 0.0 | 4.0 | 2518 |
| Ridesharing : Benefits greater than Costs | 2.14 | 2.00 | 0.91 | 0.0 | 4.0 | 1252 |
| Ridesharing : Benefits less than Costs | 1.79 | 2.00 | 0.98 | 0.0 | 4.0 | 1264 |
| Credit Card | 1.60 | 2.00 | 0.95 | 0.0 | 4.0 | 2514 |
| Abortion : Positive Impact | 2.12 | 2.00 | 1.19 | 0.0 | 4.0 | 1241 |
| Abortion : No Impact | 1.60 | 2.00 | 1.10 | 0.0 | 4.0 | 1277 |
| Tax Policy : Pro-Tax Finding | 1.48 | 2.00 | 1.14 | 0.0 | 4.0 | 1271 |
| Tax Policy : Anti-Tax Finding | 1.92 | 2.00 | 1.14 | 0.0 | 4.0 | 1247 |

Table 5: **Likert Scale to CoI Trust Reduction Mapping**

Table 5 shows our mapping from the Likert Scale to CoI Trust Reduction. Respondents see only the values in the column *Impact on Prior Trust*, which are then mapped to a *CoI Trust Reduction*

| *Impact on Prior Trust* | *Score* | *CoI Trust Reduction* |
|---|---|---|
| It completely makes me trust the results | 3 | 100% increase |
| It seriously increases my trust in the results | 2 | 50% increase |
| It somewhat increases my trust in the results | 1 | 20% increase |
| It does not impact my trust in the results | 0 | No change |
| It somewhat undermines my trust in the results | -1 | 20% decrease |
| It seriously undermines my trust in the results | -2 | 50% decrease |
| It completely makes me distrust the results | -3 | 100% decrease |

# Table 6: **CoI Trust Reduction and Prior Trust: Vignettes Results**

Table 6 presents regression results where we regress CoI Trust Reduction for each vignette individually on Prior Trust and its interaction with the two economist categories. Prior trust refers to the pre-disclosure trust in the findings of the paper. We take only the randomization that constitute a conflict of interest from each of the vignettes. Therefore, all randomization are kept from Industrial Disaster, Alzheimer, Trading, Supermarket and CEOs vignettes. Only private data, benefits greater than costs is kept from the ridesharing vignette, only right to review enforced randomization is kept from the credit card vignette, only democrat author, democrat finding, republican author, republican finding randomization are kept in the abortion vignette. Only pro-tax finding, pro-tax book randomization kept in the Tax vignette. These regressions were run on the full sample and include respondent type fixed effects.

| | Trust Reduction | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | (1) Industrial Disaster | (2) Alzheimer | (3) Trading | (4) Supermarket | (5) CEOs | (6) Rideshareing | (7) Credit Card | (8) Abortion | (9) Tax |
| Prior Trust | -0.141*** | -0.131*** | -0.175*** | -0.130*** | -0.182*** | -0.0947*** | -0.196*** | -0.169*** | -0.177*** |
| | (0.0102) | (0.0138) | (0.0128) | (0.0119) | (0.0124) | (0.0227) | (0.0170) | (0.0142) | (0.0140) |
| Prior Trust X Ordinary Economists | 0.0411** | 0.0357* | 0.0588*** | 0.0494*** | 0.0549*** | 0.0143 | 0.0731*** | 0.0287 | 0.109*** |
| | (0.0184) | (0.0196) | (0.0183) | (0.0177) | (0.0173) | (0.0393) | (0.0251) | (0.0221) | (0.0195) |
| Prior Trust X Selected Economists | 0.0311 | 0.0976*** | 0.105*** | 0.0816*** | 0.0859*** | 0.0681* | 0.113*** | 0.0584** | 0.160*** |
| | (0.0232) | (0.0232) | (0.0228) | (0.0211) | (0.0202) | (0.0355) | (0.0314) | (0.0258) | (0.0201) |
| Constant | 0.779*** | 0.570*** | 0.590*** | 0.406*** | 0.498*** | 0.434*** | 0.750*** | 0.530*** | 0.277*** |
| | (0.0194) | (0.0205) | (0.0162) | (0.0180) | (0.0122) | (0.0391) | (0.0210) | (0.0230) | (0.0172) |
| Respondent Type Fixed Effects | X | X | X | X | X | X | X | X | X |
| Full Sample | X | X | X | X | X | X | X | X | X |
| N | 2515 | 2515 | 2515 | 2514 | 2516 | 618 | 1251 | 1253 | 1271 |

