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

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

March 21, 2025

We conducted a randomized, controlled survey of economists and a representative sample of Americans to infer how conflicts of interest (CoI)–that is, the presence of authors with a financial, professional, or ideological stakes in the outcome of the results–impact the trustworthiness of the underlying academic findings. On average, a CoI decreases trust in the conclusions of an economics paper by 28%, though there is significant variation according to the type of conflict. We model this change in trust as reflecting a combination of the frequency of conflicted papers and the bias of papers when they are conflicted. To isolate the second term, we introduce a key construct: the *CoI Discount*, which measures the change in the value of a conflicted paper relative to a non-conflicted one. We show that, on average, conflicted papers are worth only 61% of non-conflicted ones. We validate our survey-based estimates by comparing them to actual biases observed in conflicted papers in economics, conducting a large meta-review of the medical literature on the topic, and by using LLM simulations.

*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 Jonathan Lee and Henry Cummins for their excellent research assistance. We also thank Amit Seru and the participants in seminars in Brussels, Chicago, Columbia, ETH Zurich, George Mason, Harvard, Imperial, and the University of Southern California 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 Requirements, the authors disclose the following potential conflicts of interest over 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:* is 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, 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 primary role of academic researchers is to generate new ideas that compete in the intellectual marketplace. These ideas vary not only in their novelty but also in their trustworthiness—the degree to which a paper influences individuals' prior beliefs toward its conclusions. Trustworthiness, in turn, depends on multiple factors, one of the most significant being the presence of an author with a conflict of interest (CoI).

CoIs 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: pharmaceutical firms might be more likely to retain the services of a consultant-scholar who has previously found results that redound to the company's benefit; ideological or others. Since empirical and theoretical analyses often involve subjective judgments, the greater a researcher's interest in a specific direction of the 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 findings to align them with the desired outcome (Dawson et al., 2002). However, despite widespread recognition that conflicts of interest can impact academic research–almost all major economic journals currently have a CoI disclosure policy–there is only scant evidence on what types of practices lead to a meaningful conflict and how different types of conflicts affect the integrity of research findings.

This paper examines how different types of CoIs shape perceptions of trustworthiness and value of research. To this purpose, we survey both academic economists and a representative sample of the U.S. population. Economists are further split into two samples, those at the top of the profession (denoted as selected economists) and the rest (denoted as ordinary economists). Respondents were initially presented with vignettes describing findings from both economics and science and asked to provide their initial trust on the described results. We then treated respondents with a randomized exposure to a disclosure statement that presented different types of CoI, asking respondents how much this added information impacted their initial beliefs. Each vignette focused on one 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, which arise when research outcomes may influence an author's professional trajectory. We then consider data conflicts, where study data are controlled by a public or private organization with a vested interest (or not) in the direction of the results. Additionally, we analyze academic conflicts, which occur when an author's reputation is linked to a particular set of findings, and political conflicts, which emerge when an author's political ideology aligns with specific results.[1]

Thanks to the answers we receive, we can calculate how the presence of a CoI impacts trust. On average, CoIs reduce trust in the results of economics articles by 30%, so we call this variable *CoI Trust Reduction*. This CoI Trust 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%. Career-related CoIs reduce trust by 36%, while data-access conflicts result in reductions between 20% and 52%, depending on the level of control over the data. Yet, the use of publicly available data in an article, for example, *increases* trust in the results by 4%. Academic CoIs, by contrast, have minimal effects (and average Americans do not distinguish between a conflicted paper and a non-conflicted one). Economists' trust *increases* by 4% when the finding is against the academic interest of the author and decreases by 5% when it is aligned with the author's interest. Ideological (political) CoIs reduce trust by 17% on average, but this average masks significant variation between how much democrat readers trust republican authors and vice-versa–signaling an increasing 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. While financial incentives are important, the specific amount of compensation matters less than expected. For example, trust falls by 44% when an author receives \$1M, compared to 37% for an author receiving \$10,000 as a consulting fee. This result is hard to

[1]While we do not directly 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 the evolution of science–albeit distinct in nature from the class of CoI-induced biases analyzed in this paper. We plan to explore this issue in future work.

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 the bias arising from motivated reasoning. Career-related conflicts also significantly impact trust, particularly when professional advancement opportunities are involved. Respondents reduce their trust by 42% after learning that an author is seeking employment with an entity that benefits from their research findings. However, if the appointment occurs two years after publication of the article, the trust reduction is lower, at 30%. Additionally, the nature of the relationship also matters: respondents are more skeptical of a CoI involving a private company than one involving the government. For example, trust in the results of a paper written by a former expert witness for a private plaintiff decreases by 28%, whereas the reduction is only 15% when the author served as an expert witness for the Department of Justice.

We find that the control and availability of data play a critical role in shaping trust. We observe significant differences in trust based on whether the data used in a study are proprietary or public. Regardless of the result of a study, proprietary data leads to a 20% reduction in trust, while 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, 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 reductions in trust and respondents' expertise. Experts tend to trust the result more to begin with, and 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 well-trained PhD 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 CoIs erode trust in research findings by influencing respondents' perceptions of credibility. While understanding this trust reduction is important, it represents only part of the broader issue. A reduction in trust implies a deeper concern: that CoIs not only affect how research is perceived but also diminish its overall value in the intellectual marketplace. 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 framework 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 pre-existing prior. 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 our respondents did not anticipate any potential presence of a conflict of interest when providing their initial answers on how much they trusted the headline results of a given article.

To measure the expected frequency of conflicts of interest, we estimate the frequency of such conflicts by analyzing the disclosures from all papers 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 conflicted. Theory papers are less conflicted – 19% of AER theory papers (out of a total of 135 papers) and 12% of RAND theory papers (out of a total of 103 papers) present some form of conflict. While, for empirical papers, these percentages are respectively 45% (out of a total of 416 papers) and 57% (out of a total of 61 papers). The most common CoIs are the reliance on discretionary, gated data and political appointments or affiliations.[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

[2]We measure political appointments by creating a dataset of economists with formal federal appointments. Similarly, we use the FEC data to map AER/Rand authors in the 2019-2023 quinquennium to their donations (cumulatively greater than $10,000) to federal election candidates between 2000-2024.

a right to review (CoI Discount of 58%), followed by articles written by highly paid consultants (CoI Discount of 44%), scholars seeking positions in boards of public companies (CoI Discount of 47%) and highly paid private expert witnesses (CoI Discount of 37%). The CoI Discount is less-pronounced for scholars with a history of publications in a given direction (13%) or with an open political leaning (17%). These results highlight the extent to which CoIs can affect not only the trust in research but also its overall value in the marketplace of ideas–though, again, the specific types of CoIs present in a given article matter.

One potential criticism of our findings is that they simply reflect a reduction in trust due to *perceived* conflicts of interest rather than *actual* ones. Even if this were a mere perception, it would be important. If economists and the public perceive certain results as both being less valuable and are less convinced of them, these results are inferior quality products in the marketplace of ideas, becoming less likely to influence public policy and human or corporate behavior. This happens regardless of the correctness of the perception of bias. Nevertheless, we run 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 the effects of CoI in economics and medicine. In economics, there are only two papers (Asatryan et al., 2020, and Fabo et al., 2021) that study the effect of CoI with observational data. They both find that conflicts tend to inflate the results in the direction that favors the interested party. In medicine, after aggregating all meta-studies published in reputable medical journals from 1998 to 2023, we find that industry-sponsored studies are, on average, 1.4 times more likely to report statistically significant findings that benefit the sponsor compared to non-sponsored studies. This average, however, conceals considerable heterogeneity across methodologies. Among observational studies, the odds of industry-sponsored studies reporting significant results are 22.4 times higher than those of non-sponsored studies. By contrast, randomized control trials (RCTs) show a much smaller difference, with the odds of statistically significant results being 1.3 times higher in industry-sponsored trials than in non-sponsored ones. All in all, improvements in methodology are important, but CoIs can still generate bias even in medical RCTs, our current methodological gold-standard.

Second, we calibrate our model by using the findings from Fabo et al., 2021. We find that the

optimal belief updating process mirrors the trust reductions observed in our survey given the CoI in Fabo et al., 2021 (employees of a public entity).

Third, we extend our analysis by simulating an additional 1,000 survey responses per randomization using GPT-4 Omni, a large-scale language model trained to replicate human decision-making patterns. The results generated by GPT-4 Omni are similar to those obtained from our three survey instruments, further validating the robustness of our findings. Notably, GPT-4 Omni exhibits a CoI trust reduction pattern more aligned with the responses of economists than with those of the general population, suggesting that it mimics the more skeptical approach of experts when evaluating research with potential conflicts of interest.

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.

This article contributes to many inter-disciplinary streams of scholarship. The medical field pioneered the study of conflicts of interest in academic research. Numerous studies and meta-analyses, such as Ioannidis, 2005, have demonstrated that studies with strong financial backing are more prone to bias due to explicit monetary CoI. This research has been critical in revealing how funding sources can systematically skew outcomes in favor of sponsors, thereby 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 Long and Lang, 1992, who highlighted the prevalence of insignificant results being unpublished (see Christensen and Miguel, 2018 for a comprehensive survey). Ioannidis et al., 2017 finds that nearly 80% of the reported effects in the empirical economics literature are exaggerated, typically by a factor of two and with one-third inflated by a factor of four or more. 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 (Saint-Paul, 2012; Jelveh et al., 2024). As mentioned above, only two studies in economics parallel the extensive work on CoI in the medical field. Asatryan et al., 2020 finds that government-funded studies tend to report larger fiscal multipliers, suggesting a bias towards favorable outcomes that support governmental policies. Fabo et al., 2021 demonstrates that central bank-affiliated papers report significantly larger effects of quantitative easing (QE) on both output and inflation, indicating potential institutional biases.

To the best of our knowledge, our work is the first study in economics to systematically quantify the *CoI Discount* toward academic papers generated by disclosing different conflicts of interest of the authors to readers. 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 economic research. Our data also allows us to understand how individuals' political and economic priors–or personal biases and beliefs–influence both the interpretation of economics research and the potential CoI associated therewith.

Our survey approach to identifying the effect of conflicts is novel in economics but not in medicine. Kesselheim et al., 2012 surveys 500 board-certified internists and find 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. They look only at one type of conflict, while we compare different varieties and distinguish between the reduction in trust and the reduction in the value of a paper. Østengaard et al., 2020 uses a survey not to uncover the effect of conflicts but to identify how industry funding influences drug trials.[3]

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 policy (McGarity and Wagner, 2008; McIntire and Kantor, 2024; Michaels, 2008). Policymakers

[3]Surveys have also been used to assess patients' willingness to participate in industry-sponsored trials. Snyder et al., 2009 finds that a major barrier to recruiting elderly participants in Alzheimer's studies is the lack of trust in research being conducted in this area, with industry-sponsored trials being the least trusted of all. By contrast, Hampson et al., 2006 finds that industry ties do not impact cancer patients' willingness to enroll in experimental drug trials.

and academics are increasingly concerned about the potential dissolution of the lines that divide independent academia from business (or other) advocacy (Kanter, 2024, Nevo, 2024,Valletti, 2024, Lianos, 2024). The public's trust in science and universities is also declining (Gallup Inc, 2023), and commentators have identified academics appearing to be "self-interested rather than disinterested", or insulated as some reasons behind this decline (Mills, 2023; Rosenbluth, 2024). At the same time, modern academic research is becoming more expensive and requiring access to data increasingly controlled by private parties (Edelson et al., 2023; Wagner, 2023). In many areas, high-quality research requires and directly benefits from interaction with industry participants, government agents, think-tanks and others. This project does not study the *benefits* to science that arise out of relationships that are considered CoIs. Therefore, our article should ***not*** be interpreted as recommending a ban on such relationships. We do not know the socially-optimum amount of industry-academia and other relationships.

That notwithstanding, our data fill a significant gap, and points to many problems in current academic practices. For example, our analysis demonstrates that the current disclosure system lacks credibility–only 2% of our economists respondents believe that academic economists appropriately disclose their conflicts of interest. This lack of trust creates a negative externality born by independent researchers. In addition, the relationships that led to the largest CoI Trust Reduction involve private data access–a growing practice that deserves much more academic attention. By not openly discussing how CoIs can impact science, academics risk exposing our work to criticism and undermine the overall trustworthiness of scientific findings among both experts and the public.

In particular, our results indicate that while stricter enforcement of disclosure rules could mitigate some of these negative effects, it would not address the underlying problem: academics often internalize the private benefits of conflicts— such as higher payments or access to exclusive data— while the broader social costs, including diminished overall trust and distorted research priorities, remain unaddressed. Our findings show that disclosure alone is insufficient to realign private incentives with public welfare. To correct this, conflicted papers must be discounted in literature reviews, legal proceedings, and academic promotions. Our analysis provides initial baselines for this *CoI Discount*, demonstrating that its magnitude varies based on the type and severity of the conflict, with financial and data-related conflicts resulting in the largest credibility

gaps. More work, however, is needed to better calculate discount factors and tie them to real-world situations.

Above all, we hope this article will help ignite an open conversation about the benefits and downsides of Conflicts of Interest in economics and academia more broadly. Without a structural change in how conflicted research is assessed, including better measurements of what benefits it brings to science and what are potential downsides, the misalignment between private incentives and public trust will continue. This will further erode the long-term credibility of academia and expertise more broadly.

The rest of the paper proceeds as follows: Section 2 describes the survey. Section 3 presents the main empirical findings. Section 4 introduces a simple framework to interpret the data in light of a 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. Section 8 concludes.

2 Survey Design and Implementation

To empirically assess how CoI affect trust in economic research findings, we designed a survey featuring a series of vignettes. Each vignette simulated real-world research scenarios in which potential conflicts of interest—such as financial, career, data-related, academic, or political conflicts—could arise. By presenting respondents with these scenarios, we aimed to capture how different types of conflicts affect their perceptions of research trustworthiness.

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 standards.

Participants were presented with nine vignettes in a randomized order. Each vignette began

with a statement describing the key finding of the article. 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, one 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 bold parts indicate the random variation within the vignette. Each participant only saw each vignette once.

Subsequently, respondents were randomly treated with information about a potential CoI related to the authors of the study. They were then asked how this CoI disclosure changed their trust in the paper's results using a symmetrical 7-point Likert scale ranging from "It completely makes me distrust the results" to "It completely makes me trust the results." For instance, the treatment in the abortion vignette was about political conflict 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 from the author.[4] Appendix A.1 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 between economics and physical sciences (a drug test and a chemical test; see Panel A). Of the seven economic vignettes, five of them addressed economic CoI, specifically related

[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.

to data access, funding, career opportunities and consulting opportunities for the government and private parties (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 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 (formerly Initiative on 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 about ideology in economics publication. 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 and other failed emails, we successfully 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 to a final dataset of 1,500 answers. This group is representative in relation to gender, age, race, education, family income, and political preferences.[5]

These three samples represent different segments. The Selected Economists group encompasses mostly 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 both 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.

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 amount of questions on conflicts present in the survey, there was no chance that the respondents would not eventually 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 introduced two potential costs: selection bias and priming. Both are possibilities, meaning that their presence (or lack thereof) must be empirically tested. We conducted as many tests as our data permit, and our results do not indicate the presence of strong selection or priming.

Starting with the risk of 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

[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. However, despite multiple attempts, we did not manage 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.

percentage in this type of survey. Even if our results were to be read as the average change in trust among one out of five of the top economists, it would still be interesting and relevant. Still, additional tests do not indicate the presence of selection bias. We collected data on the year of PhD graduation for our selected economists sample (in 5 years interval). 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 allay 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 the sample characteristics and these participants were not told 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-interest. The average reduction in trust in science among economists is identical to the average reduction in trust among average Americans. Thus, there is no evidence of selection in this sense.

Moving to potential priming connected to the fact that respondents are aware that the survey focuses on CoIs. It is not theoretically clear how priming would impact our results. Arguably, because we first ask respondents about their initial trust in a paper's conclusions and then rely on disclosure of a CoI as the relevant treatment, priming could actually lead us to underestimate the impact of CoIs 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 disclosure. That nonetheless, because our survey was fully randomized–meaning that respondents saw questions in random orders–we can build on that variation to empirically test for priming across questions. Figure 2 shows the average CoI Trust Reduction per vignette position for both prior trust levels and CoI Trust Reduction. As can be seen, the differences are not statistically significant, diminishing concerns that we elicited an experimenters' effect.

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.[6] We used Qualtrics to administer the survey to both Selected and Ordinary Economists. We employed 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, a few 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 economic 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.[7] For our main analysis, we only utilize 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.[8]

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 experi-

[6]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).

[7]To check if the respondents were paying attention and not haphazardly responding to questions, we incorporated an attention test question, randomly in between vignettes, 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.*

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

ences 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 delving into the specifics of how conflicts of interest 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 economic 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 those in the economics vignettes. This is true for all three subsamples, albeit selected economists are always more trusting, followed by the ordinary economists, and the average Americans last.

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 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 interpret. 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.

To validate our approach, we used GPT-4 Omni to generate a semantic mapping of the Likert scale descriptions to numerical discount values. We simulated this task 1,000 times (see online appendix, Section 1.3). GPT-4 Omni assigned a 75% change in trust to the "serious discount" and a 50% change to the "somewhat discount," both of which are higher than our thresholds of 50% and 20%, respectively. Moreover, GPT-4 Omni underestimates the positive impact on trust post-disclosure. Despite these differences, we opted for our more conservative approach, which biases our results downward and thereby strengthens our conclusions. Importantly, these different mappings do not affect the cross-sectional variation or statistical significance of our findings, even if they impact the absolute levels.

Under most treatments, the level of trust declines when the CoI is disclosed. As a result, we define CoI Trust Reduction the negative change in trust. In other words, a positive trust reduction represents a decrease in trust following the disclosure, while a negative trust reduction indicates an increase in trust. This convention simplifies 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 CoIs reduce trust in research across different fields and types of conflicts? Figure 3 (Panel B) presents the CoI trust reduction averaged across groups of vignettes and in different sub-samples.

All three groups of respondents exhibit a similar trust reduction in science vignettes (42-44%).

[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.