# Table 7: **CoI Trust Reduction and Prior Trust: Individual Results**

Table 7 presents regression results where we regress CoI Trust Reduction on Prior Trust for all respondents and then decompose the results by respondent types. We add individual respondent-level fixed effects in the even columns. Prior trust refers to the pre-disclosure trust in the findings of the paper (designated as $P_o$ in our model). We take only the randomization that constitute a conflict of interest from each of the vignettes. Therefore, all randomization are kept from Industrial Disaster, Alzheimer, Trading, Supermarket and CEOs vignettes. Only private data, benefits greater than costs is kept from the ridesharing vignette, only right to review enforced randomization is kept from the credit card vignette, only democrat author, democrat finding, republican author, republican finding randomization are kept in the abortion vignette. Only pro-tax finding, pro-tax book randomization kept in the Tax vignette.

| | Trust Reduction | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| Prior Trust | -0.530*** | -0.459*** | -0.635*** | -0.514*** | -0.436*** | -0.408*** | -0.273*** | -0.313*** |
| | (0.0126) | (0.0139) | (0.0189) | (0.0197) | (0.0209) | (0.0241) | (0.0283) | (0.0336) |
| Constant | 0.520*** | 0.492*** | 0.560*** | 0.518*** | 0.485*** | 0.473*** | 0.379*** | 0.398*** |
| | (0.00610) | (0.00632) | (0.00849) | (0.00825) | (0.0107) | (0.0118) | (0.0149) | (0.0169) |
| Respondent Fixed Effects | | X | | X | | X | | X |
| Sample | All Respondents | All Respondents | Average Americans | Average Americans | Ordinary Economists | Ordinary Economists | Selected Economists | Selected Economists |
| N | 11938 | 11933 | 6064 | 6064 | 3668 | 3665 | 2206 | 2204 |

Table 8: **Potential impact of CoIs on findings: Paper Level (2019-2023)**

Table 8 presents the results of our manual coding exercise on CoIs. We analyzed 551 papers from the *American Economic Review* (135 theoretical, 416 empirical) and 164 papers from the *RAND Journal of Economics* (103 theoretical, 61 empirical), identifying disclosures related to research grants, consulting fees, employers, data sources, and editorial rights to review. For political appointments, we constructed a separate dataset covering all federal political appointments for academic economists and merged it with our AER and RAND data. Similarly, for political affiliations, we created a separate dataset of academics' donations to candidates running for federal office. For the latter, we use the publicly available donations data from the FEC (we only consider donors with cumulative donations above $10,000 between 2000-2024). We classified potential CoI based on whether these disclosures could have influenced a paper's findings. The row *Overall Conflicted* presents the proportion of papers with at least one potential conflict.

| **Conflict** | **AER** | | | **RAND** | | |
|---|---|---|---|---|---|---|
| | **Theory** | **Empirical** | **All** | **Theory** | **Empirical** | **All** |
| Consulting Academic | 0.015 | 0.051 | 0.042 | 0.000 | 0.000 | 0.000 |
| Research Grant | 0.008 | 0.056 | 0.044 | 0.010 | 0.067 | 0.031 |
| Primary Employment | 0.135 | 0.166 | 0.159 | 0.078 | 0.213 | 0.128 |
| Explicit Editorial Right to Review | 0.007 | 0.058 | 0.045 | 0.000 | 0.000 | 0.000 |
| Political Appointment or Affiliation | 0.030 | 0.046 | 0.042 | 0.029 | 0.033 | 0.030 |
| Discretionary Gated Data Sources | 0.000 | 0.284 | 0.215 | 0.000 | 0.459 | 0.171 |
| Overall | 0.187 | 0.454 | 0.389 | 0.117 | 0.574 | 0.287 |
| **N** | 134 | 416 | 551 | 103 | 61 | 164 |

## Table 9: **CoI Discounts: AER, Empirical Papers' $\lambda$s**

Table 9 maps the Trust Reduction into the CoI Discount. Column *Trust Reduction* reports the average reduction in trust of subjects in our survey when they learned about the existence of a CoI. The *CoI Discount* columns report the CoI discount. These values are computed from Equation (9), using two different estimates of the frequency of conflicted papers $\lambda$. $\lambda_{specific}$ denotes the conflict-specific $\lambda$ value taken from Table 8 (only AER papers). $\lambda_{overall}$ is the overall $\lambda = 0.389$ taken from Table 8. *Average* denotes the average of all vignette randomizations that disclose a conflict ex-ante. All randomizations except Rideshare Public Data satisfy this constraint. The rest of the rows decompose the results by different randomization of vignettes. For readability, the table below excludes some randomizations, though they were included in the calculation of the *Average*. For Trading : Research Grant vignette we use the *research grant* specific $\lambda$, for Trading : Consulting fee, we use the *consulting fee* specific $\lambda$, for Supermarket vignettes we use the *consulting fee* specific $\lambda$, for CEO vignettes we use the *employer sources*, for rideshare we use the *discretionary gated data sources* specific $\lambda$ and for Credit Card vignette we use the *editorial right to review* specific $\lambda$. We use the political appointment/affiliation specific $\lambda$ for Tax and Abortion vignettes.