Overall, we see a 30% trust reduction in economics vignettes, but with different levels in the subsamples. The average is 34% among ordinary Americans, 29% among ordinary economists, and 25% among selected economists. This lower reduction in trust among the most selected economists can be due to their knowledge that conflicts are more pervasive in economics (and thus they are less surprised) or their perception that biases are less severe in economics. We will return to this issue soon.

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 they do not trust from the beginning.

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 a dummy 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. Interestingly, this effect is smaller for ordinary economists and especially so for selected economists.

This lower reduction in trust can be due to a higher expectation that CoIs 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 review the presence of a conflict contradicts the first hypothesis–if they expected CoIs 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 more benign treatment can be objective or behavioral. Our results will sort this out 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 results above show that, on average, awareness about authors' CoI significantly decreases the trustworthiness of economic articles. The seven economic vignettes, however, present different types of conflicts and are subjected to different treatments. We now provide better context to these discussions by studying each of the CoI separately, notably: (i) monetary incentives, (ii) career incentives, (iii) access to data, (iv) academic conflict, and (v) ideological conflict.

3.4.1 Monetary Incentives

Monetary incentives feature prominently in the American Economic Association (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 that deal with CoI related to such monetary incentives.

Figure 5 Panel A shows the results for the vignette describing a study on the profitability of a financial trading strategy based on investors' reactions to the news. There are two randomizations. The first regards whether the support comes in the form of a grant or a consulting fee. The second regards the amount: we tell the subjects that the author of the study received $10K, $100K, or $1M from a company employing such strategies.

The support through a grant engenders lower mistrust than the support through a consulting fee. If we pool across amounts and samples, the average CoI trust reduction for research grants is smaller (34%) than the one for consulting fees (40%). Yet, the most surprising result is the level of the trust reduction for a grant. Even among selected economists, this trust reduction averages 28%. Higher payments lead to larger reductions in trust, but the differences are smaller. The average reduction for a $10K consulting fee payment is 37%, while the discount for a $1M payment is 44%. Selected economists do not seem to respond much to the amount at stake when it is paid in the form of a consulting fee, while they react more significantly when it is in the form of a grant. A $10K grant leads to a 21% reduction in trust, while a $1M grant, a 32% one.

Panel B further decomposes the effect of expertise further by dividing economists by their area of expertise, and comparing Industrial Organization (IO) and financial economists (remember that this vignette is focused on the profitability of financial trading strategies). IO economists' trust drops by 36% if the author has received $1M in consulting fees, while financial economists' trust drops by only 15%. Do financial economists consider all financial conflicts as more benign, or is it unique to the conflicts where they are more likely to be involved? A comparison with another vignette helps us answer this question.

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 large 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 result is always in the direction of 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 the 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 the awareness about the amount paid and the ultimate employer matter.

First, the results obtained by a former DoJ expert witness are more believed than those obtained by a former expert witness for a private defendant. This is true for all three subsamples, regardless of whether the amount paid is revealed. The reduction in trust toward a former DoJ expert witness is only 8% if no compensation is mentioned. When the compensation is mentioned, the average reduction rises to 21%, driven mostly by ordinary Americans, who discount it by 25%.

When respondents were informed that the authors received $400K in compensation, they discounted the results more, regardless of whether the client was the DoJ (a reduction of 21%) or a private entity (a reduction of 36%). Yet, the discount is particularly pronounced when the client is private.

Panel B breaks down the economists' sample based on their expertise. In this IO-focused vignette, IO economists exhibit a smaller reduction in trust than their counterparts in finance. For example, for when they learn that the author of a study was previously an expert witness for a

private defendant who received $400K for his expert testimony, the average IO economist discount is 23%. Financial economists report a reduction in trust of 39%. This result suggests that financial economists are not particularly unaffected by conflicts of interest; they are only less affected by their own conflicts of interest (and so are IO economists).

This differential behavior across economists in different sub-disciplines can be driven by two factors. First, people who are familiar with a field have a better estimate of the frequency of conflicts their field, and thus, they are less surprised when the conflict is revealed. Second, researchers who are more familiar with a field underestimate the bias of their closest colleagues because this would be tantamount to admitting that they are biased themselves.

We can test the first hypothesis because we elicited trust in the result before the conflict was revealed. The trust of IO economists, financial economists, and the rest of the economists towards a finding that supermarket mergers are welfare decreasing is equal (1.97, 1.91, 1.98). When the finding is that the supermarket merger is welfare-increasing, the trust of IO economists is higher than the trust of financial economists and of the rest of the economists (2.56, 2.15, and 2.28). In the case of the profitable trading strategy, the initial trust of financial economists is 2.01–higher than that of IO economists (1.86) and of the rest of the economists (1.82). Thus, economists with more expertise trust the results in their field more than their peers (not less). This higher trust is inconsistent with their perceiving of a higher frequency of conflicts to start with. By exclusion, the lower trust reduction should be caused by economists underestimating the bias that conflicts generate in their own field rather than in other similar fields.

Overall, monetary incentives reduce trust by 27%, where this number is the average of the trading strategy vignette (with a reduction of 37%) and the supermarket vignette (with a trust reduction of 21%).

3.4.2 Career Incentives

CoIs are not limited to consulting fees and grants. Career concerns–which include monetary rewards such as higher salaries but also more prestige and influence– can also distort the incentives of academics; quid-pro-quo can be instantaneous or delayed. Such incentives are also hard to disclose. In such a complex situation, the survey method is superior to archival data analysis.

Our vignette described an article finding that CEOs are underpaid relative to the value they render to their companies. The main variation 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. Respondents distrusted an author with such clear career goals (an average trust reduction of 36%, see Figure 7), but the 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.

The reduction in trust is consistently larger for average Americans than for economists and is also larger for ordinary economists than for selected ones (who themselves are more likely to sit on corporate boards). Nevertheless, even selected economists severely distrust (36%) the results of an author potentially motivated by certain career aspirations.

3.4.3 Access to Data

The AEA disclosure policy requires authors to disclose if an interested party provided data for a research project. For this reason, in Figure 8, we test whether access to confidential data (and the forms of this access) matters for the trust readers have in the results of a paper.

The first treatment we consider is access to data. The specific vignette concerns a paper studying the welfare effects of the entry of a ride-sharing company in a given city. We introduce two randomized components. The first randomization concerns the content of the initial 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, and the results are of particular interest. 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 borderline statistically significant. On the other hand, 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; one with results favorable to the industry suffers twice as much–a 27% drop in trust. The trust result here is driven by economists. Average Americans do not seem to increase their trust when the results come from publicly available data, while economists do.

Potentially, the trust in public data might be driven by ideological factors. To investigate this aspect, in Figure 9, we decompose the revision in trust according to the economic ideology of the respondents (results are similar when we use the social ideology). We present the results only for selected economists (the others can be viewed 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 shows that the largest gain in trust among economists is from those who have previously used private data themselves, suggesting that they understand the downsides of working with private data.

The ride-sharing vignette compares publicly owned and publicly accessible data with privately owned and not publicly accessible data. While the two characteristics (ownership and access) tend to go hand-in-hand, there are plenty of cases 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).

Our conjecture is that a crucial dimension of trustworthiness is data accessibility: the extent to which a party with interest in the direction of the results can deny researchers access to the data. This restricted access generates the possibility of explicit and implicit quid-pro-quo. To test this conjecture we use a different vignette. The initial prompt describes the results of a paper studying the efficiency of a digital payments platform. To make it salient that the data are proprietary and investigate the effect of different access policies, we state that the data were given by a credit card

company. In one treatment, we reveal that the data provider has waived the right to review the paper before circulation; 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 in the two conditions. When a subject learns that the data provider has waived the rights, the trust in the results of the paper *increases* by an average of 12%. This increase is entirely 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 (55%) than for ordinary people (46%). The online appendix, Section 1.2 shows that those who increased their trust in the results the most are economists who had previously engaged in private consulting or used private data themselves (20% and 19% *increase* in trust, respectively), again suggesting that they understand the downsides of working with private data.

Overall, data-access conflicts result in reductions between 20% (an average of rideshare vignette, private data randomization, irrespective of findings) and 52% (an average of credit card vignette, right to review randomization).

These results underscore the significant impact that data access conditions can have on perceived trustworthiness. Yet, disclosure practices often lack specificity regarding data access terms. 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 normally subject to disclosure requirements. Yet, there are other important sources of external and internal pressure that can lead to bias. The first source, which we call "academic," regards 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 tell respondents about 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. Then, 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.

As Figure 11 shows, average Americans do not respond to the treatment. Their trust is reduced by 19% regardless of the treatment. By contrast, economists seem to respond to the treatment. If the paper finds results in contradiction with the previous work of the same author, trust *increases* by 4%, while it finds results in line with the previous work trust decreases by 5%. Therefore, economists are more likely to perceive an academic conflict of interest as being somewhat problematic than average Americans. Nevertheless, the magnitude of the conflict is an order of magnitude smaller than the results involving economic CoIs. Overall, in academic conflicts trust reduces by 12% when the results are in-line with the author's previous work.

3.4.5 Ideological Conflicts

Our final vignette examines the effect of ideological conflicts of interest. Previous research, Javdani and Chang, 2023 use a standard survey to show that ideology can bias perceptions in economics. 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 requirements to disclose this conflict in academic research, and such conflicts are rarely voluntarily disclosed by authors.

This vignette included two levels of randomization. First, respondents were presented with an article finding that access to abortion had either a positive impact on women's lifelong earnings or no impact. Then, the disclosure informed respondents that the author was either a Democrat or a Republican, having made significant donations to candidates from their respective party in recent election cycles.

As Figure 12 shows, the trust in the results drops an average 32% when a Republican author finds no impact of access to abortion on women's lifelong earnings (what we call a "Republican" finding). The CoI trust reduction is smaller, but still, an average of 19%, when a Democrat author

finds a positive impact (what we call a "Democratic" finding). Surprisingly, when a Democratic author presents a "Republican" finding, there is still an average trust reduction of 13%, but when a Republican author presents a "Democratic" finding, the average trust reduction is not statistically different from zero.

Examining this trust reduction across our different subgroups, we find that the general public has the highest trust reductions (from an average CoI trust reduction of 11% for a Republican author with a Democrat finding, to a 33% for a Republican author with a Republican finding). In contrast, trained economists experienced an *increase* in trust when a Republican author had a Democratic finding while discounting a Republican result of a Republican author by an average of 32%. In the case of Democratic author's findings, trained economists experienced an 11% decrease in trust when the Democratic author had a Democratic finding and a non-significant decrease in trust when a Democratic author had a Republican finding.

Panels A and B in Figure 13 decompose these results based on the respondent's political leaning. For brevity, we report only the results for Selected Economists. If the paper finds that access to abortion improves career outcomes, the fact that the author is Republican increases everybody's trust, particularly at the extreme of the political spectrum. If the author is a Democrat, it decreases everybody's trust, particularly among very conservative people. If the paper finds that access to abortion has no effect on career outcomes, the fact that the author is Republican decreases everybody's trust a great deal (on average, 30%). Yet, the fact that the author is a Democrat does not increase people's trust; in fact, there is still a mild reduction for the more conservative respondents. Thus, the effect is not perfectly symmetric. Conservatives seem to mistrust Democrats, no matter what the results are. While ultraliberals revise positively their prior when a Republican author finds a result more favorable to the Democratic perspective. Overall, ideological CoIs reduce trust by 17% on average.

In sum, while political conflicts of interest of the author are important, respondents' ideology affects trust regardless of any conflict of interest. On the one hand, Republicans distrust Democrats regardless of the result they find. On the other hand, Democrats believe the results they like even in the face of a conflict of interest. The implication is that both Republicans and Democrats think conflicts of interest are important. Yet, for Republicans, ideological conflicts are more important

than economic ones, while for Democrats, economic conflicts are more important than political ones.

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 the 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 CoIs. 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 the expected bias from the likelihood of conflicts, the CoI Discount allows us to evaluate the broader impact of CoIs 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 α is the typical significance level (usually α = 0.05), $1-\beta$ denotes the typical power of an experimental design, and π 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 increases my prior that quantitative easing increases GDP) or negative (e.g., a paper decreases 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, as a decrease in the prior that quantitative easing increases GDP can be cast as an increase in the prior that quantitative easing does not increase GDP.

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

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

Conflicted research, however, where researchers have financial, professional, or ideological stakes in the outcome, may be more likely to report significant findings—not because the underlying effects are stronger, but rather due to subtle biases introduced at various stages of the research process. Bias arises from a combination of factors related to study design, data collection, analysis, and presentation that can generate significant findings in cases where there should not be any (Ioannidis, 2005). Equation (1) does not account for these potential biases and needs to be adjusted. To address this, 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). We modify the model to capture this increased likelihood as:

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

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

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 CoIs, which they took into account when forming their beliefs. Let λ 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:

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

P_0 is the first response to each vignette. 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:

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

r corresponds to what we called earlier the CoI Trust Reduction and is the second response to each vignette. It measures participants' perceived loss of trust when informed of the specific conflict. This trust reduction combines the impact of conflicts of interest on the posterior (PSP^{bias}) with the expected frequency of these conflicts (λ). Ideally, we want to separate the two. As the formula shows, r is decreasing in λ; thus, the higher the subjects expect the presence of a conflict, the lower r will be.

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 the empirical evidence of a non-conflicted paper increases readers' prior beliefs. For non-conflicted research, the value is proportional to how strongly the PSP shifts the reader's priors.

For a conflicted paper, however, the presence of a conflict introduces potential 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^{bias}. The value of a conflicted paper is thus:

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

To estimate the extent to which a conflicted paper is worth, we calculate its relative value

compared to that of a non-conflicted paper (the ratio of Equation (6) to Equation (2)):

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

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

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

The $CoIDiscount$ quantifies the proportion of the paper's value lost due to the presence of a CoI. This discount is directly tied to the reduction in trust, r, which we derive from our survey data in Equation (5). From Equation (5), we can rewrite the $CoIDiscount$ as:

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

where λ represents the expected frequency of conflicted papers in the relevant academic literature.

Intuitively, 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. It then 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.

One critical insight from this framework is that the CoI Discount typically exceeds the reduction in trust r whenever the expected frequency of conflicted papers λ is greater than zero. This suggests that, in environments where conflicts of interest are prevalent, even small reductions in trust can lead to significant devaluations of research papers. The higher the frequency of conflicted papers, the more heavily discounted each individual conflicted paper becomes, compounding the overall loss of credibility in the field.

Our survey provides direct measures of the reduction in trust r. By combining these trust reductions with estimates of the frequency of conflicted papers λ, we can calculate the CoI Discount for various types of conflicts. We do this in Section 5, after we test the model predictions in our

data.

4.1 Bringing the Data to the Model

The responses to each vignette in our survey are mapped directly into the model. The first answer (Prior Trust) corresponds to P_0 and the second answer (CoI Trust Reduction) corresponds to r, given by Equations (4) and (5) respectively. Since we record several answers provided by the same individual, we can run regressions between our observables, P_0 and r, that fully account for individual fixed effects. But what does the model predict about this association?

In principle, the shocks (between vignettes but within individual) could come from any parameter in the model, that is, α, β, π, u, or λ. Since we told participants that every result was published in a highly reputable journal, and subject to standard refereeing standards, we can safely assume that α and β do not change between vignettes, for a given individual. Similarly, the fraction λ of conflicted papers can also be reasonably held constant for a certain respondent. Instead, the shocks to P_0 and r between the vignettes certainly come from the different CoIs that were revealed in the second question (which would affect u), and possibly also from the prior beliefs that an individual has over the association they have to evaluate in the first question (which would affect π).

The model gives precise predictions. An increase in u reduces PSP^{bias} (see Equation (3)), and thus it reduces P_0. But the reduction in the latter is moderated by $\lambda < 1$ (Equation (4)). Thus, from Equation (5), we would see an increase in r. Hence the model predicts a negative correlation between P_0 and r when the same individual evaluates vignettes with different intensity of CoI. Similar considerations can be made when the shock between vignettes concerns π, also leading to a negative correlations between our observables. Because of the randomization in the survey, the distributions of the variables u and π are orthogonal to each other, hence no correlation between them is expected. Thus the model has an unambiguous prediction: the responses given by a survey respondent, P_0 and r, should be *negatively correlated*.

We already knew of this negative correlation from the aggregate data (see Section 3.3.1). This prediction arises now directly from the model, and we can test it in a much more powerful way. In Table 7 we present the results. The negative correlation between P_0 and u is extremely robust,

both overall and in each sample separately, even after including full individual FEs. The effect is largest in absolute value for Americans, followed by ordinary economists, and finally by selected economists.

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

So far, we have shown that the presence of a conflict of interest reduces the trustworthiness of a conflicted paper by roughly one-third. Now, we want to calculate how much CoIs reduce the overall value of a research paper. As outlined in Section 4, our theoretical framework allows us to translate reductions in trust into a measurable reduction in the value of a paper, which we refer to as the CoI Discount. To calculate this discount, we first need to estimate the frequency of conflicted papers, as this directly impacts the expected likelihood that a given research paper is conflicted.

Since the survey took place in October 2023, we assume that participants, if acting rationally, form their expectations about the presence of conflicts based on the frequency of conflicted studies in economics at that time. To estimate this frequency, we analyzed papers published in the American Economic Review (AER) and The RAND Journal of Economics (RAND) between 2019-2023, focusing on author's voluntary disclosure of conflicts of interest and the conditions under which the accessed databases used in their projects.

To identify the presence of conflicts of interest in published articles we relied on two different data sources, which differ in their quality. For the AER, we collected the disclosure statements available in the AER website for all papers published between 2019-2023 (1280 individual disclosure statements), the *Readme* files for their replication packages, and the first footnote of each paper. We then manually reviewed each disclosure statement to identify the source of research grants [10], consulting fees, employment outside of academia, explicit editorial rights to review the

[10] We coded potential conflicts conservatively and only considered as leading to conflict research grants by parties directly interested in the results of the article, therefore excluding sources such as the National Science Foundation or private organizations such as JPAL, Russell Sage and others. We consider grants as conflicted when they are given by a private or public party with an interest in the results of the article–e.g. grants by Google on topics related to digital

results of the article before publication and discretionary and gated data sources (meaning data that cannot be accessed freely or upon payment) [11]. Access to the Readme files is important because, most scholars do not disclose the terms of access to data in their disclosure statements, but they have to list the availability of data in their replication packages. For the RAND journal, only the first footnote is available–so we replicate the same analysis restricting ourselves to this footnote. Finally, to understand political affiliations, we first collect data on academic economists who were appointed to high-level federal offices.[12] We then collected data on donations by academics and economists to candidates running for federal office, between the years 2000-2024. The source of the latter was the FEC. We then merge both datasets with our AER and RAND data. We only include academics who cumulatively donated more than $10,000 to candidates running for federal offices, to keep the analysis in line with our survey questionnaire.

As Table 8 shows, overall, 38.8% of the articles published in the AER and 28.7% of the articles published in RAND over the 2019-2023 period present some form of conflict. The two most common sources of conflict are the reliance on gated data sources that cannot be shared or replicated (21.4% of AER articles and 17.1% of RAND articles) and authors formally employed by a company with an interest in the results of the article (15.8% of articles in the AER and 12.8% of articles in RAND).

Table 9 reports the CoI Discount under both assumptions about the frequency of conflicted papers–as per our model in Part II. On average, the CoI discount is 39%. However, this average masks substantial variation depending on the type of conflict.

For example, research grants by a directly interested party reduce a paper's value by 34%,

markets, the Kayser Foundation on topics related to healthcare

[11] for example, we did not consider as "gated and discretionary" private but paid databases such as Compustat, because anyone can access them upon payment independent of the research topic. On the other hand, we included databases where there is vetting of research topic before access, such as private companies such as UBER, Microsoft or Meta, or public databases such as the CENSUS bureau or the FED)

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

while consulting arrangements with a directly interested party result in a discount of 35%. A paper authored by an antitrust expert who has previously testified for the Department of Justice retains 80% of the value of a non-conflicted paper. However, if the expert testified for a private defendant and was highly paid, the value drops significantly to 63%, reflecting a 37% discount.

Career concerns, particularly those involving aspirations for board membership or executive roles, are equally problematic. A paper on executive compensation written by an author seeking a board position is worth only 53% of what a non-conflicted paper would be worth (47% discount). This effect is mitigated (but not eliminated) if the position is obtained two years after the publication of the article (34% discount).

The use of confidential data also leads to a significant reduction in value. A paper using proprietary data is worth only 76% of a paper using publicly available data. When the provider of the confidential data retains the right to review results prior to publication, the value of the paper falls dramatically to just 42% of a similar paper where the right to review has been waived. This is the largest drop in our estimates–reflecting the increasingly crucial role data plays in modern economics research.

Finally, ideological conflicts were considered less severe than financial or career-related conflicts, but could still result in meaningful discounts. A paper written by a prominent author with a big reputation at stake on a topic should be discounted by 13%. A paper on abortion written by an author with a political interest at stake in this debate should be discounted by 16% or 18%, depending on whether the author is Democrat or Republican.

6 Perception and Reality

In the prior sections, we established that research produced by authors with a CoI is met with reduced trust by both economists and the general public. This CoI trust reduction is consistent across different groups, indicating that it is not merely a reflection of generalized mistrust in science among certain segments of the population. Nevertheless, the evidence presented thus far does not allow us to determine whether this CoI trust reduction is warranted. To assess the justification for this discount, we next examine the actual bias in different studies involving CoI in different fields

of economics and in the medical literature–which has been studying the topic for a long while–and then compare it with the levels of discounting we observe.

6.1 Bias in Research with CoI: Economics

The practice of CoI disclosure is much more recent in economics than in medicine–where it started half a century ago. Except for the Review of Financial Studies, which began requiring disclosure of authors' conflicts in 2006, the first leading economic journal to do so was the American Economic Review in 2012. The others followed shortly afterward. As a result, there are much fewer studies of the effects of conflicts of interest in economics compared to medicine.

Zingales, 2013 looks at the effect of conflicts of interest on economists' opinions on executive compensations. By looking at the opinions of experts included in the University of Chicago Kent A. Clark Center for Global Markets' panel, he finds that experts who served on a corporate board were four times more likely than those who didn't to disagree with the statement, "The typical chief executive officer of a publicly traded corporation in the U.S. is paid more than his or her marginal contribution to the firm's value." He also finds that articles published in major economics journals are more likely to find that the current level of CEO compensation is justified and less likely to conclude that the CEO salaries are too high. Unfortunately, these findings cannot easily be mapped into a CoI trust reduction.

Asatryan et al., 2020 report that, on average, government-financed research projects find 33% larger fiscal multipliers than unfunded studies. Fabo et al., 2021 show that papers co-authored by central bank employees estimate the effect of QE on both output and inflation to be larger than papers where no author was affiliated with a central bank. All central-bank-affiliated papers report a significant effect of QE on output, while only 50% of the unfunded ones did: Central-bank-affiliated papers report a standardized peak effect of QE on output equal to 0.28, while non-affiliated ones only of 0.11.

6.2 Bias in Research with CoI: Medicine

The medical field has a far more extensive body of research examining the effects of CoIs compared to economics. By 2006, the year when the Review of Financial Studies was the first major economics journal to establish a CoI disclosure policy, about 90% of all high-impact medical journals already had a CoI disclosure policy for authors (Blum et al., 2009). Beyond influencing research findings on the effectiveness of industry-sponsored drugs and medical devices (Lundh et al., 2017), CoIs also affect favorable recommendations in clinical guidelines, advisory committee reports, opinion pieces, and narrative reviews (Nejstgaard et al., 2020).

To systematically identify relevant studies, we performed a comprehensive literature search following a structured methodology. We started from a group of seed articles identified on Google Scholar and in the citation network of Oostrom, 2024.[13] Then, using the PubMed API, we performed a systematic literature search of all the Medical Subject Headings (MeSH) keywords in the seed articles. The first such search yielded 27,140 publications.

To refine our sample, we applied a three-step filter process. First, we dropped all papers published in journals with a Scimago score below the score of the British Medical Journal (BMJ).[14] This narrowed down our sample to 4,394 articles. Second, we asked GPT-4o to filter out papers that did not focus on our area of interest.[15] GPT returned a total of 224 articles. Finally, we manually reviewed these 224 articles to confirm their relevance, ultimately narrowing our selection to 38 articles that specifically addressed the impact of industry sponsorship on research findings. Using this refined set of 38 articles as a new seed, we extracted all their MeSH keywords and restarted the search of all articles containing a keyword that appeared in at least 20% of the new seed articles. We repeated the procedure until no new keywords emerged, concluding with a final set of 53 articles, including Oostrom, 2024. To the best of our knowledge, this is largest meta-analysis of conflicts of interest in medical publications to date.

[13] Cho and Bero, 1996, Kjaergard and Als-Nielsen, 2002, Lexchin et al., 2003, Bero et al., 2007, Bourgeois et al., 2010, Knox et al., 2000.

[14] For more information, see Scimago Journal Rankings.

[15] The prompt is "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 does study this correlation, return 1. If it does not, return 0."

These 53 articles can be broadly grouped into four categories: Observational Studies (ObS), Meta-analyses of Meta-analyses (MoM), Pharmacoeconomics (Ph-E), and Randomized Clinical Trials (RCTs). Within these broad categories, there are differences in the underlying research topic (Oncology, Cardiovascular, Psychiatric, Diabetes, etc.).

As Figure 15 shows, the conflicts of interest in medicine are mainly of two types: funding of a study by a drug manufacturer with an interest in the drug or authors' financial interest in the success of a drug. For the purpose of the following analysis, we treat the two as equivalent. Figure 15 reports the odds ratios for all these studies' findings. The odds ratio is 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.

In ObS, the odds ratio is 22.4. In other words, conflicted observational studies have odds of a significant result that are 22.4 times those of non-conflicted ones. To have a more intuitive measure, we need to make some assumptions about the odds of non-conflicted studies. If we assume that the probability that a non-funded study finds a significant result is 50% (and thus the odds are 1), then the odds of a funded study is 95.7%, thus almost 46 p.p. more.[16]

In MoM studies, the odds ratio is 1.37, which means that the odds of a conflicted study are 1.37 times the odds of a non-conflicted one. Again, if we assume that the probability that a non-funded

[16] If $OR = \frac{\frac{P(S|C)}{1-P(S|C)}}{1} = 22.4$, then $P(S|C) = 0.957$.

study finds a significant result is 50% (and thus the odds are 1), then the odds of a funded study is 58% – almost 8 percentage points more. Ph-E studies are also prone to be biased in favor of drug manufacturers sponsoring the study with an odds ratio of 2.09, which translates to a 67% chance of a pharmacoeconomic study reporting relatively lower costs to life (in terms of Quality Adjusted Life Years) when a drug manufacturer funds it.

The bias in drug manufacturer-sponsored RCTs is smaller than in MoM, and Ph-E studies. The combined odds ratio here is close to 1.26, resulting in a 56% chance of pro-manufacturer findings, i.e., only 6 percentage points more than the non-funded studies. One potential reason behind this could be the objectivity of RCTs, where results are less susceptible to underlying assumptions since RCTs need to be pre-registered using accepted protocols. Ph-E, MoMs, and to some extent, ObS have subjective components that can be tweaked by slightly perturbing the framework.

Overall, combining the results of all papers across the categories, we find an odds ratio of 1.39.[17] In other words, on average, conflicted studies are 8% p.p. more likely to obtain a pro-manufacturer finding than non-conflicted studies.

6.3 Bayesian Updating

One way to further validate our results is to compare the CoI Trust Discount 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 Fabo et al., 2021, the conflict 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 (π) 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

[17]Note that combining the odd ratio does not simply mean averaging them, since they have different degrees of precision. The metan package in Stata was utilized to combine the odds ratios from multiple studies due to its ability to perform inverse-variance weighting, a precise method for meta-analysis. Inverse-variance weighting calculates a pooled odds ratio by assigning weights to each study's effect size based on the inverse of its variance, giving more weight to studies with greater precision (Borenstein et al., 2009). This method ensures that larger, more reliable studies have a greater influence on the combined effect size.

frequency of conflicted studies (λ) 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 (α) of 0.05. Substituting these values into Equation (5), we calculate a CoI trust reduction equal to 16%.[18]

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

In both the computer science and economics literature, there is an active discussion about how GPT and other large language models can act as a rational economic agent (Horton, 2023, Chen et al., 2023, Korinek, 2023, Kim et al., 2024). These articles find that, when properly prompted, GPT models "are mostly rational and even score higher than human decisions" when making risk, time, social and food decisions (Chen et al., 2023). Thus, we use GPT-models as an alternative benchmark to check the consistency of our results.

In Table 12, we report the CoI trust reduction computed by GPT-4 Omni when we confront it with the same questions we asked in the survey. We report the CoI trust reduction by treatment and we compare the GPT-4 Omni result with the average survey result and with the average result in each of the three subsamples. The GPT-4 Omni CoI trust reduction is a bit smaller than that of average Americans but aligned with that of economists. The only exception is the vignette about the Alzheimer's drug. GPT-4 Omni seems more optimistic about the cost of CoI in that situation.

Then, we construct a vignette and question based on Fabo et al., 2021 and interrogate GPT-4 Omni. We start by informing GPT-4 Omni about a paper reporting a significant impact of QE on output and asking for its initial beliefs on the results according to the 5-point Likert scale used in our

[18] Conducting a robustness check based on the values of β reported by Maniadis et al., 2014, we obtain CoI trust reductions in the range of 16% and 18%.

survey.[19] Then, we follow up with the disclosure that the authors were affiliated with a central bank that implemented quantitative easing, and finally ask for the resulting impact on trustworthiness according to our 7-point Likert scale.

Following 1,000 simulations of the vignette and question, GPT-4 Omni's CoI trust reduction is approximately 20%, or around the "somewhat undermines my trust in the results". This is in line with the 16.4% result obtained through Bayesian updating.

6.5 Comparing Trust Reduction with Reductions in Citations

As a final validation test, we contrast our findings with those of Leuz et al., 2023. They investigate whether the disclosure of financial ties with industry impacts the citation rates of medical journal articles. They utilize a large dataset comprising over 17,000 research and review articles published in seven prominent medical journals between 1988 and 2008. To address the issue of selection bias, where higher-quality researchers may be more likely to have industry ties, they implement several strategies, including controlling for observable measures of article quality and analyzing the effect of disclosure within the same author's work.

One of their key analyses involves examining the impact of CoI disclosures on the likelihood of an article being recommended by the University of Chicago's Priority Updates from the Research Literature (PURL) program. This program identifies and disseminates important research studies to family physicians. Leuz et al., 2023 findings indicate that articles with disclosed conflicts are less likely to be recommended in PURL. Specifically, they find that disclosing industry ties decreases the likelihood of a conflicted paper being included in the PURL program by 7% to 16.5%, depending on the controls applied. Assuming that the evaluators at PURL also adjust their beliefs in response to a conflict of interest, similar to our survey respondents (which include academic experts), we can infer that the CoI trust reduction by the evaluators ranges between 7% and 16.5%. These figures are lower than the estimates produced by our survey–where the average CoI Discount for the economics vignettes is of 39%.

[19] We follow our survey's structure, including the provision to the GPT model of the same assumptions about the quality of the paper as in the main survey–paper published in a highly reputed peer-reviewed journal.

However, there are reasons to believe that the estimates provided by Leuz et al., 2023 understate the true impact of CoI disclosures. First, while the recommendation system used by the authors is a clever way to account for quality differentials between articles, it is not a perfect control–meaning that the articles not included may still be of a lower quality. In contrast, we can control for article quality in our survey through randomization of only the potential conflict of interest. Second, the decision to exclude a paper from the PURL list is an observable action, potentially exposing decision-makers to retaliation from the excluded paper's authors. Given that authors with disclosed conflicts often wield significant influence, this concern might lead evaluators to be more cautious in applying the trust reduction, thereby reducing the observed effect. Considering these factors, it is reasonable to suggest that the true impact of CoI disclosures on the perceived trustworthiness of research may be greater than what is observed in Leuz et al., 2023's PURL test.

If our interpretation for the difference between our finding and Leuz et al., 2023 is correct and reputation considerations prevent citations from fully adjusting for the trustworthiness of the paper, then citations discount underestimates the drop in perceived quality caused by a CoI.

7 Implications

Our findings have implications that extend beyond academic circles. Scholars have long documented the influence of conflicted research on policy decisions in areas such as drug approval processes (McGarity and 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 studies to render judicial decisions that have far-reaching effects on public life (e.g., McIntire and Kantor, 2024). Companies have also sponsored judicial training programs and engaged with academics to influence courts and public policy more broadly (Ash et al., 2022; Lancieri et al., 2023

Given the influence of many academic studies, it is unsurprising that there are strategic efforts to shape their results. For instance, a Congressional investigation revealed that oil companies partnered with universities to influence research recommendations. Internal Shell emails disclosed that the company's funding of the Futures Lab at Imperial College London was part of a "Global

Methane Communications Plan," overseen by Shell's general manager for gas advocacy, aimed at producing research that would "underpin the role for gas" in energy transitions.[20]

Such influence often extends beyond financial support, and increasingly rely on the control of access to data. Companies and public institutions, like the Federal Reserve, offer selective access to proprietary databases, often with the goal of promoting particular narratives that align with their interests (Berg and Johnston, 2019) while restricting access to data that might support conflicting perspectives (Horan, 2019; Wagner, 2023; Zingales, 2019). A whistleblower leak of Uber's internal communications offers a telling example: after completing a commissioned study, an academic requested access to Uber's data for an independent, unpaid study. In internal discussions,Uber executives expressed concerns that granting access would cause them to "lose editorial control," but a senior staff member reassured them, saying, "*We see low risk here because we can work with [the person] on framing the study and we also decide what data we share with him*" (Lawrence, 2022).

Thus, the key question is not whether strategic attempts to bias research exist, nor whether these biases matter for policy— these are well-documented. The real concern is whether such biases are limited to policy papers written by think tanks or whether they permeate peer-reviewed academic research, particularly in economics, just as they do in medicine. Studies by Asatryan et al., 2020 and Fabo et al., 2021 provide concrete evidence of bias in peer-reviewed economic research, and our work builds on this by showing that both economists and the general public recognize this bias as pervasive. The result is a systemic erosion of trust in academic research as a whole (Lipton et al., 2016), as noted by former U.S. Assistant Attorney General for Antitrust Jonathan Kanter in a recent speech. He pointed out that conflicts of interest among antitrust scholars have contributed to a "*pervasive breakdown in the distinction between expertise and advocacy in competition policy*," fueling a "*seeping distrust of expertise by the courts and by law enforcers*" (Kanter, 2024). Ironically, this growing distrust disproportionately harms independent scholars engaged in unbiased research who do not obtain the benefits associated with CoIs (funding, better data, etc.) but are impacted by the decreases in the trust in science that they generate (a topic we

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

return below).

The question, then, becomes: what can be done to mitigate conflicts of interest and the biases they introduce? Our paper does not undertake a full cost-benefit analysis of conflicted research. In many instances, this relationships generate benefits to science: research investigations can be so expensive or specialized that they are only feasible with private funding, proprietary data or the knowledge generated by these relationships–the benefits of conflicted research remain an open question. Access to private data, for example, can provide significant insights that would otherwise remain out of reach for the scientific community (e.g. González-Bailón et al., 2023; Guess et al., 2023. This means that banning conflicted research outright is probably not socially optimal–our results should *not* be interpreted as a call to prohibit private funding, consulting, or the use of confidential data in research.

Yet, even without conducting a full cost-benefit analysis, our results and analytical framework indicate ways to reduce the negative impact of conflicts of interest on trust in academic research.

7.1 Credibility of Disclosure

Our findings suggest that detailed and credible disclosure of CoIs can reduce, and in some cases eliminate, the "cross-subsidization" between conflicted and non-conflicted papers. For instance, in our credit card vignette, respondents were initially unaware of the conditions under which the credit card company provided the data to the researchers. When it is revealed that the company waived its right to review the paper before publication, trust in the paper's finding increased by 17%; however, when it was revealed that the company retained that right, trust fell by 52%. This demonstrates that without detailed and credible disclosure, trust in conflicted papers is artificially inflated, effectively "subsidized" by the trust in non-conflicted articles. This dynamic is not lost on conflicted parties themselves. A Wall Street Journal investigation revealed that some conflicted authors deliberately withheld disclosure of their CoI in many articles to maintain the perceived credibility of their research and preserve their ability to influence policymakers (Mullins, 2024).