| *Vignette* | *Trust Reduction* | *Specific* | | *Overall* | |
|---|---|---|---|---|---|
| | | $\lambda$ | *CoI Discount* | $\lambda$ | *CoI Discount* |
| Average | 0.28 | | | 0.39 | 0.39 |
| Trading : Research Grant | 0.34 | 0.04 | 0.35 | 0.39 | 0.46 |
| Trading : Consulting Fee | 0.4 | 0.04 | 0.41 | 0.39 | 0.52 |
| Trading : 1M Consulting Compensation | 0.44 | 0.04 | 0.45 | 0.39 | 0.56 |
| Supermarket : Public Expert Witness, No mention of Compensation | 0.08 | 0.04 | 0.09 | 0.39 | 0.13 |
| Supermarket : Public Expert Witness, 400K Compensation | 0.21 | 0.04 | 0.22 | 0.39 | 0.3 |
| Supermarket : Private Expert Witness, No mention of Compensation | 0.2 | 0.04 | 0.21 | 0.39 | 0.29 |
| Supermarket : Private Expert Witness, 400K Compensation | 0.36 | 0.04 | 0.37 | 0.39 | 0.48 |
| CEO : Nominated 2 Years Post | 0.3 | 0.16 | 0.34 | 0.39 | 0.41 |
| CEO : Immediately Nominated | 0.36 | 0.16 | 0.41 | 0.39 | 0.48 |
| CEO : Seeking Position | 0.42 | 0.16 | 0.47 | 0.39 | 0.55 |
| Rideshare : Private Data | 0.2 | 0.21 | 0.24 | 0.39 | 0.29 |
| Credit Card : Right to Review Enforced | 0.52 | 0.23 | 0.58 | 0.39 | 0.64 |
| Pro Tax Book : Pro Tax Finding | 0.12 | 0.04 | 0.13 | 0.39 | 0.19 |
| Abortion : Democrat Author | 0.16 | 0.04 | 0.16 | 0.39 | 0.23 |
| Abortion : Republican Author | 0.18 | 0.04 | 0.18 | 0.39 | 0.26 |

Table 10: **Average Citation Count by Conflict Status and Journal Category**

Table 10 shows the average citation counts of conflicted vs. non-conflicted papers in the AER and RAND (2019-2023). In the final row, we report the p-value of the difference between the average citation counts of the conflicted v/s non-conflicted papers within each journal. We collect the citations from official journal sources for each article (AER and RAND).

| | **AER** | | **RAND** | |
|---|---|---|---|---|
| | **Theory** | **Empirical** | **Theory** | **Empirical** |
| Overall Conflicted | 95.200 | 208.571 | 56.556 | 43.200 |
| Overall Non-Conflicted | 143.122 | 185.628 | 26.617 | 41.269 |
| p-value | 0.428 | 0.402 | 0.192 | 0.889 |
| **N** | 134 | 416 | 103 | 61 |

Table 11: **GPT-4 Omni v/s Human Respondents' Trust Reduction**

Table 11 compares the mean trust reduction in our original survey (first column) with the mean trust reduction when the same survey is run with GPT-4 Omni (second column). The last two columns report their absolute difference and statistical difference. The trust reduction for each vignette can be found in the first four columns of Table 7 in the online appendix.

| *Respondent Type* | $r$ | $r_{GPT}$ | *Diff.* | *P-value* |
|---|---|---|---|---|
| Average Americans | 0.286 | 0.303 | 0.017 | 0.398 |
| Ordinary Economists | 0.214 | 0.303 | 0.089 | 0.000 |
| Selected Economists | 0.184 | 0.303 | 0.119 | 0.000 |

Table 12: **Within-Subject Survey v/s Between-Subject Survey Trust Reduction**

Table 12 compares the mean trust reduction in our original *within-subject* survey (first column) with the mean trust reduction when the *between-subject* survey is run using GPT-4 Omni (second column). The last two columns report their absolute difference and statistical difference. The trust reduction for each vignette can be found in Table 7 in the online appendix.

| *Respondent Type* | $r_{within}$ | $r_{between(GPT)}$ | *Diff.* | *P-value* |
|---|---|---|---|---|
| GPT-4 Omni | 0.303 | 0.319 | 0.016 | 0.832 |
| Average Americans | 0.286 | 0.319 | 0.033 | 0.605 |
| Ordinary Economists | 0.214 | 0.319 | 0.105 | 0.030 |
| Selected Economists | 0.184 | 0.319 | 0.135 | 0.008 |

**Appendix**

# A Survey

Figure A1 is a QR code which acts as hyperlink to a Google drive that contains the two versions of our surveys, one for economists and one for the general public. The only difference between the two surveys lies in the non-vignette questions, for example (i) PhD Year of Economist, (ii) Area of Specialization, (iii) Tenure status, (iv) External pressure to change results, (v) Adequate self-disclosure, (vi) Career Aspirations.



Figure A1: QR Code : Access to Survey Instrument PDFs