There are no restrictions on disclosing the presence or absence of CoI. Thus, non-conflicted researchers could theoretically deviate from this equilibrium by unilaterally disclosing a lack of

conflict (e.g., stating that the company providing data had no right to review the paper). Such disclosure, however, only works if it is credible. To assess the credibility of CoI disclosures, we surveyed economists with two specific questions. First we asked, "Do you think that your past disclosure of potential conflicts of interest has been adequate?" Respondents could choose from: Yes, Mostly, Sometimes, Occasionally, and No. As reported in Figure 16, 78% of respondents answered "Yes" and 17% "Mostly." Thus, 95% of economists believe their own disclosure practices are largely effective.

Next, we asked, "Do you think that, *at present*, academic economists generally disclose potential conflicts of interest adequately?" The response options were the same, plus "I do not know." Only 2% answered "Yes," while 43% selected "Mostly," 26% chose "Sometimes," 13% said "Occasionally," and 5% responded "No." An additional 10% answered, "I do not know." This contrast indicates that while economists feel confident about their own disclosures, they have little trust in the disclosure practices of their peers. This disconnect highlights the current system's lack of credibility.

We suspect that the root of this issue lies in the lack of enforcement. Punishments for non-disclosure are rare and not well-publicized. Without credible enforcement, conflicted papers continue to impose a negative externality on non-conflicted ones. It is therefore crucial that journals and the broader academic community develop a credible enforcement mechanism to ensure the integrity of CoI disclosures.

7.2 Disclosure of the Terms of Data Access

The conditions under which researchers gain access to such data are often unclear. While some academic institutions and journals have formal rules prohibiting data providers from retaining the right to block publication, it is uncertain whether these rules are consistently followed. The data agreements between the researchers/universities and the data providers are often confidential (many times even to those within the university who are working with the datasets). This lack of transparency undermines trust in the results of papers that rely on proprietary data, as it leaves open the possibility of selective data sharing or result shaping based on data provider preferences.

Confidentiality in data access agreements is a direct threat to the credibility of academic research

and should not be tolerated. Academic journals should not accept papers where the data providers have a right to review the paper for reasons other than maintaining the confidentiality of the data.

Academic journals should also require that data agreements (with the possible exception of payment terms) be made public as a pre-condition for the publication of the article. Such a policy would not only enhance transparency but also increases the bargaining power of researchers vis-à-vis data-supplying institutions–and in doing so, ultimately safeguard the integrity of academic research.

7.3 Prospective Conflicts

Our paper highlights an important and often-ignored conflict of interest: future career prospects. These conflicts are particularly difficult to disclose and therefore pose a more severe problem, especially when potential employers are concentrated in a few institutions. As Kempf, 2020 demonstrates, career concerns do not distort behavior when multiple employers can observe a potential hire's past actions. In the absence of these alternatives, however, such as in the case of concentrated industries or academia, career concerns can exert significant pressure on individuals to skew their results. This highlights a major cost of industrial concentration: "He who pays the piper calls the tune." When there is only one potential payer, there is only one tune, which ultimately destroys the diversity of ideas in the marketplace of knowledge.

7.4 Discounting Publications

Finally, our analysis reveals that, in the marketplace of ideas, a conflicted paper is worth, on average, 61% of a non-conflicted one. This problem is exacerbated by the asymmetric incentives created by conflicts of interest. Academics internalize the benefits of the conflicted research, such as the prestige of publishing papers based on proprietary data or the monetary gains from consulting and research grants. However, the trust deficit created by these conflicts is not internalized, leading to a systemic problem in which the negative externalities of conflicted research affect the entire academic community. One particularly troubling consequence is how conflicts of interest shape the direction of academic research. Papers that could be highly impactful and in the public interest

may never be written because they would alienate potential funders or require access to private data that is not made available.

Despite this imbalance, in literature reviews, legal settings, and academic promotions papers are currently evaluated based mostly on journal prestige and citation counts. 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 leads to an overproduction of conflicted papers from a societal perspective. However, there is a straightforward remedy: systematically discounting conflicted papers in literature reviews, legal proceedings, and academic promotions. By applying a consistent discount to conflicted research, we can better align the private incentives of academics with the broader public interest and ensure that academic research continues to serve as a reliable source of knowledge. Our results provide an initial benchmark for how such discounting could be implemented in practice, offering a path forward for reducing the overproduction of conflicted research.

8 Conclusion

In this paper, we examine how various types of conflicts of interest impact the perceived trustworthiness and overall value of economic research. Our survey demonstrates that both the general public and leading economists view conflicted research with skepticism—even when it is published in prestigious journals. We formalize the relationship between trust and the value of a paper in the presence of CoIs. We quantify the reduction in value due to conflicts, the CoI Discount, which reflects the broader social cost of conflicted research. Given the increasing mistrust in expertise, this critical issue cannot be fully addressed through disclosure alone. Our findings suggest that the current disclosure system, while necessary, is insufficient to mitigate the negative effects of CoIs on the credibility of academic research.

One benign interpretation of our results is that users fully discount the effect of CoI; thus, we should not worry about it. This conclusion is inaccurate for several reasons. First, non-

credible disclosure allows conflicted research to impose a large externality on non-conflicted work by diminishing trust in the field as a whole. Second, cognitive dissonance biases distort rational discounting precisely where it is most needed: among the experts of a specific sub-field. Finally, in many areas (from promotions to literature reviews), papers are generally counted, not weighed. If they are weighed, they are weighed by citations, which do not seem to properly discount conflicted papers. Hence, the marketplace of ideas overvalues conflicted research. As the reliance on evidence-based policy continues to grow (Haskins, 2018), these distortions can have first-order welfare effects.

In computing our estimates of the social costs of conflicted research, we assumed fully rational and perfectly informed actors. This assumption, which is conservative, likely leads to an underestimation of the true societal costs of conflicted research. In practice, conflicted research is frequently employed to "fool" people (Michaels, 2008; Valletti, 2024), suggesting that the broader societal costs may be greater. Future work should aim to quantify these broader impacts, further illuminating the consequences of CoIs on public trust and decision-making.

Finally, conflicts of interest represent just one form of distortion in the marketplace of ideas more broadly and the academic sphere more specifically. Identifying and addressing these distortions is a complex yet essential task. Failing to tackle these issues will continue to erode trust in both the general public and expert communities regarding the credibility and value produced by academic research (Kanter, 2024). This reinforces the need to address the distortions caused by CoIs and other biases in academic work: above all, we hope that the results of this article will trigger a continued, open conversation on how CoIs impact science and what scholars and society more broadly can do to mitigate this negative effects.

References

Adda, J., Decker, C., & Ottaviani, M. (2020). "P-hacking in clinical trials and how incentives shape the distribution of results across phases". *Proceedings of the National Academy of Sciences*, *117*(24), 13386–13392.

Adekunle, L., Chen, R., Morrison, L., Halley, M., Eng, V., Hendlin, Y., Wehner, M. R., Chren, M.-M., & Linos, E. (2020). "Association between financial links to indoor tanning industry and conclusions of published studies on indoor tanning: Systematic review". *The British Medical Journal*, *368*.

Ahn, R., Woodbridge, A., Abraham, A., Saba, S., Korenstein, D., Madden, E., Boscardin, & Keyhani, S. (2017). "Financial ties of principal investigators and randomized controlled trial outcomes: Cross sectional study". *The British Medical Journal*, *356*.

Als-Nielsen, B., Chen, W., Gluud, C., & Kjaergard, L. L. (2003). "Association of Funding and Conclusions in Randomized Drug Trials : A Reflection of Treatment Effect or Adverse Events?" *The Journal of the American Medical Association*, *290*(7), 921–928.

Aneja, A., Esquitin, R., Shah, Iyengar, R., Nisenbaum, R., Melo, M., Matthewkutty, S., Sethi, S. S., Mamdani, M., & Farkouh, M. E. (2013). "Authors' self-declared financial conflicts of interest do not impact the results of major cardiovascular trials". *Journal of the American College of Cardiology*, *61*(11), 1137–1143.

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*.

Bariani, G. M., Ferrari, A. C. R., Hoff, P. M., Krzyzanowska, M. K., & Riechelmann, R. P. (2013). "Self-reported conflicts of interest of authors, trial sponsorship, and the interpretation of editorials and related phase iii trials in oncology". *Journal of Clinical Oncology*, *31*(18), 2289–2295.

Barnes, D. E., & Bero, L. A. (1998). "Why Review Articles on the Health Effects of Passive Smoking Reach Different Conclusions". *The Journal of the American Medical Association*, *279*(19), 1566–1570.

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*.

Bekelman, J. E., Li, Y., & Gross, C. P. (2003). "Scope and impact of financial conflicts of interest in biomedical research: A systematic review". *The Journal of the American Medical Association*, *289*(4), 454–465.

Bell, C. M., Urbach, D. R., Ray, J. G., Bayoumi, A., Rosen, A. B., Greenberg, D., & Neumann, P. J. (2006). "Bias in published cost effectiveness studies: Systematic review". *The British Medical Journal*, *332*(7543), 699–703.

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, Anglemyer, A., Vesterinen, H., & Krauth, D. (2016). "The relationship between study sponsorship, risks of bias, and research outcomes in atrazine exposure studies conducted in non-human animals: Systematic review and meta-analysis". *Environment International*, *92-93*, 597–604.

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.

Bes-Rastrollo, M., Schulze, M. B., Ruiz-Canela, M., & Martinez-Gonzalez, M. A. (2013). "Financial conflicts of interest and reporting bias regarding the association between sugar-sweetened beverages and weight gain: A systematic review of systematic reviews". *PLoS Medicine*, *10*(12).

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.

Booth, C. M., Cescon, D. W., Wang, L., Tannock, I. F., & Krzyzanowska, M. K. (2008). "Evolution of the randomized controlled trial in oncology over three decades". *Journal of Clinical Oncology*, *26*(33), 5458–5464.

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.

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.

Brown, A., Kraft, D., Schmitz, S. M., Sharpless, V., Martin, C., Shah, & Shaheen, N. J. (2006). "Association of industry sponsorship to published outcomes in gastrointestinal clinical research". *Clinical Gastroenterology and Hepatology*, *4*(12), 1445–1451.

Chartres, N., Fabbri, A., & Bero, L. A. (2016). "Association of industry sponsorship with outcomes of nutrition studies: A systematic review and meta-analysis". *JAMA Internal Medicine*, *176*(12), 1769–1777.

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).

Chiu, K., Grundy, Q., & Bero, L. (2017). "'spin' in published biomedical literature: A methodological systematic review". *PLoS Biology*, *15*(9).

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 Souza Gutierres, B., Aguiar, P. N., Dourado, B. B., Alves, A. L., Matsas, S., Simões, A. R., & Del Giglio, A. (2020). "Evidence strength of pharmaceutical industry-funded clinical trials in metastatic nsclc: A comparison with other sources of funding". *Journal of Thoracic Oncology: Official Publication of the International Association for the Study of Lung Cancer*, *15*(7), 1170–1176.

Djulbegovic, Lacevic, M., Cantor, A., Fields, K., Bennett, C., Adams, J., Kuderer, N., & Lyman, G. (2000). "The uncertainty principle and industry-sponsored research". *Lancet*, *356*(9230), 635–638.

Downing, N. S., Cheng, T., Krumholz, H. M., Shah, & Ross, J. S. (2014). "Descriptions and Interpretations of the ACCORD-Lipid Trial in the News and Biomedical Literature: A Cross-Sectional Analysis". *The Journal of the American Medical Association : Internal Medicine*, *174*(7), 1176–1182.

Dunn, A. G., Arachi, D., Hudgins, J., Tsafnat, G., Coiera, E., & Bourgeois, F. T. (2014). "Financial conflicts of interest and conclusions about neuraminidase inhibitors for influenza: An analysis of systematic reviews". *Annals of Internal Medicine*, *161*(7), 513–518.

Dunn, A. G., Zhou, Hudgins, J., Arachi, D., Mandl, K. D., Coiera, E., & Bourgeois, F. T. (2016). "Financial competing interests were associated with favorable conclusions and greater author productivity in nonsystematic reviews of neuraminidase inhibitors". *Journal of Clinical Epidemiology*, *80*, 43–49.

Ebrahim, S., Bance, S., Athale, A., Malachowski, C., & Ioannidis, J. P. A. (2016). "Meta-analyses with industry involvement are massively published and report no caveats for antidepressants". *Journal of Clinical Epidemiology*, *70*, 155–163.

Edelson, L., Graef, I., & Lancieri, F. (2023, February). "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.

Flacco, M. E., Manzoli, L., Boccia, S., Capasso, L., Aleksovska, K., Rosso, A., Scaioli, G., De Vito, C., Siliquini, R., Villari, P., & Ioannidis, J. P. (2015). "Head-to-head randomized trials are mostly industry sponsored and almost always favor the industry sponsor". *Journal of Clinical Epidemiology*, *68*(7), 811–820.

Friedberg, M., Saffran, B., Stinson, T. J., Nelson, W., & Bennett, C. L. (1999). "Evaluation of Conflict of Interest in Economic Analyses of New Drugs Used in Oncology". *The Journal of the American Medical Association*, *282*(15), 1453–1457.

Fugh-Berman, A., McDonald, C. P., Bell, A. M., Bethards, E. C., & Scialli, A. R. (2011). "Promotional tone in reviews of menopausal hormone therapy after the women's health initiative: An analysis of published articles". *PLoS Medicine*, *8*(3).

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

Garne, D., Watson, M., Chapman, S., & Byrne, F. (2005). "Environmental tobacco smoke research published in the journal indoor and built environment and associations with the tobacco industry". *Lancet (London, England)*, *365*(9461), 804–809.

Gartlehner, G., Morgan, L., Thieda, P., & Fleg, A. (2010). "The effect of study sponsorship on a systematically evaluated body of evidence of head-to-head trials was modest: Secondary analysis of a systematic review". *Journal of Clinical Epidemiology*, *63*(2), 117–125.

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.

Guelimi, R., Afach, S., Bettuzzi, T., Meyer, A., Padern, G., Yiu, Z., Naudet, F., Sbidian, E., & Le-Cleach, L. (2024). "Funding and conclusions of network meta-analyses on targeted therapies in inflammatory diseases: An overview". *Journal of Clinical Epidemiology*, *172*.

Guess, A. M., Malhotra, N., Pan, J., Barberá, P., Allcott, H., Brown, T., Crespo-Tenorio, A., Dimmery, D., Freelon, D., & Gentzkow, M. (2023). "How do social media feed algorithms affect attitudes and behavior in an election campaign?" *Science*, *381*(6656), 398–404.

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.

Hartmann, M., Knoth, H., Schulz, D., & Knoth, S. (2003). "Industry-sponsored economic studies in oncology vs studies sponsored by nonprofit organisations". *British Journal of Cancer*, *89*(8), 1405–1408.

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.

Jagsi, R., Sheets, N., Jankovic, A., Motomura, A. R., Amarnath, S., & Ubel, P. A. (2009). "Frequency, nature, effects, and correlates of conflicts of interest in published clinical cancer research". *Cancer*, *115*(12), 2783–2791.

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". *The British Medical Journal*, *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.

Krauth, D., Anglemyer, A., Philipps, R., & Bero, L. (2014). "Nonindustry-sponsored preclinical studies on statins yield greater efficacy estimates than industry-sponsored studies: A meta-analysis". *PLoS Biology*, *12*(1).

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).

Lane, P. W., Higgins, J. P. T., Anagnostelis, B., Anzures-Cabrera, J., Baker, N. F., Cappelleri, J. C., Haughie, S., Hollis, S., Lewis, S. C., Moneuse, P., & Whitehead, A. (2013). "Methodological quality of meta-analyses: Matched-pairs comparison over time and between industry-sponsored and academic-sponsored reports". *Research Synthesis Methods*, *4*(4), 342–350.

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

Lesser, L. I., Ebbeling, C. B., Goozner, M., Wypij, D., & Ludwig, D. S. (2007). "Relationship between funding source and conclusion among nutrition-related scientific articles". *PLoS Medicine*, *4*(1), 41–46.

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". *The British Medical Journal*, *326*(7400), 1167–1170.

Liang, F., Zhu, J., Mo, M., Zhou, C. M., Jia, H. X., Xie, L., Zheng, Y., & Zhang, S. (2018). "Role of industry funders in oncology rcts published in high-impact journals and its association with trial conclusions and time to publication". *Annals of Oncology: Official Journal of the European Society for Medical Oncology*, *29*(10), 2129–2134.

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

Lipton, E., Confessore, N., & Williams, B. (2016). "Think Tank Scholar or Corporate Consultant? It Depends on the Day". *The New York Times*.

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

Lundh, A., Lexchin, J., Mintzes, B., Schroll, J. B., & Bero, L. (2018). "Industry sponsorship and research outcome: Systematic review with meta-analysis". *Intensive Care Medicine*, *44*(10), 1603–1612.

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". *The British Medical Journal*, *371*.

Neumann, P. J., Sandberg, E. A., Bell, C. M., Stone, P. W., & Chapman, R. H. (2000). "Are pharmaceuticals cost-effective? a review of the evidence". *Health Affairs (Project Hope)*, *19*(2), 92–109.

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*.

Østengaard, L., Lundh, A., Tjørnhøj-Thomsen, T., Abdi, S., Gelle, M. H. A., Stewart, L. A., Boutron, I., & Hróbjartsson, A. (2020). "Influence and management of conflicts of interest in randomised clinical trials: Qualitative interview study". *BMJ*, *371*.

Peppercorn, J., Blood, E., Winer, E., & Partridge, A. (2007). "Association between pharmaceutical involvement and outcomes in breast cancer clinical trials". *Cancer*, *109*(7), 1239–1246.

Perlis, R. H., Perlis, C. S., Wu, Y., Hwang, C., Joseph, M., & Nierenberg, A. A. (2005). "Industry sponsorship and financial conflict of interest in the reporting of clinical trials in psychiatry". *American Journal of Psychiatry*, *162*(10), 1957–1960.

Ridker, P. M., & Torres, J. (2006). "Reported Outcomes in Major Cardiovascular Clinical Trials Funded by For-Profit and Not-for-Profit Organizations: 2000-2005". *The Journal of the American Medical Association*, *295*(19), 2270–2274.

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.

Schillinger, D., Tran, J., Mangurian, C., & Kearns, C. (2016). "Do sugar-sweetened beverages cause obesity and diabetes? industry and the manufacture of scientific controversy". *Annals of Internal Medicine*, *165*(12), 895–897.

Siena, L. M., Papamanolis, L., Siebert, M. J., Bellomo, R. K., & Ioannidis, J. P. A. (2023). "Industry Involvement and Transparency in the Most Cited Clinical Trials, 2019-2022". *The Journal of American Medical Association Network Open*, *6*(11).

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.

Stelfox, H. T., Chua, G., O'Rourke, K., & Detsky, A. S. (1998). "Conflict of interest in the debate over calcium-channel antagonists". *New England Journal of Medicine*, *338*(2), 101–106.

Tibau, A., Bedard, P. L., Srikanthan, A., Ethier, J.-L., Vera-Badillo, F. E., Templeton, A. J., Ocaña, A., Seruga, B., Barnadas, A., & Amir, E. (2015). "Author financial conflicts of interest, industry funding, and clinical practice guidelines for anticancer drugs". *Journal of Clinical Oncology*, *33*(1), 100–106.

Valachis, A., Polyzos, N. P., Nearchou, A., Lind, P., & Mauri, D. (2012). "Financial relationships in economic analyses of targeted therapies in oncology". *Journal of Clinical Oncology*, *30*(12), 1316–1320.

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.

Wang, McCoy, C. P., Murad, M. H., & Montori, V. M. (2010). "Association between industry affiliation and position on cardiovascular risk with rosiglitazone: Cross sectional systematic review". *The British Medical Journal*, *340*.

Xie, F., & Zhou, T. (2022). "Industry sponsorship bias in cost effectiveness analysis: Registry based analysis". *BMJ (Clinical research ed.)*, *377*.

Yank, V., Rennie, D., & Bero, L. A. (2007). "Financial ties and concordance between results and conclusions in meta-analyses: Retrospective cohort study". *BMJ (Clinical research ed.)*, *335*(7631), 1202–1205.

Zhang, L., Dai, F., Brackett, A., Ai, Y., & Meng, L. (2018). "Association of conflicts of interest with the results and conclusions of goal-directed hemodynamic therapy research: A systematic review with meta-analysis". *Intensive Care Medicine*, *44*(10), 1638–1656.

Zhou & Xie, F. (2023). "Sponsorship bias in oncology cost effectiveness analysis". *Journal of Clinical Epidemiology*, *156*, 22–29.

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 v/s 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 (formerly IGM).

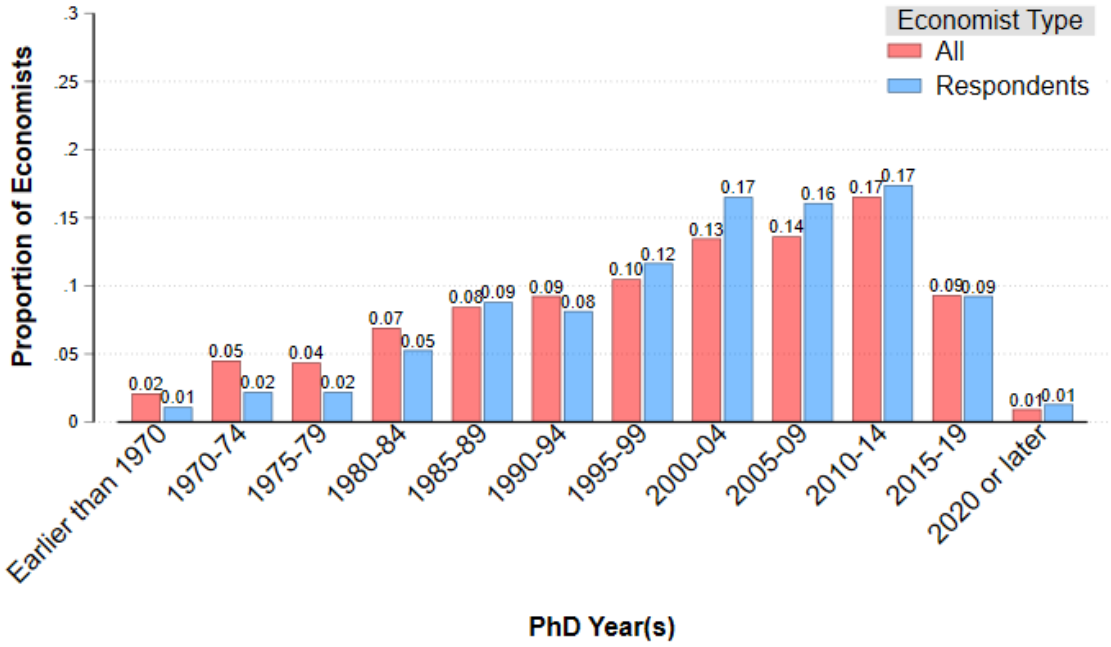


Figure 2: **Priming Bias Conditional on Position of the Vignette**

Figure 2 Panel A exhibits the average prior trust contingent on vignette positioning in the survey. Panel B shows average trust reduction contingent on vignette positioning. Overall in both panels, there is no significant difference between average responses conditional on vignette positioning which hints at minimal bias, if any.

A. Prior Trust



B. CoI Trust Reduction

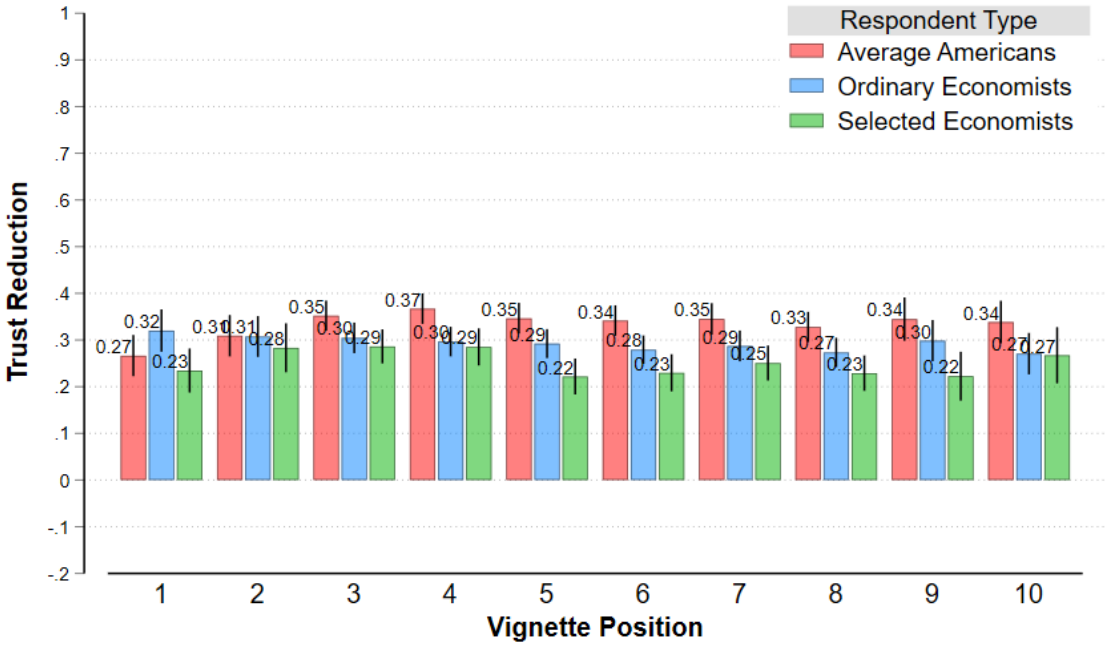
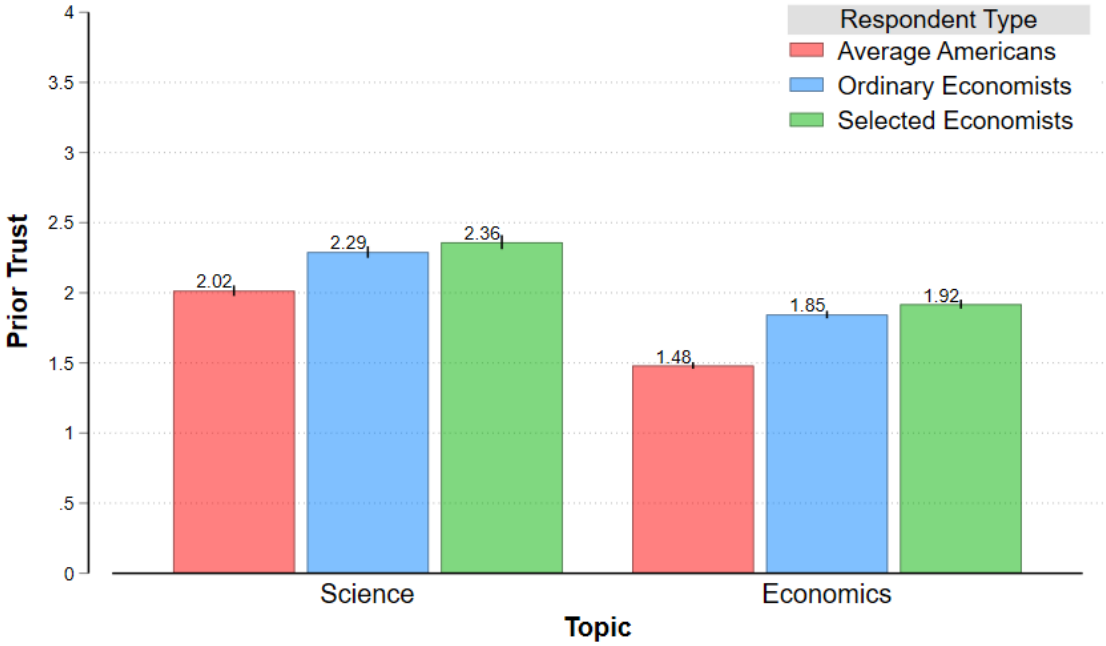


Figure 3: **Prior Trust and CoI Trust Reduction Across Topics and Respondent Types**

Figure 3, Panel A illustrates respondents' pre-conflict of interest (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-conflict of interest (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

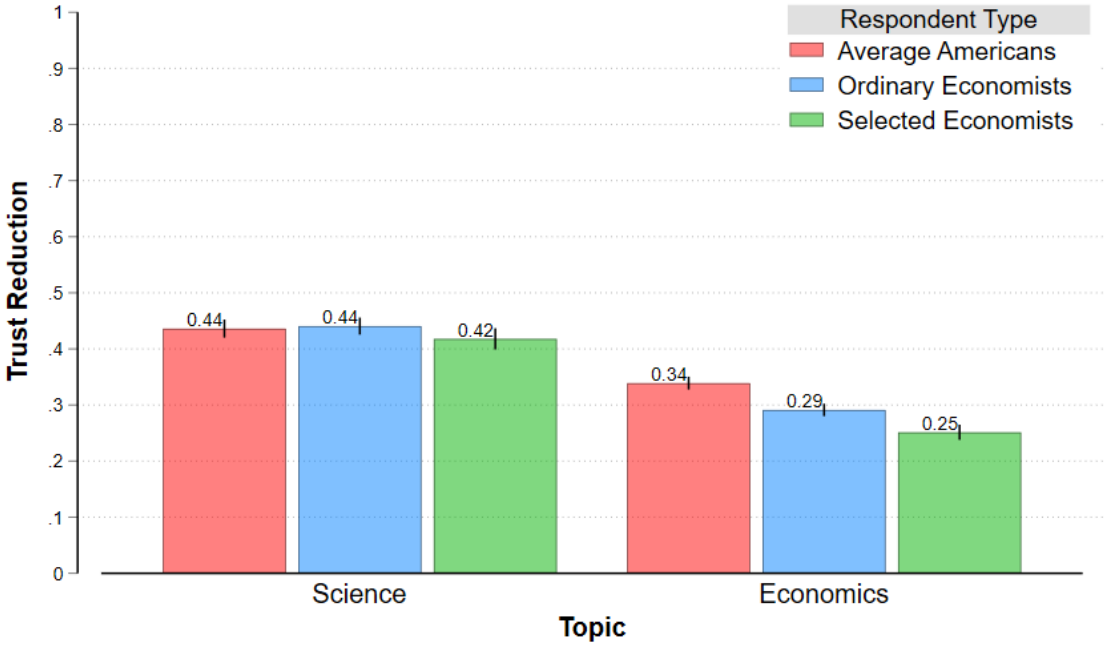


Figure 4: **CoI Trust Reduction v/s Prior Trust**

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-conflict of interest (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.

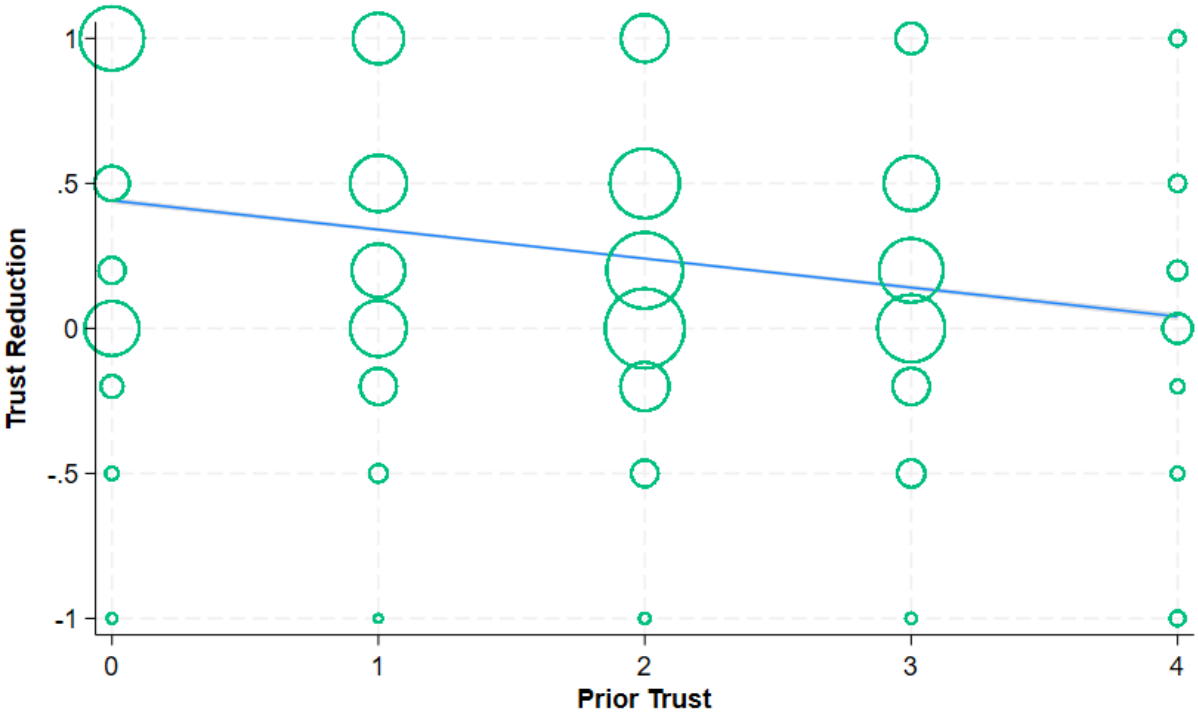
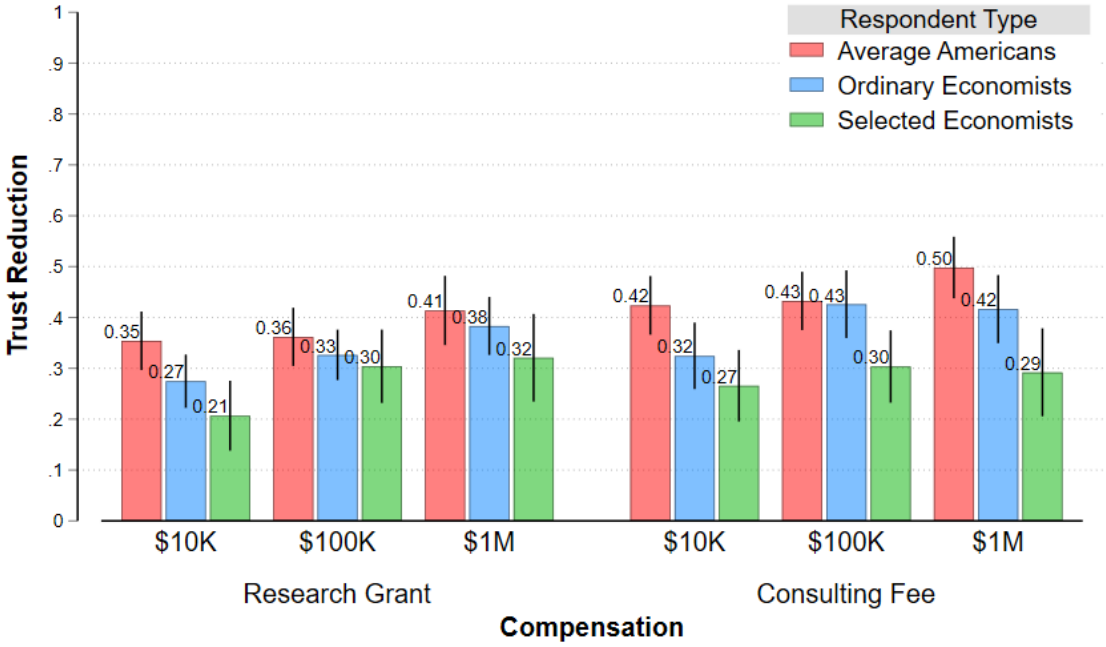


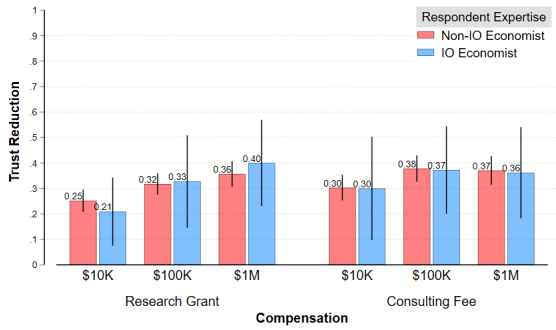
Figure 5: **CoI Trust Reduction in the Presence of Monetary Incentives: Trading Strategy**

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 exhibits decomposes the results for economists by their self reported expertise. For the sake of the trading strategy vignette, we only focus on IO and Finance experts.

A. All Three Samples



B1. Only IO Economists



B2. Only Financial Economists

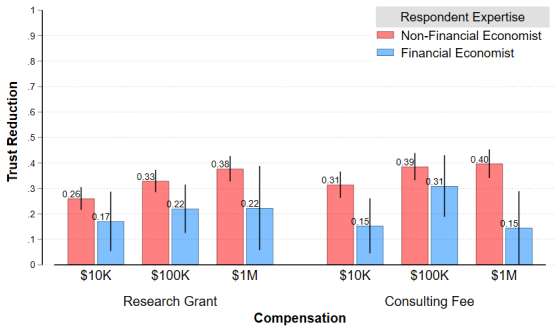
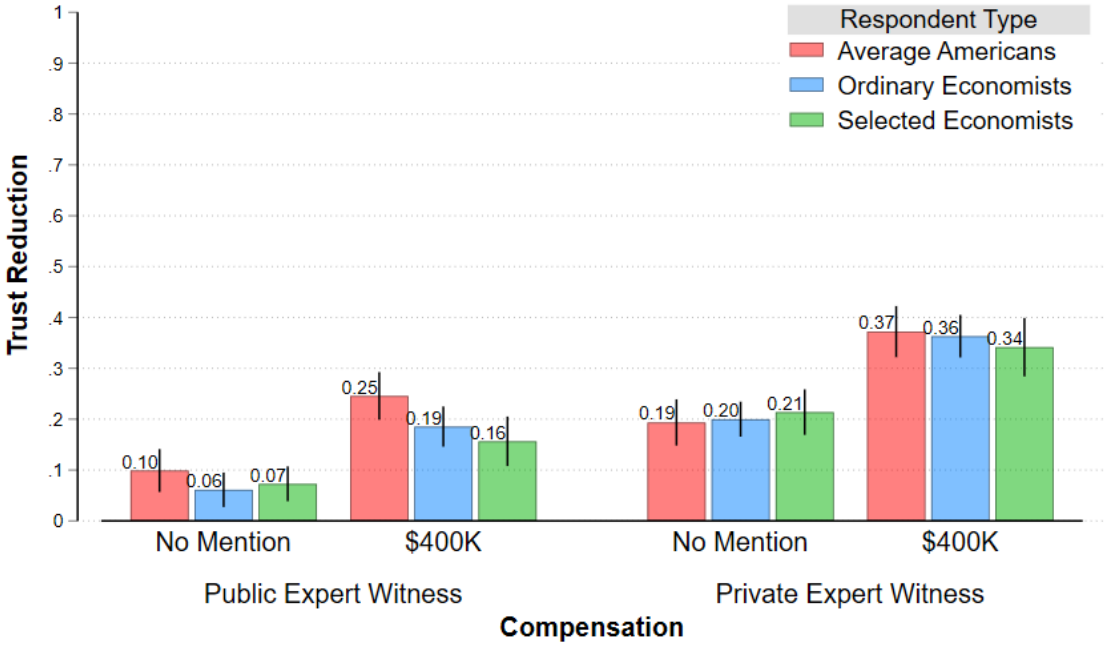


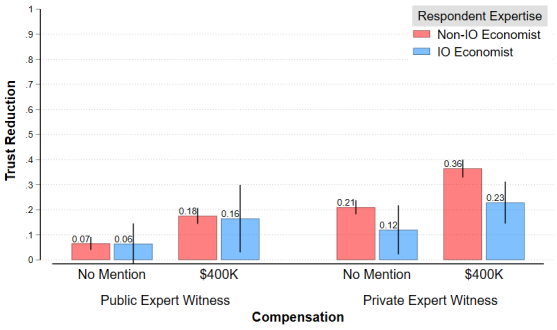
Figure 6: **CoI Trust Reduction in the Presence of Monetary Incentives: Supermarket Vignette**

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 for two regional supermarkets (Department of Justice Antitrust Division) involved in a merger challenged by the Department of Justice Antitrust Division. Additionally, we inform the respondent that the author was either not compensated or compensated $400,000 for the testimony. Figure 6 Panel B explores the monetary variation in the Supermarket Vignette across economists' self-reported expertise/specialization. For the sake of the supermarket vignette, we only focus on IO and Finance experts.

A. All Three Samples



B1. Only IO Economists



B2. Only Financial Economists

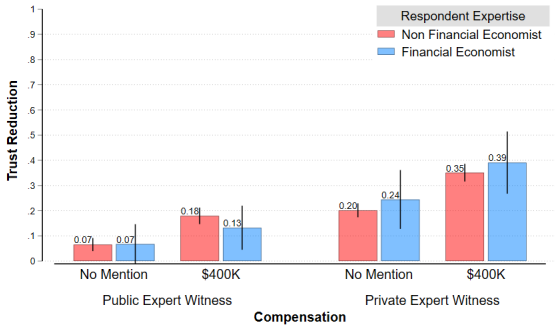


Figure 7: **CoI Trust Reduction in the Presence of Career Incentives**

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.

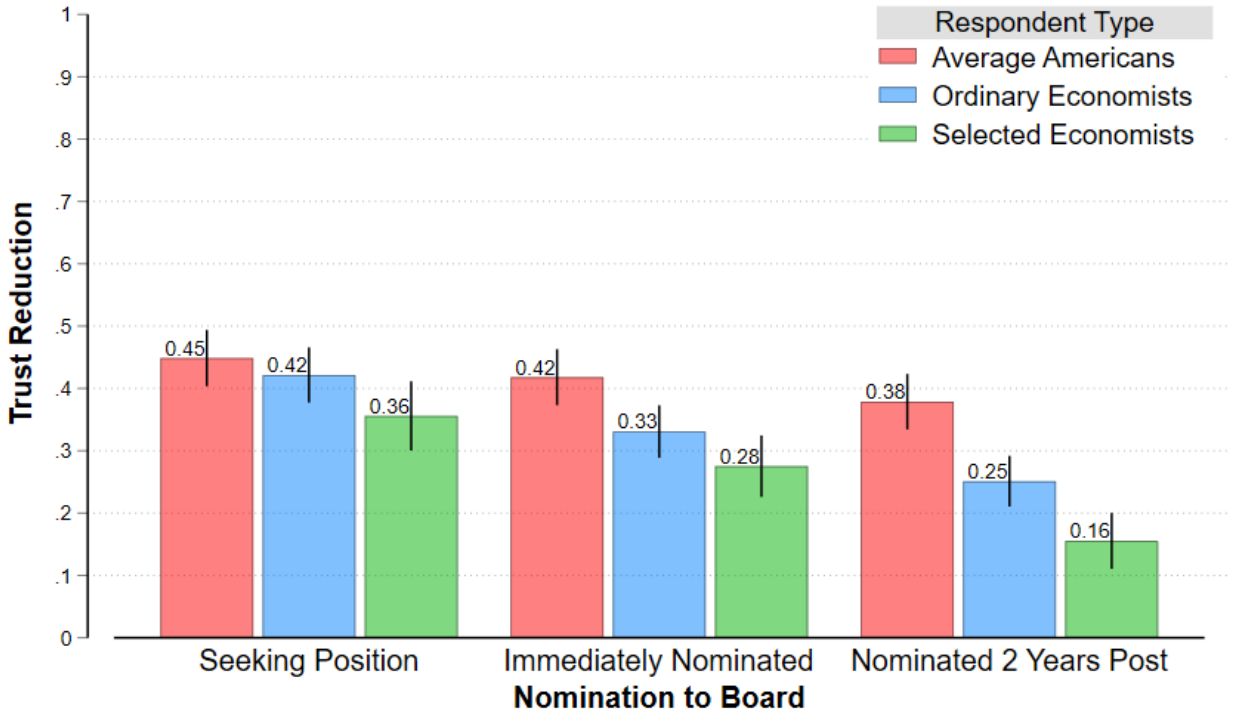


Figure 8: **CoI Trust Reduction in the Presence of Data Restrictions : Ridesharing**

Figure 8, Panel A shows the average 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.

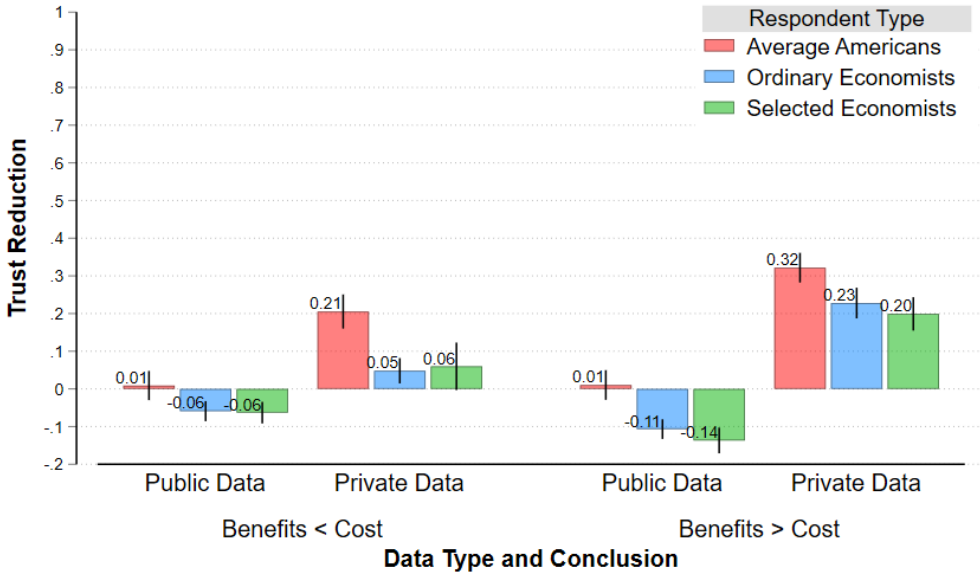
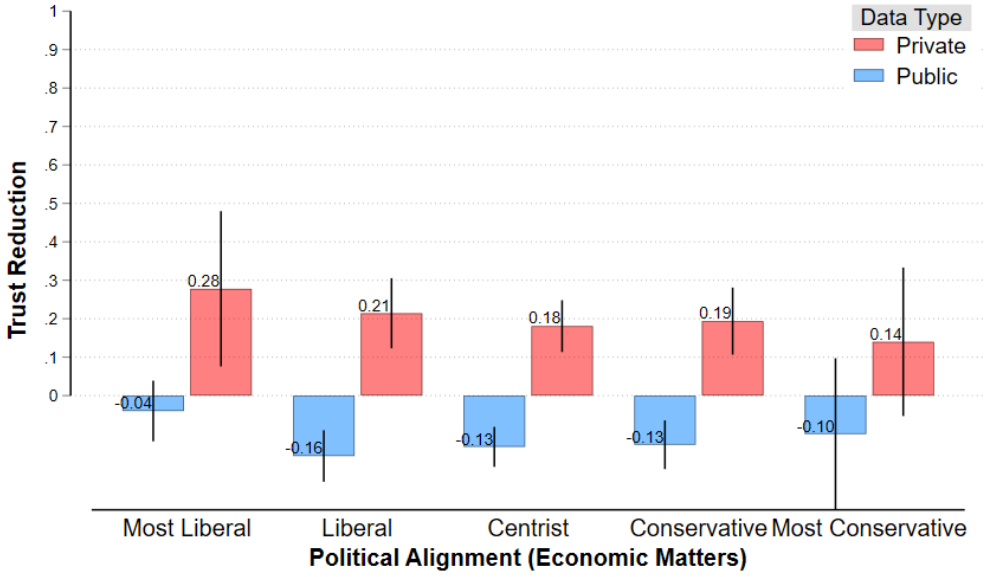


Figure 9: **Political Alignment and CoI Trust Reduction: Ridesharing x Selected Economists**

Figure 9, Panel A illustrates the CoI trust reduction in the ridesharing vignette across the political alignment of selected economist respondents on economic matters. 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, Panel B exploits the latter randomization and 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

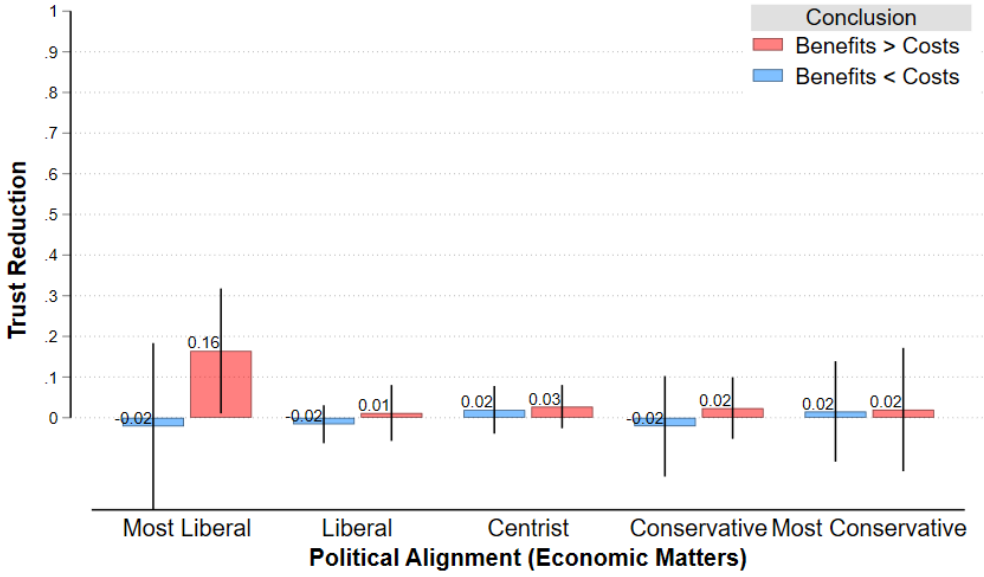


Figure 10: **CoI Trust Reduction in the Presence of Data Restrictions: Credit Card**

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 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

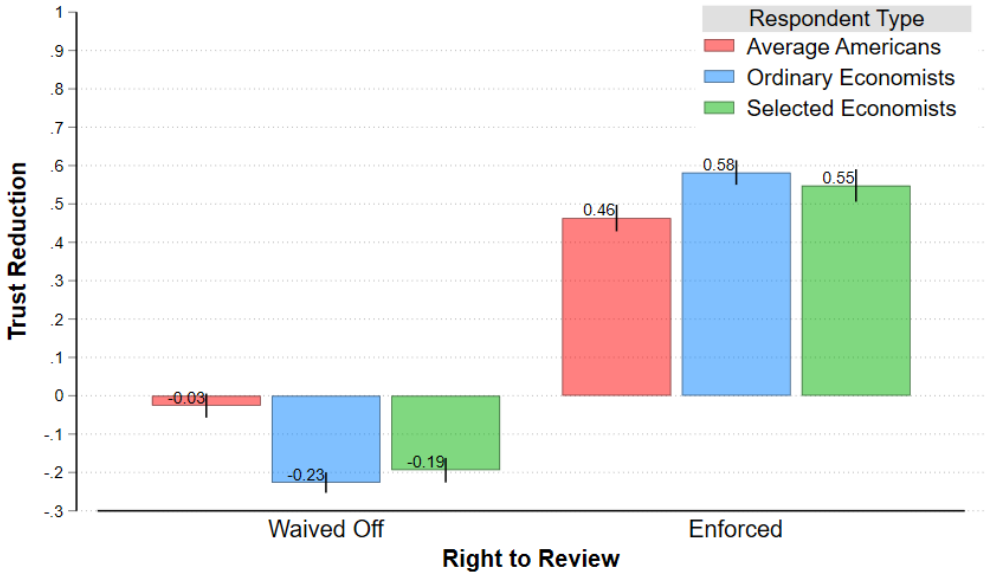


Figure 11: **CoI Trust Reduction in the Presence of Academic Conflicts: Tax Policy**

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.

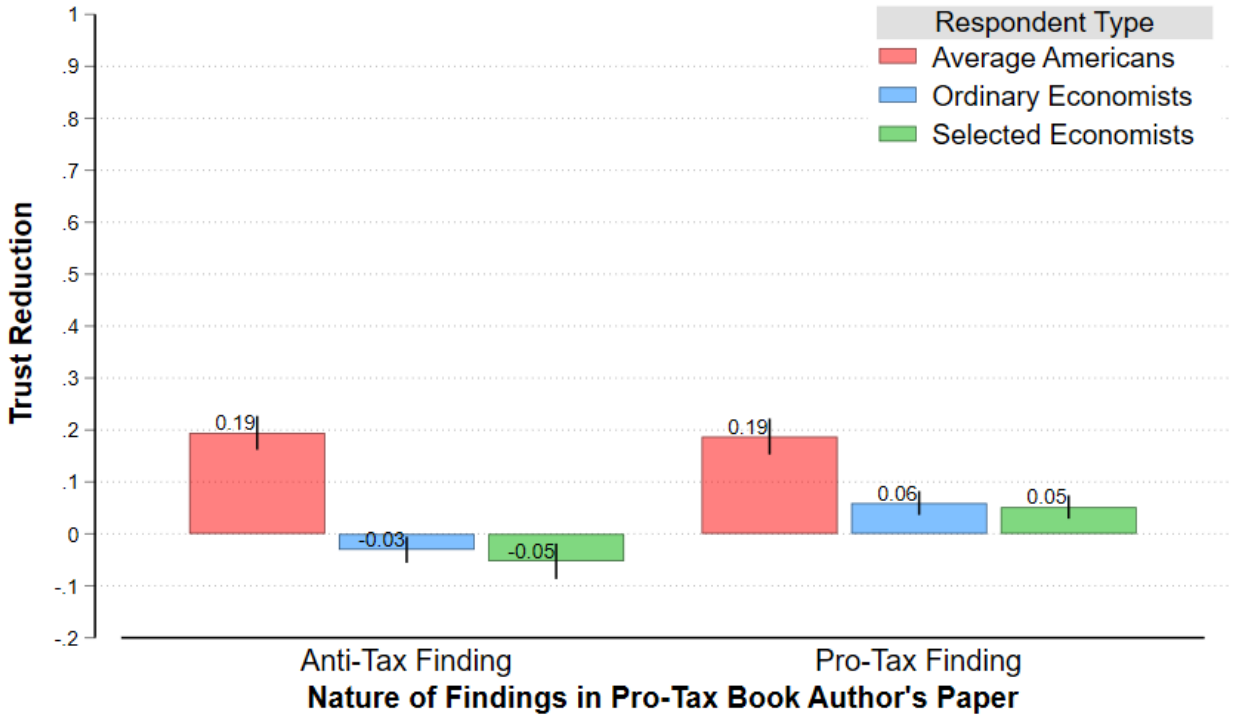


Figure 12: **CoI Trust Reduction in the Presence of Political Conflicts**

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.

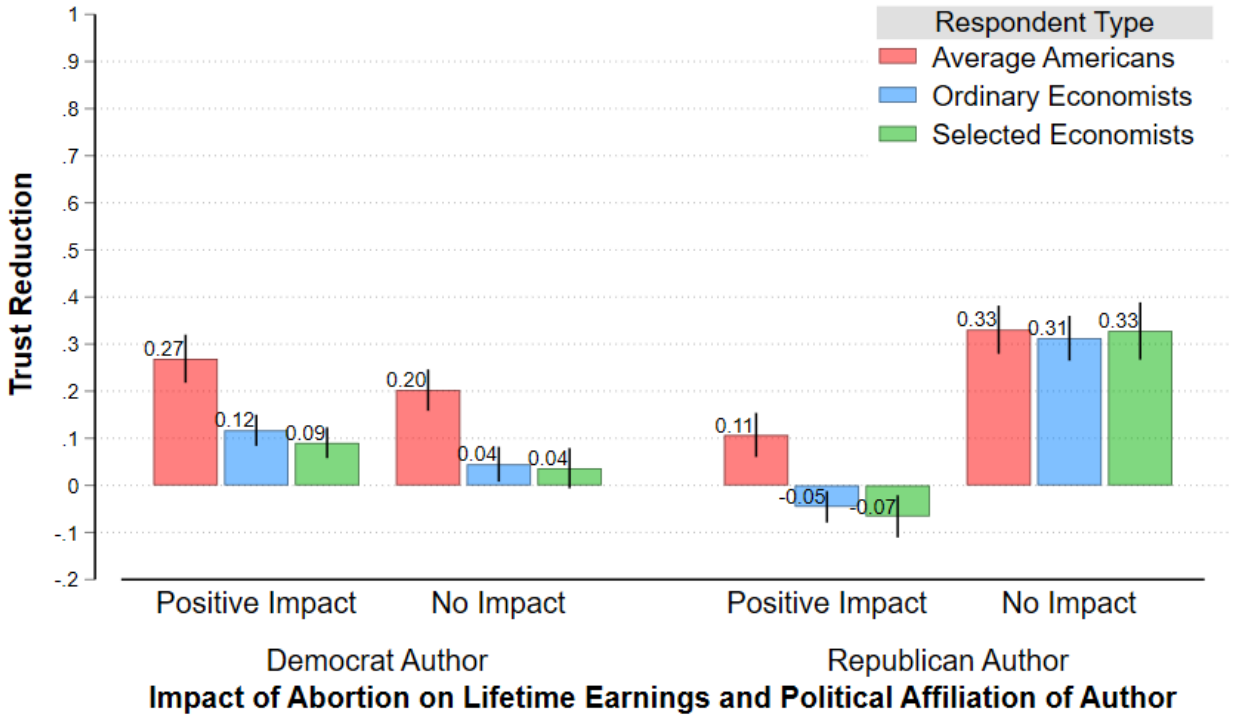
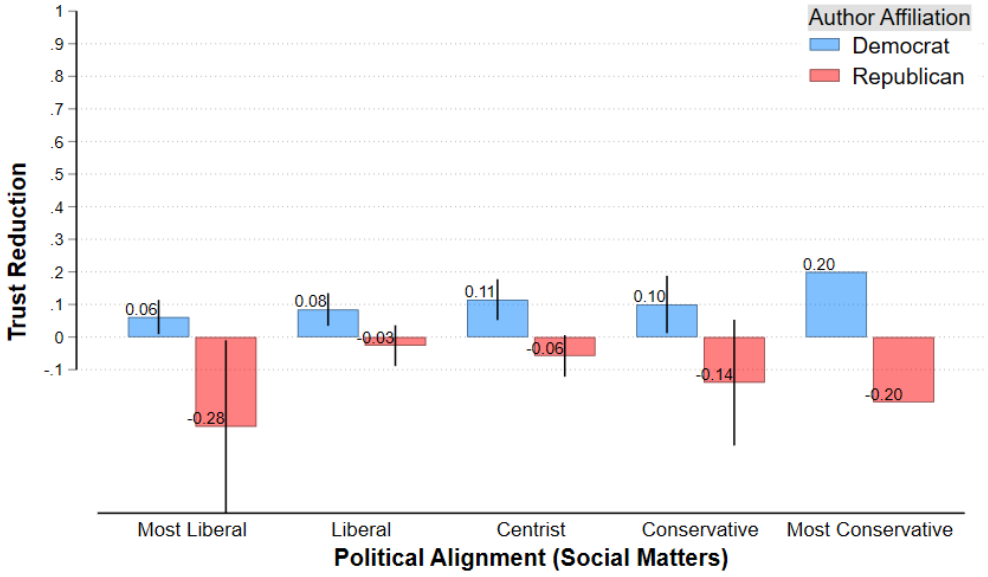


Figure 13: **Political Alignment and Trust Reduction: Abortion Policy, Selected Economists**

Figure 13 illustrates the CoI trust reduction in the abortion vignette across selected economist respondents' political alignment on social matters. In this vignette, respondents are initially presented with the findings of a study featuring 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, irrespective of the study's findings. 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

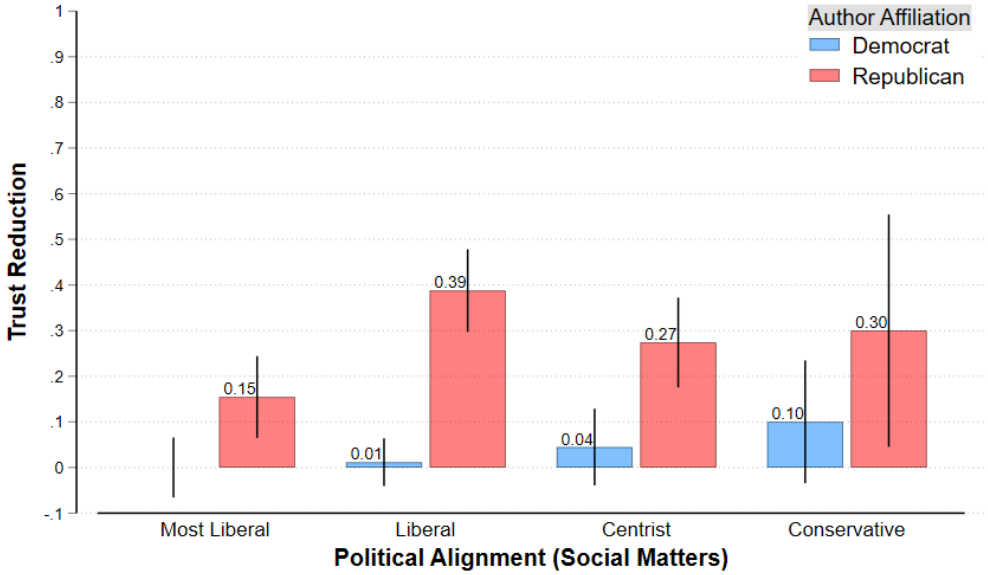


Figure 14: **Data Exemptions in the American Economic Review**

Figure 14 graphs the number of empirical papers published in the American Economic Review, papers that used any kind of data in their analysis (blue line). It also shows the fraction papers that use data and got exempted from sharing the data as a part of replication package. It is worth noting that American Economic Review stopped giving exemptions starting 2022. For the year 2022, we report the number of papers with restricted data instead. As per the new policy, starting 2022, papers with restricted data need to provide the path by which others can access the data – this might involve applying for the data or negotiating a data use agreement (DUA) with the provider.

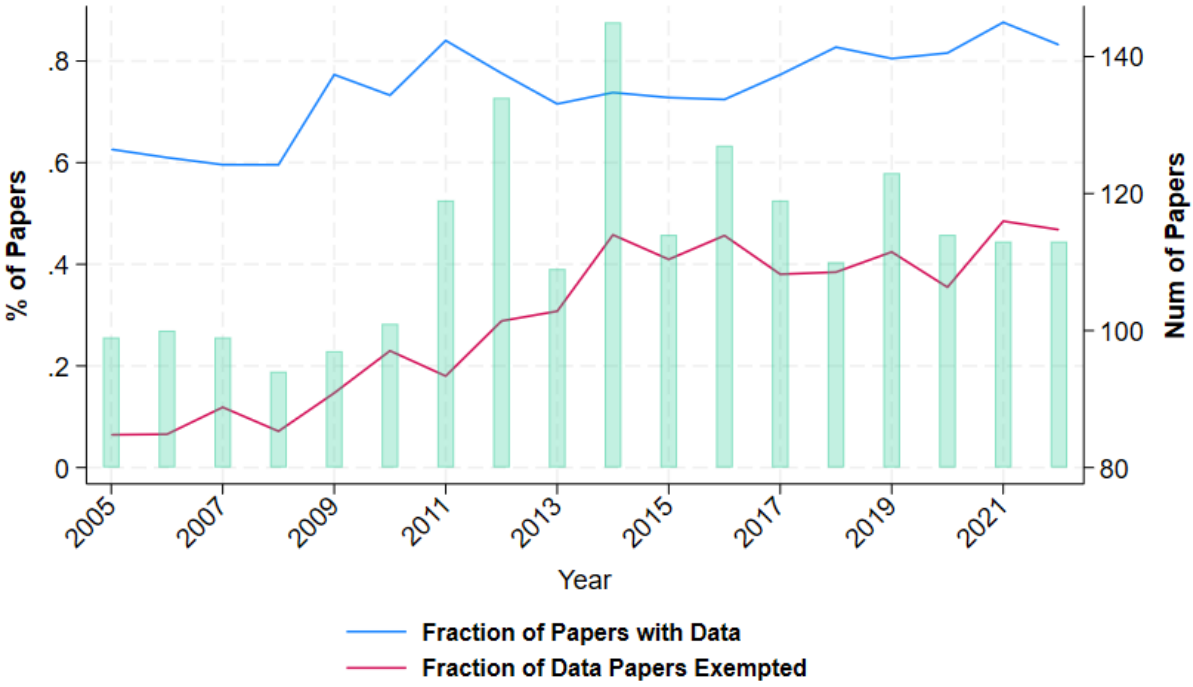


Figure 15: **CoI and Bias in Results: Medicine**

Figure 15 plots the odds ratios and confidence intervals of various studies from Table 11 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 11.

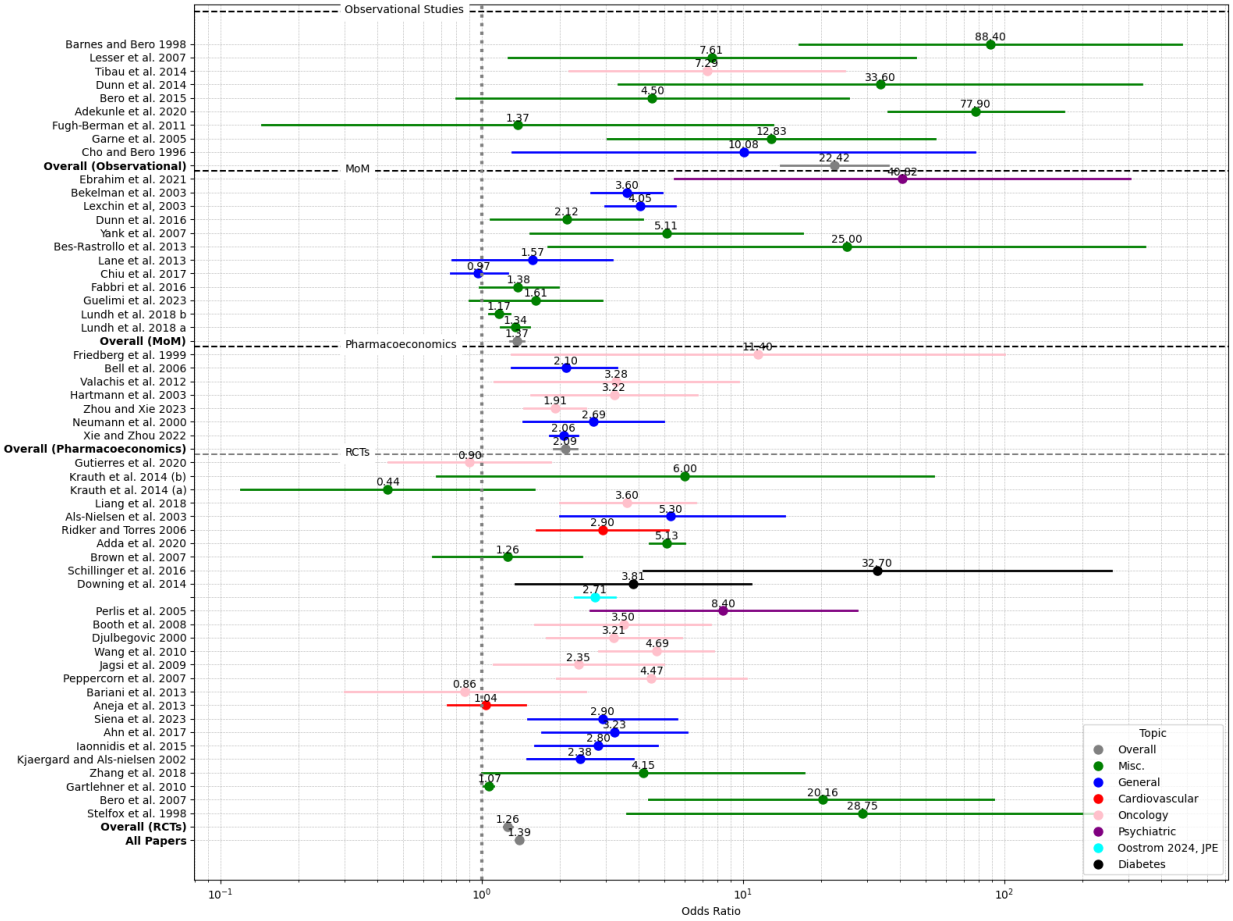
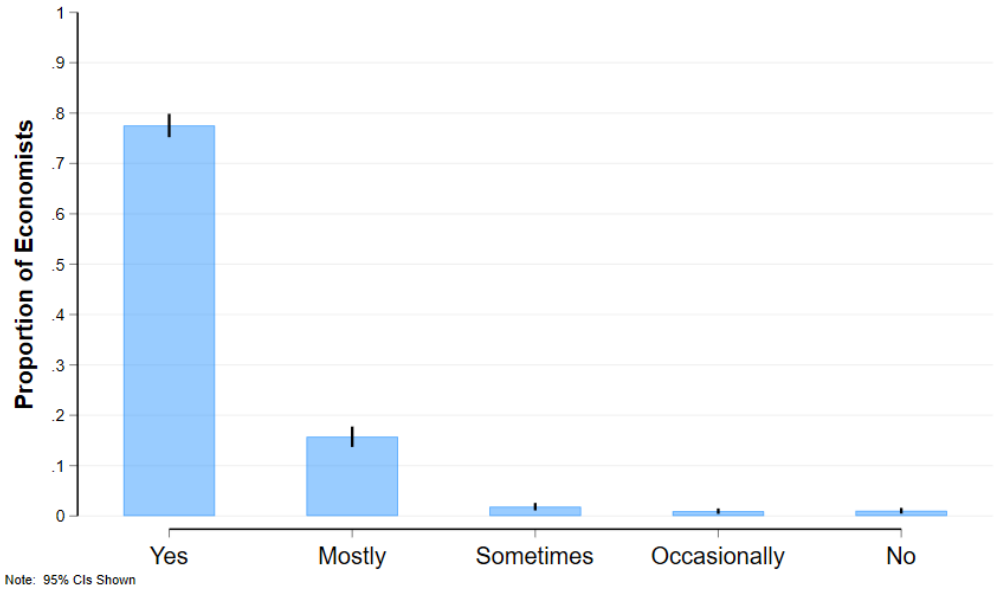


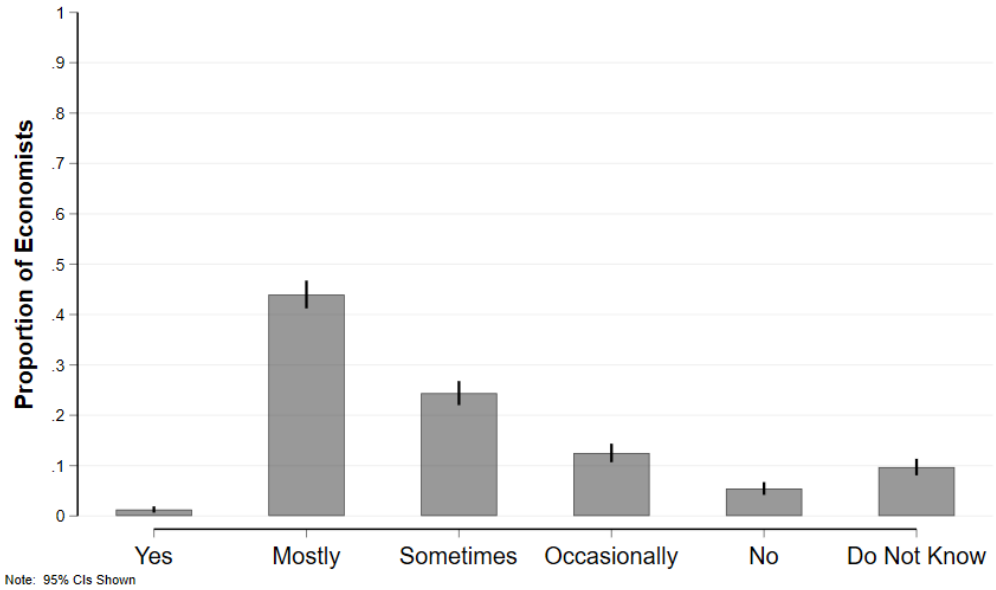
Figure 16: **Economists' Disclosure**

Figure 16 illustrates Economist's responses to our interrogative question about conflicts of interest 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, broadly divided into Science and Economics topics. We divide our economics vignettes into two categories – namely, Non-ideological and Ideological. The first column *Vignette* highlights the short-hand 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 the summary statistics of the interrogative questions asked 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 periods 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 conflicts of interest. 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 USD between 2000-2024). We classified potential conflicts of interest 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	0.134	0.166	0.158	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	0.214	0.000	0.459	0.171
Overall	0.185	0.454	0.388	0.117	0.574	0.287
N	135	416	551	103	61	164

Table 9: **CoI Discounts : AER, Empirical Papers' λs**

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, i.e., $1 - \frac{Value\ Conflicted}{Value\ Non-Conflicted}$. These values are estimated from equation (9), using two different estimates of the frequency of conflicted papers λ. $\lambda_{specific}$ denotes the conflict specific λ value taken from Table 8 (only AER papers). $\lambda_{overall}$ is the overall $\lambda = 0.388$ 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 λ, for Trading : Consulting fee, we use the *consulting fee* specific λ, for Supermarket vignettes we use the *consulting fee* specific λ, for CEO vignettes we use the *employer sources*, for rideshare we use the *discretionary gated data sources* specific λ and for Credit Card vignette we use the *editorial right to review* specific λ. We use the political appointment/affiliation specific λ for Tax and Abortion vignettes.

Vignette	*Trust Reduction*	*Specific*		*Overall*	
		λ	*CoI Discount*	λ	*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 v/s 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 column.

	AER		**RAND**	
	Theory	**Empirical**	**Theory**	**Empirical**
Overall Conflicted	95.200	208.571	56.556	43.200
Overall Non-Conflicted	145.806	185.628	26.617	41.269
p-value	0.401	0.402	0.192	0.889
N	135	416	103	61

Table 11: **CoI and Bias in Results: Literature Review in Medicine**

Table 11 summarizes the results of various studies on the association between Industry Sponsorship or Financial Ties and reported outcomes. OR (Odds Ratio) values are provided along with confidence intervals. Next, we show the topic and type of the study. We divide the type of the study into broadly three categories - RCTs (Randomized Clinical Trials), Ph-E (Pharmo-economics) ,MoM (Metanalyses of Metanalyses), ObS (Observational Studies). We also show a breakdown of the papers analyzed in the Data column. The values in the Data column are structured as follows - [A,B,C,D] where A is the number of industry-funded studies with statistically significant outcomes supporting the underlying drug; B is the number of industry-funded studies that do not report statistically significant results in favor of the underlying drug; C is the number of non-industry-funded studies that report statistically significant results in favor of the underlying drug and D is the number of non-industry-funded studies that do not report statistically significant results in favor of the underlying drug. *Agg. OR* in the Data column means that the paper reports an aggregated OR from other meta-analyses. The value *N/A* in the data column denotes papers that did not report the breakdown of studies they analyzed, and the authors do not have access to the data used in the paper anymore. *N/A*** denotes papers that did not report the breakdown of studies they analyzed, and the authors did not respond to our query about the data. † denotes the papers that did not report the breakdown of studies they analyzed, but we accessed these numbers from the replication package. We also show the source of the OR, whether it was reported in the body of the paper or calculated by us using data from the paper. Calculated OR Source means that we used the values in the Data column to calculate the OR and its confidence interval (more details for each paper in Appendix section X). Lastly, we report the type of CoI studied in each paper. Financial Ties: Analyzes financial relationships of authors with the drug manufacturer who funded the study, Industry Sponsorship: Analyzes if the study was funded by industry (independent of the Financial Ties of authors). Asterisk (*) denotes that the outcome of interest was the interpretation/conclusion of the results by the author.

	OR	**Topic**	**Type**	**Data**	**OR Source**	**Conflict**
Stelfox et al., 1998*	28.75 [3.6, 231.7]	Cardiovascular	RCTs	[23,1,20,25]	Calculated	Financial Ties
Bero et al., 2007	20.16 [4.37, 90.98]	Statins	RCTs	[54,41,39,58]	Text/Abstract	Industry Sponsorship
Gartlehner et al., 2010	1.07 [1.02, 1.11]	SSRIs	RCTs	[13,7,0,6]	Text/Abstract	Industry Sponsorship
Zhang et al., 2018*	4.15 [1.00, 17.11]	GDHT	RCTs	[16,3,18,14]	Text/Abstract	Industry Sponsorship, Financial Ties
Kjaergard and Als-Nielsen, 2002*	2.38 [1.49, 3.80]	General	RCTs	N/A	Calculated	Industry Sponsorship, Financial Ties
Flacco et al., 2015	2.80 [1.6, 4.7]	General	RCTs	[125,57,111,28]	Text/Abstract	Industry Sponsorship
Ahn et al., 2017	3.23 [1.7, 6.1]	General	RCTs	[98,36,38,23]	Text/Abstract	Financial Ties
Siena et al., 2023*	2.90 [1.50, 5.60]	General	RCTs	[279,24,85,21]	Text/Abstract	Industry Sponsorship, Financial Ties
Aneja et al., 2013	1.04 [0.74, 1.47]	Cardiovascular	RCTs	[169,181,130,239]	Text/Abstract	Industry Sponsorship, Financial Ties
Bariani et al., 2013*	0.86 [0.30, 2.50]	Oncology	RCTs	[71,32,24,9]	Text/Abstract	Industry Sponsorship, Financial Ties
Peppercorn et al., 2007*	4.47 [1.94, 10.30]	Oncology	RCTs	[86,17,56,38]	Calculated	Industry Sponsorship
Jagsi et al., 2009*	2.35 [1.11, 4.96]	Oncology	RCTs	[48,19,29,27]	Calculated	Industry Sponsorship, Financial Ties
Wang et al., 2010*	4.69 [2.82, 7.72]	Oncology	RCTs	[24,11,3,52]	Text/Abstract	Financial Ties
Djulbegovic et al., 2000*	3.21 [1.77, 5.81]	Oncology	RCTs	[74,26,47,53]	Calculated	Industry Sponsorship
Booth et al., 2008*	3.50 [1.60, 7.50]	Oncology	RCTs	N/A	Text/Abstract	Industry Sponsorship

Continued on next page

Table 11 – continued from previous page

	OR	Topic	Type	Data	OR Source	Conflict
Perlis et al., 2005	8.40 [2.60, 27.30]	Psychiatric	RCTs	[93,20,37,12]	Text/Abstract	Industry Sponsorship, Financial Ties
Oostrom, 2024	2.71 [2.27, 3.24]	Psychiatric	RCTs	[526,657,262,887]†	Calculated	Industry Sponsorship
Downing et al., 2014*	3.81 [1.35, 10.73]	Diabetes	RCTs	[23,62,5,86]	Calculated	Financial Ties
Schillinger et al., 2016*	32.70 [4.15, 257.27]	Diabetes	RCTs	[25,26,1,34]	Calculated	Industry Sponsorship
Brown et al., 2006*	1.26 [0.65, 2.42]	Gastroenterology	RCTs	[38,27,83,36]	Calculated	Industry Sponsorship
Adda et al., 2020	5.13 [4.40, 5.98]	General	RCTs	[4194,1763,272,586]†	Calculated	Industry Sponsorship
Ridker and Torres, 2006	2.90 [1.62, 5.20]	Cardiovascular	RCTs	[92,45,51,53]	Calculated	Industry Sponsorship
Als-Nielsen et al., 2003*	5.3 [2.00, 14.40]	General	RCTs	[75,72,11,56]	Text/Abstract	Industry Sponsorship
Liang et al., 2018*	3.6 [2,6.6]	Oncology	RCTs		Text/Abstract	Industry Sponsorship
Krauth et al., 2014	0.44 [0.12,1.60]	Statins	RCTs	[9,8,18,7]	Calculated	Industry Sponsorship
Krauth et al., 2014*	6 [0.67,53.5]	Statins	RCTs	[18,1,21,7]	Calculated	Industry Sponsorship
de Souza Gutierres et al., 2020	0.90 [0.44,1.83]	Oncology	RCTs	[67,85,31,37]	Table 1	Industry Sponsorship
Xie and Zhou, 2022	2.06 [1.82, 2.33]	General	Ph-E		Text/Abstract	Industry Sponsorship
Neumann et al., 2000	2.70 [1.45,5.00]	General	Ph-E	[16,61,51,519]	Calculated	Industry Sponsorship
Zhou and Xie, 2023	1.91 [1.45, 2.51]	Oncology	Ph-E	[183,204,476,674]	Text/Abstract	Industry Sponsorship
Hartmann et al., 2003*	3.22 [1.55, 6.70]	Oncology	Ph-E	[27,17,35,71]	Calculated	Industry Sponsorship
Valachis et al., 2012*	3.28 [1.12, 9.65]	Oncology	Ph-E	[32,13,9,12]	Calculated	Industry Sponsorship, Financial Ties
Bell et al., 2006	2.1 [1.3, 3.3]	General	Ph-E	N/A**	Text/Abstract	Industry Sponsorship
Friedberg et al., 1999*	11.40 [1.3, 100.25]	Oncology	Ph-E	[19,1,15,9]	Calculated	Industry Sponsorship
Lundh et al., 2018	1.34 [1.19,1.52]	General	MoM	[1198,1776,576,1147]	Calculated	Industry Sponsorship

Continued on next page

Table 11 – continued from previous page

	OR	Topic	Type	Data	OR Source	Conflict
Lundh et al., 2018*	1.17 [1.07,1.28]	General	MoM	[1601,2127,1581,2456]	Calculated	Industry Sponsorship
Guelimi et al., 2024*	1.61 [0.90,2.89]	Inflammatory Diseases	MoM	[39,28,69,80]	Text/Abstract	Industry Sponsorship
Chartres et al., 2016*	1.38 [0.98,1.96]	Nutrition	MoM	[140,190,80,150]	Calculated	Industry Sponsorship
Chiu et al., 2017*	0.97 [0.76,1.26]	General	MoM	[137,486,180,624]	Calculated	Industry Sponsorship
Yank et al., 2007*	5.11 [1.54,16.92]	Hypertension	MoM		Text/Abstract	Industry Sponsorship
Lane et al., 2013*	1.57 [0.77, 3.17]	General	MoM	[37,26,30,33]†	Text/Abstract	Industry Sponsorship
Dunn et al., 2016	2.12 [1.08, 4.14]	Neuraminidase Inhibitors	MoM	[54,17,63,42]	Calculated	Financial Ties
Lexchin et al., 2003	4.05 [2.98, 5.50]	General	MoM	Agg. OR	Text/Abstract	Industry Sponsorship
Bes-Rastrollo et al., 2013*	25.00 [1.8, 346.71]	Obesity	MoM	[5,1,2,10]	Calculated	Industry Sponsorship
Bekelman et al., 2003*	3.60 [2.63, 4.91]	General	MoM	Agg. OR	Calculated	Industry Sponsorship, Financial Ties
Cho and Bero, 1996	10.07 [1.31,77.30]	General	Ob.S.	[39,1,89,23]	Calculated	Industry Sponsorship
Garne et al., 2005	12.83 [3.03,54.42]	Tobacco	Ob.S.	[28,6,4,11]	Calculated	Industry Sponsorship, Financial Ties
Fugh-Berman et al., 2011	1.37 [0.144,13]	Menopause	Ob.S.	[6,35,1,8]	Calculated	Financial Ties
Adekunle et al., 2020*	77.90 [36.0, 168.57]	Indoor Tanning	Ob.S.	[39,11,27,593]	Calculated	Industry Sponsorship
Ebrahim et al., 2016*	40.82 [5.48, 304.18]	Psychiatric	Ob.S.	[53,1,74,57]	Calculated	Industry Sponsorship , Financial Ties
Bero et al., 2016	4.50 [0.8, 25.35]	Animal Toxicology	Ob.S.	[9,2,12,12]	Calculated	Industry Sponsorship
Dunn et al., 2014*	33.60 [3.35, 337.23]	Neuraminidase Inhibitors	Ob.S.	[7,1,5,19]	Calculated	Industry Sponsorship , Financial Ties
Tibau et al., 2015*	7.29 [2.17, 24.49]	Oncology	Ob.S.	N/A	Text/Abstract	Industry Sponsorship, Financial Ties
Lesser et al., 2007*	7.61 [1.27, 45.73]	Nutrition	Ob.S.	[14,8,24,28]	Text/Abstract	Industry Sponsorship
Barnes and Bero, 1998*	88.40 [16.40, 476.50]	Smoking	Ob.S.	[29,2,10,65]	Text/Abstract	Financial Ties

Table 12: **GPT-4 Omni CoI Trust Reduction by Treatment**

Table 12 exhibits the average CoI Trust Reduction by GPT-4 Omni vis-a-vis our survey respondents, with the number of observations in the parenthesis. First column denotes the average CoI Trust Reduction in our GPT simulations. Note that we run 1000 simulations per randomization. Second column shows the average CoI Trust Reduction of our survey respondents. In subsequent columns three, four and five we break down our survey respondents into the three broad categories of Average Americans, Ordinary Economists and Selected Economists.

	GPT-4 Omni	Full Sample	Average Americans	Ordinary Economists	Selected Economists
Industrial Disaster	0.50	0.51	0.52	0.51	0.49
	(1000)	(2515)	(1280)	(771)	(464)
Alzheimer Drug : Consulting Fee	0.21	0.36	0.36	0.37	0.35
	(1000)	(1267)	(650)	(385)	(232)
Alzheimer Drug : Reasearch Funding	0.20	0.35	0.34	0.38	0.34
	(1000)	(1248)	(630)	(386)	(232)
Trading : Consulting Fee	0.42	0.40	0.45	0.39	0.29
	(3000)	(1257)	(639)	(384)	(234)
Trading : Research Grant	0.40	0.34	0.38	0.33	0.28
	(3000)	(1258)	(641)	(385)	(232)
CEO : Seeking Position	0.50	0.42	0.45	0.42	0.36
	(1000)	(836)	(419)	(265)	(152)
CEO : Immediately Nominated	0.50	0.36	0.42	0.33	0.28
	(1000)	(848)	(433)	(253)	(162)
CEO : Nominated 2 Years Post	0.49	0.30	0.38	0.25	0.16
	(1000)	(832)	(428)	(254)	(150)
Credit Card : Right to Review Enforced	0.50	0.52	0.46	0.58	0.55
	(1000)	(1251)	(631)	(389)	(231)
Credit Card : Right to Review Waived	-0.20	-0.12	-0.03	-0.23	-0.19
	(1000)	(1261)	(649)	(380)	(232)
Supermarket : Private Expert Witness, 400K	0.32	0.36	0.37	0.36	0.34
	(2000)	(617)	(309)	(193)	(115)
Supermarket : Public Expert Witness, 400K	0.04	0.18	0.25	0.11	0.09
	(2000)	(649)	(335)	(199)	(115)
Supermarket : Private Expert Witness, No Mention of Compensation	0.21	0.20	0.19	0.20	0.21
	(2000)	(610)	(310)	(185)	(115)
Supermarket : Public Expert Witness, No Mention of Compensation	0.04	0.11	0.16	0.06	0.04
	(2000)	(638)	(326)	(194)	(118)
Abortion : Positive Impact, Democrat Author	0.20	0.19	0.27	0.12	0.09
	(1000)	(617)	(319)	(183)	(115)
Abortion : Positive Impact, Republican Author	0.21	0.03	0.11	-0.05	-0.07
	(1000)	(623)	(308)	(195)	(120)
Abortion : No Impact, Democrat Author	0.20	0.13	0.20	0.04	0.04
	(1000)	(639)	(332)	(192)	(115)
Abortion : No Impact, Republican Author	0.23	0.32	0.33	0.31	0.33
	(1000)	(636)	(321)	(200)	(115)
Pro-Tax Book : Pro-Tax Finding	0.20	0.12	0.19	0.06	0.05
	(1000)	(1271)	(639)	(394)	(238)
Pro-Tax Book : Anti-Tax Finding	-0.16	0.08	0.19	-0.03	-0.05
	(1000)	(1247)	(641)	(378)	(228)
Rideshare : Private Data (Benefits greater than Costs)	0.43	0.27	0.32	0.23	0.20
	(1000)	(618)	(314)	(190)	(114)
Rideshare : Private Data (Benefits less than Costs)	0.49	0.13	0.21	0.05	0.06
	(1000)	(627)	(323)	(189)	(115)
Rideshare : Public Data (Benefits greater than Costs)	0.07	-0.05	0.01	-0.11	-0.14
	(1000)	(634)	(321)	(194)	(119)
Rideshare : Public Data (Benefits less than Costs)	0.15	-0.03	0.01	-0.06	-0.06
	(1000)	(636)	(322)	(198)	(116)
Trading : 10k Compensation	0.31	0.34	0.39	0.30	0.24
	(2000)	(842)	(439)	(254)	(149)
Trading : 100k Compensation	0.42	0.37	0.40	0.38	0.30
	(2000)	(824)	(418)	(251)	(155)
Trading : 1M Compensation	0.50	0.41	0.46	0.40	0.31
	(2000)	(849)	(423)	(264)	(162)

Appendix

A Survey Sample

A.1 Survey Instrument

Figure A1 is a QR code which acts as a 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